sustainability

MDPI

Review

Metaheuristic Optimization and Artificial Intelligence in Wind Energy Systems for Sustainable Environment: Forecasting, Control, Design, and Diagnostics—A Critical Review

Ahmed M. Nassef [1,*], Walid Merrouche [2] and Mostafa A. Elhosseini [3]

[1] Department of Electrical Engineering, College of Engineering in Wadi Alddawasir, Prince Sattam bin Abdulaziz University, Wadi Alddawasir 11991, Saudi Arabia
[2] Center for Renewable Energies Development, CDER, Algiers 16340, Algeria; w.merrouche@cder.dz
[3] Computers and Control Systems Engineering Department, Faculty of Engineering, Mansoura University, Mansoura 35516, Egypt; melhosseini@mans.edu.eg
* Correspondence: a.nasef@psau.edu.sa

Abstract

Wind power is the largest renewable energy source used at the sustainable utility scale. However, wind turbine design, siting, forecasting, control, and maintenance face real challenges, mainly because wind speed is variable and poorly predictable. Over the last ten years, researchers have applied both artificial intelligence (AI) models and metaheuristic optimization algorithms across all phases of the wind energy value chain. This review examines these two methodological approaches across four application fields, with particular attention to evaluation practice, validation level, and implications for sustainability. First, we review wind speed and power forecasting, from statistical baselines to machine learning and deep learning architectures, including models that rely on metaheuristics to tune their hyperparameters and decomposition stages. Second, we examine control applications, focusing on metaheuristic tuning of generator and converter controllers (mainly rotor-side control of doubly-fed induction generators), pitch and torque control, and maximum power point tracking (MPPT). Third, we discuss design-level optimization, covering blade aerodynamic design and wind farm layout under wake effects. Fourth, we examine machine-learning-based supervisory control and data acquisition (SCADA) monitoring and fault diagnosis, including normal-behavior modelling and anomaly detection. Finally, as floating offshore wind turbines (FOWTs) are a cross-cutting emerging field, with coupled hydro-aero-elastic problems in control co-design and AI-based monitoring, the paper proposes a research agenda for the next sustainable generation of onshore and offshore wind systems, focused on standardized open benchmarks, uncertainty-aware forecasting, physics-informed learning, and hardware-validated optimization.

Keywords: sustainable wind energy; metaheuristic optimization; artificial intelligence; wind power forecasting; MPPT; wind farm layout; condition monitoring; SCADA

Check for updates

Academic Editors: Quetzalcoatl Hernández-Escobedo and Jesus Alejandro Franco

Received: 11 September 2026
Revised: 18 October 2026
Accepted: 21 October 2026
Published: 22 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

As global installed capacity is expanding, wind power has emerged as a key component of decarbonization initiatives [1]. However, several critical engineering challenges remain in wind energy development. The resource is sporadic and partially predictable. Turbines are large, flexible, nonlinear electromechanical systems that operate under turbulent inflow conditions. Wind farms couple tens of machines through wake interactions. Moreover,

operations and maintenance (O&M) account for about 25% to 30% of the levelized cost of energy, especially for offshore installations [2,3].

These hurdles can be treated as optimization problems, learning problems, or both. The layout of wind farms is a combinatorial optimization problem subject to weak constraints [4]. Controller tuning is a multi-objective optimization problem. It targets objectives such as maximizing energy capture and minimizing structural fatigue loads [5]. Forecasting is a supervised learning problem that uses non-stationary time series [6]. Fault diagnosis requires a supervised approach that detects anomalies based on high-dimensional supervisory control and data acquisition (SCADA) streams [7]. Consequently, the two predominant methodological trends in computational intelligence, metaheuristic optimization and machine/deep learning, have been extensively applied in the wind energy literature.

These applications support sustainability through different pathways: forecasting informs renewable integration; control balances energy capture and equipment stress; design addresses production and resource requirements; and monitoring supports timely maintenance. Reported indicators include forecast quality, energy yield, fatigue loads, blade mass, and diagnostic lead time. However, improvements in these indicators do not automatically establish lower life-cycle environmental impacts. This review therefore examines both the potential sustainability benefits and the evidence supporting their practical realization.

Optimization is a branch of computer science and mathematics. It aims to identify the optimal solution to a given optimization problem by minimizing or maximizing an objective's function through dependent optimization variables. Classical mathematical programming includes linear programming [8], solved by simplex or interior-point methods, nonlinear programming [9], typically solved with gradient-based algorithms, and dynamic programming (DP) [10], which decomposes sequential decisions recursively and does not rely on gradients. These methods are efficient, and they can offer optimality guarantees when the problem structure suits them (linearity or convexity for LP and convex NLP, a decomposable stage structure for DP). However, gradient-based local solvers can be trapped in local optima on nonconvex problems, and all three require an explicit model with the corresponding structure, so they do not apply directly to black-box objectives. On the contrary, metaheuristic optimization algorithms (MOAs), such as genetic algorithms (GAs) [11], particle swarm optimization (PSO) [12], differential evolution (DE) [13], grey wolf optimizer (GWO) [14], and whale optimization algorithm (WOA) [15], do not require gradients, convexity, or a decomposable structure; their derivative-free, stochastic search makes them applicable to nonconvex and black-box problems and helps them escape local optima, at the cost of many function evaluations and no optimality guarantee [16,17]. Despite improvements, the stochastic nature of MOAs and their random initialization make it difficult for researchers to validate MOAs' performance. Hence, this raises concerns about the statistical significance of published findings. Due to their excellent performance in solving complex problems, AI models, starting from support vector machines (SVMs) [18] and random forests (RFs) [19] to long short-term memory (LSTM) networks [20], temporal convolutional networks (TCN) [21], and Transformer networks [22], are increasingly being used in engineering applications like the wind energy field. These tools offer the approximation required for forecasting, surrogate modeling, and diagnostics [6,23]. Yet, these models still depend on data quality and their interpretability. MOAs are increasingly integrated with AI models: learnt surrogates speed up metaheuristic design loops [4], and metaheuristics serve to adjust the decomposition phases and hyperparameters of AI models [24].

Consequently, methods can only be compared meaningfully on comparable problems, against baselines of comparable strength, with a reported computational effort and tuning

budget, and under the same validation conditions; Section 8.5 applies these criteria to the studies reviewed here.

Existing wind-energy-oriented studies typically concentrate on a single algorithm class or application, such as MPPT control techniques [25], forecasting [6,26], turbine design optimization [5], layout optimization [4], or monitoring [7,27].

Previous surveys have mapped artificial intelligence (AI) applications across the wind energy sector [28] and explored machine learning in offshore wind systems [29]. These works primarily catalogue methods as well as the performance metrics reported by their respective authors. In contrast, this review differs in two key aspects. First, it considers metaheuristics and AI together, recognizing them as complementary and, at times, competing approaches to similar problems. Second, it systematically examines the evaluation practices underlying reported performance gains, including data and benchmarks, baselines, validation protocols, statistical testing, computational budgets, reproducibility, and validation levels. The same evaluation criteria are applied consistently across all application domains (see evaluation caveats in Sections 4–6; cross-domain comparison in Section 8.5).

The closest recent reviews each cover part of this scope. Rezk et al. [5] and Zhang et al. [2] review optimization methods for wind turbines and offshore wind systems (design, layout, and control), without covering AI-based forecasting or condition monitoring. Moshtaghi et al. [28] review AI applications across wind power generation, while Wan et al. [29] provide a systematic review of machine learning in offshore wind for 2022–2025; yet, both focus on AI rather than metaheuristic optimization. Sharma et al. [16] and Manzano et al. [30] limit their reviews to forecasting, whereas Pandit et al. [7] and Sheng et al. [27] concentrate on SCADA-based condition monitoring. To our knowledge, none of these reviews treats metaheuristics and AI together across forecasting, control, design, and monitoring, or applies common evaluation criteria across these domains; these two elements, together with the problem–decision–method framework of Section 8.6, constitute the added value of the present review.

This review addresses three questions: (RQ1) In which wind-energy tasks do metaheuristics and AI models provide measurable gains over conventional methods? (RQ2) How rigorously are these gains evaluated, in terms of data, baselines, validation protocol, and validation level? (RQ3) What methodological gaps limit their transfer to field operation and the demonstration of sustainability benefits? The scope covers four application fields (wind speed and power forecasting; turbine and farm control; layout and blade design; and SCADA-based condition monitoring and fault diagnosis), with floating offshore wind as a cross-cutting topic. Excluded are studies on grid integration and power-system operation (e.g., unit commitment and electricity markets), vibration-, acoustic-, or image-based inspection without SCADA data, purely CFD or structural analyses without an optimization or learning component, and conceptual proposals without quantitative evaluation.

To close this gap, this review paper makes the following contributions:

- A tabulated summary of representative quantitative results and a unified taxonomy of AI and metaheuristic techniques translated onto the four primary application fields of wind energy engineering.
- A cross-domain analysis of evaluation practices across forecasting, control, design, and diagnostics using consistent criteria, including datasets and benchmarks, baseline models, validation protocols, statistical testing, computational budgets, reproducibility, and validation level.
- A research agenda emphasizing open benchmarks, uncertainty quantification, physics-informed learning, and the specific challenges of floating offshore wind.

This review adopted a narrative methodology. Studies were selected based on relevance, recency, and methodological significance; exhaustive coverage is not claimed. The

studies were identified through search queries in Scopus, Web of Science, IEEE Xplore, and Google Scholar. Search terms combined methodological keywords (metaheuristic, evolutionary, swarm, machine learning, deep learning, reinforcement learning, physics-informed) with application-specific terms (wind forecasting, MPPT, pitch control, wind farm layout, blade design, SCADA, fault diagnosis, floating offshore). The review focuses on English-language articles, reviews, and conference papers published between 2016 and 2026, as well as foundational works before 2016. The included studies addressed one of four application domains (forecasting, control, design, or monitoring) or floating offshore wind, and they reported quantitative results on an identified dataset, benchmark turbine, or experimental setup. Conceptual proposals without evaluation were excluded. Because the search results were not systematically screened, the total number of records was not reported, and neither a PRISMA flow diagram nor a meta-analysis was provided. Some sub-topics may therefore be under-represented. These limitations are discussed in Section 8. The final corpus comprises 122 references. Sections 3–7 cite approximately 24 references on forecasting, 11 on turbine and farm control, 14 on layout and blade design, 13 on SCADA-based monitoring, and 13 on floating offshore wind; some references are cited in more than one section, and the remainder cover methodological background and related reviews. Table 1 summarizes the review scope by application domain.

Table 1. Scope of the review by application domain: period of emphasis, methodological families covered, and the level of experimental evidence in the corresponding literature.

Application Domain	Period Emphasized	Dominant Method Families Reviewed	Evidence Level Reached in Literature
Forecasting	2017–2026	Statistical baselines, classical ML, deep learning (LSTM/TCN/Transformer/GNN), hybrid decomposition pipelines, probabilistic methods	Mostly single-site studies; few shared public benchmarks; probabilistic evaluation emerging but not standard
Turbine control—MPPT, pitch, generator	2018–2026	Swarm/evolutionary gain tuning, fuzzy and fractional-order controllers, MPC weight design, deep RL	Predominantly simulation; laboratory-scale validation in a few studies
Farm-level control—wake steering	2020–2026	Deep reinforcement learning, optimization-based baselines	Dynamic simulation; no field deployment
Layout and blade design	2005–2026	GA/PSO/CMA-ES layout search, NSGA-II multi-objective blade design, ML surrogates	Benchmark turbines (NREL 5 MW) and standard wake models; results optimizer- and wake-model-dependent
SCADA diagnostics and monitoring	2016–2026	Normal behavior models, deep anomaly detection, end-to-end deep diagnosis, transfer learning	Public multi-year SCADA datasets with event logs; cross-farm generalization unresolved
Floating offshore wind	2018–2026	Metaheuristic controller tuning, wave-feedforward control	Predominantly aero-hydro-servo-elastic simulation; one experimental validation [31–34]

Section 2 provides the methodological background. Sections 3–6 review forecasting, control, design, and diagnostics, respectively. Section 7 addresses floating offshore wind. Section 8 develops the critical discussion, Section 9 the future directions, and Section 10 concludes.

2. Methodological Background

Metaheuristics and AI models are widely used to design and optimize wind energy applications. Each one is divided into several classes as illustrated in Figure 1.

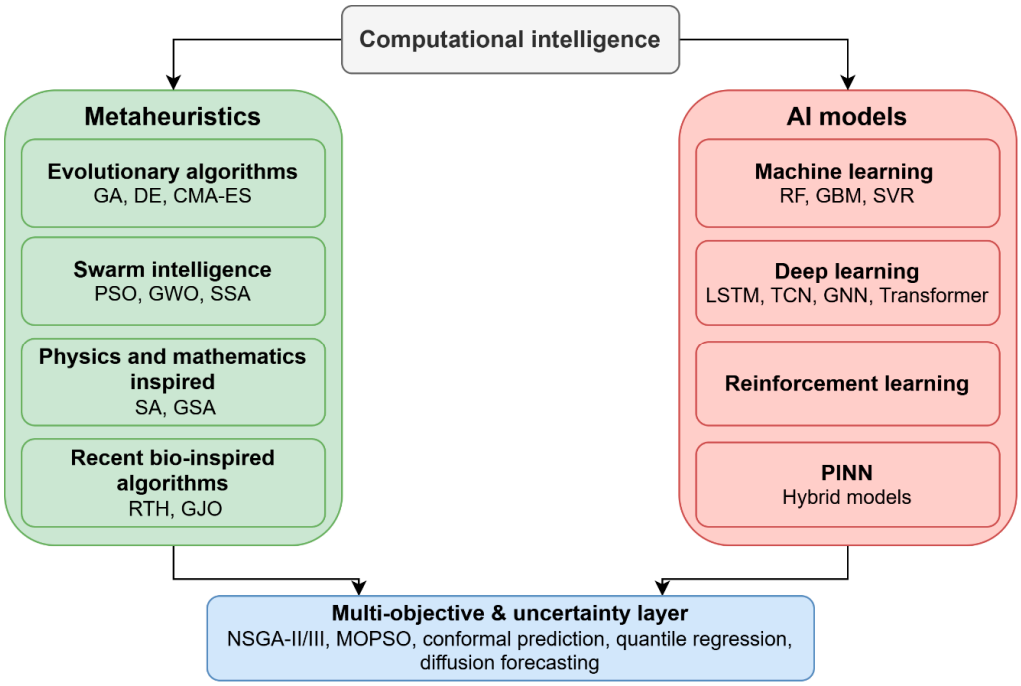


Figure 1. Twin taxonomy of metaheuristic algorithms and AI model families used in wind energy.

2.1. Metaheuristic Optimization

Wind energy optimization problems, such as layout placement, controller tuning, and blade shape design, share a common structure. They all have a cost function, $J(x)$, which is often a nonlinear and non-convex function that accepts the optimization variables as an input vector (x). It is also frequently non-differentiable or expensive to evaluate [35]. The cost function is defined over a search space that mixes continuous, discrete, and combinatorial variables [36]. Gradient-based methods have difficulty with this combination: non-differentiability rules out analytic gradients, and multimodality can trap local descent in suboptimal basins [37]. To overcome these limitations, MOAs use a group of candidate solutions (population, referred to as optimization variables, x_1,..., x_n) for gradient information. These solutions are updated at each iteration based on a mechanism that balances exploration (wide sampling to prevent premature convergence) and exploitation (search concentration around promising regions) to converge toward a global or near-global solution [38].

2.1.1. Evolutionary Algorithms

Genetic algorithms (GAs) use selection, crossover, and mutation operators to evolve populations while encoding potential solutions as chromosomes. Compared with standard grid configurations, Grady et al.'s research showed that a GA-optimized wind farm architecture could significantly reduce wake losses. As a result, GAs became the historical standard for layout issues [39]. Each candidate is perturbed by differential evolution (DE) using the scaled vector difference between other population members. As a result, it performs well with fewer control parameters than GA for continuous, high-dimensional problems such as blade design variables, blade twist, and chord distributions. A full covariance matrix is used for the multivariate Gaussian sampling distribution at each generation in CMA-ES, the most advanced member of this family for wind applications [40,41]. Without gradient information, this method successfully determines the local curvature and correlation structure of the cost landscape. For ill-conditioned, non-separable continuous problems, CMA-ES is regarded as one of the most effective derivative-free techniques. CMA-ES is a

common reference technique in comparative layout research and has been directly benchmarked against GAs, particle swarm algorithms, and random searches for wind farm layout optimization [42,43].

2.1.2. Swarm Intelligence

These techniques simulate group hunting or foraging. Particle swarm optimization (PSO) is the most often used controller-tuning methodology in wind engineering. PSO is simple to implement, converges quickly on low- to medium-dimensional problems, and can be easily hybridized with fuzzy logic controllers [44]. In PSO, each particle updates its velocity based on its own best-known position and the swarm's global best. Ant colony optimization (ACO) and artificial bee colony (ABC) draw on pheromone-trail and foraging-recruitment metaphors [45,46], respectively, and are used mainly for combinatorial variants of layout and cable-routing problems. The grey wolf optimizer (GWO) simulates the leadership hierarchy of grey wolf packs (alpha, beta, delta, and omega wolves) and their encircling hunting mechanism. The whale optimization algorithm (WOA) mimics the bubble-net hunting strategy used by humpback whales. Both the GWO and the WOA have become common baselines in DFIG/PMSG controller-gain tuning studies because their exploration–exploitation balance is controlled by a single shrinking coefficient, which simplifies convergence analysis [47–49]. Salp swarm algorithm (SSA) [50] and its multi-objective variant (MSSA) [50] have been applied to MPPT and pitch/fault-ride-through controller tuning [51], where the chain-like swarm topology and leader–follower position updates propagate sequentially down the salp chain, suited to problems with sequential constraint handling.

2.1.3. Physics and Mathematics-Inspired Algorithms

This cluster of algorithms encodes physical or mathematical processes rather than biological behavior. For example, simulated annealing (SA) mimics the slow cooling of a metal lattice by accepting moves that worsen the solution with a temperature-dependent probability to escape local minima [52]. The gravitational search algorithm (GSA) treats candidate solutions as masses that attract one another according to Newtonian gravity. Heavier masses exert a stronger pull than lighter ones [53]. The sine–cosine algorithm (SCA) and the arithmetic optimization algorithm (AOA) use trigonometric or basic arithmetic operators to update positions with a controllable exploration/exploitation ratio [54,55]. While these methods are not commonly used as primary optimizers in wind applications, they are common components of hybrid algorithms, where they refine solutions found by a swarm or evolutionary search.

2.1.4. Recent Bio-Inspired Algorithms

Bio-inspired metaheuristics applied to wind energy problems encompass both well-established and recent algorithms. Examples of older algorithms still widely used today include the crow search algorithm (2016) [56], built on memory-based food-hiding behavior, and the manta-ray foraging optimization algorithm (2020). More recent additions include the golden eagle optimizer (spiral attack trajectories) [57], the greater cane rat algorithm (applied to wind turbine MPPT [58]), the energy valley optimizer (applied to DFIG adaptive fuzzy control [59]), and the red-tailed hawk algorithm [17]. Several of these more recent additions were introduced with an energy-system case study as an early validation benchmark. This proliferation raises a methodological caveat addressed in Section 8.1: many of these algorithms differ from established methods primarily in metaphor rather than mechanism, and their claimed superiority should be considered in light of the no-free-lunch

theorem. A comparative summary of metaheuristic families that are applied in wind energy engineering is illustrated in Table 2.

Table 2. Comparative summary of metaheuristic families applied in wind energy engineering.

Family	Mechanism	Control Parameters	Representative Wind Application	Key Limitation
Evolutionary (GA, DE, CMA-ES)	Selection, crossover, mutation (GA/DE); adapted covariance-matrix Gaussian sampling (CMA-ES)	Population size, crossover/mutation rates; CMA-ES step size	GA-optimized wind farm layout [39]; CMA-ES benchmarked against layout baselines [42,43]	Costly when each evaluation requires a wake or aeroelastic simulation
Swarm (PSO, ACO, ABC, GWO, WOA, SSA)	Position/velocity updates guided by personal- and swarm-best (PSO); pheromone/foraging metaphors (ACO/ABC); leader-hierarchy encircling (GWO/WOA)	Swarm size, inertia/shrinking coefficient, topology	PSO/GWO/WOA-tuned DFIG and PMSG controller gains [47,49]; SSA-tuned MPPT and fault ride-through [51]	No single algorithm dominates across studies; rankings invert by dataset (no-free-lunch)
Physics/math-inspired (SA, GSA, SCA, AOA)	Simulated cooling (SA); Newtonian mass attraction (GSA); trigonometric/arithmetic position updates (SCA/AOA)	Cooling schedule (SA); gravitational constant decay (GSA)	Rarely primary optimizers in wind applications; used as refinement stage in hybrids	Limited standalone use; mainly a local-search complement to swarm/evolutionary search
Recent bio-inspired (cuckoo search, golden eagle, greater cane rat, energy valley, RTH)	Varied nature metaphors (memory-based caching, spiral attack, foraging, valley energy dynamics, hunting)	Algorithm-specific (e.g., awareness probability, attack angle)	MPPT via greater cane rat [58]; adaptive fuzzy DFIG control via energy valley [59]	Frequently validated on a single case study at introduction; proliferation risks metaphor-over-mechanism

2.2. AI Models

2.2.1. Machine Learning

Because of its interpretability and data efficiency, traditional machine learning is still an efficient tool for tabular, feature-engineered SCADA data and short forecasting horizons. Support vector regression (SVR) fits a maximum-margin function inside an ε-insensitive tube using kernel functions to handle nonlinearity. This method has previously provided a strong basis for wind power curve modeling [60]. Using uncorrelated decision trees, random forests generate accurate hour-ahead wind-power forecasts with integrated feature-priority ranking. This feature helps identify the SCADA channels that actually improve predictive capacity [61]. By iteratively fitting trees to residuals, gradient boosting (XGBoost and LightGBM) often outperforms deep learning on tabular data benchmarks when sample numbers are moderate. This pattern is also seen in comparative studies of SCADA fault diagnosis [62]. Gaussian process regression's cubic-in-sample-size scaling limits its use to short time series or sparse feature sets, even though it offers calibrated prediction uncertainty estimates [26].

2.2.2. Deep Learning

These architectures are more useful when features are impractical because of sequence length, spatial organization, or raw signal resolution. Multilayer perceptrons (MLPs) provide a simple, nonlinear baseline deep learning model. Convolutional networks extract local temporal or spatial patterns from wind-speed sequences or turbine-array configurations. LSTM and GRU recurrent networks yield longer-range temporal correlations in wind speed and in power time series. These networks are the most often utilized backbone for forecasting and SCADA-based fault categorization [63]. Temporal convolutional networks (TCNs) offer a parallelizable, dilated-convolution alternative to recurrence with comparable or higher accuracy and less expensive training [64]. Farms are represented as graphs in graph

neural networks, with turbines as nodes and wake or spatial-proximity correlations encoded on the edges. This enables spatiotemporal modeling by utilizing inter-turbine correlation rather than treating each turbine independently for forecasting or potential-power estimation [65,66]. To directly weigh the relevance of distant time steps or turbines, transformer architectures use self-attention and have, thus, consistently outperformed recurrent baselines in recent comprehensive renewable energy forecasting benchmarks [6,23]. However, this comes at the expense of higher data requirements and quadratic attention complexity.

2.2.3. Reinforcement Learning

RL frames control-or-predicting policy adaptation as a series of decisions made in the face of uncertainty. Based on observations of the system state, including wind conditions, the turbine operating point, and prediction inaccuracy, an agent learns a policy through interaction or high-fidelity simulation that maximizes cumulative reward, such as energy capture, forecast accuracy, or minimal load [67]. This is particularly crucial for farm-level problems such as wake steering and induction control, where coupling between turbines limits the availability or reliability of analytic gradients derived from a wake model. In these scenarios, deep reinforcement learning has been directly used for short-term wind speed forecasting as an alternative to supervised sequence models for policy learning [68]. However, sampling inefficiency and exploration-safety problems restrain the practical application of RL in deployed turbine controllers.

2.2.4. Physics-Informed and Hybrid Models

When extrapolating beyond their training distribution, purely data-driven AI models perform badly. This is a major limitation for safety- and load-related wind applications where extreme or unusual events are critical. Physics-informed neural networks (PINNs) embed governing equations (e.g., conservation laws) as soft constraints in their training loss [69], which regularizes the learned function toward physically consistent behavior even in regions with little data. In wind energy, such constraints include wake-deficit physics, structural load calculations, and simplified aerodynamic relations (e.g., power and thrust coefficients as functions of the tip-speed ratio and pitch angle).

Particularly, the sustainability relevance of these methods depends on application objectives and constraints rather than the algorithm family alone. Therefore, method selection should consider the performance gain alongside evaluation costs, training and inference requirements, and robustness. Sometimes sustainability-related objectives may conflict. In this case, multi-objective formulations can settle these conflicts.

3. Wind Speed and Power Forecasting

Each horizon has a distinct function and forecasting needs vary widely across them. Dispatch and reserve allocation are supported in the short term (30 min to 6 h); unit commitment and day-ahead market participation are supported in the medium term (6–72 h); maintenance scheduling and energy trading are supported in the long term (days to seasons); and turbine-level and storage control are supported in the very short term (seconds to 30 min), where the relevant metric is tracking error at high sampling rates [6,70]. Hybrid numerical weather prediction (NWP) plus machine-learning pipelines, in which NWP output is post-processed or bias-corrected by a learned model, are now the operational standard [26,71]. Physical NWP predominates beyond a few hours because it resolves synoptic-scale atmospheric dynamics that no purely statistical model can extrapolate. Statistical and learning methods predominate below because local site-specific patterns are more significant than large-scale weather evolution. This horizon-dependent structure is crucial for benchmarking because forecasting techniques that are hailed as state-of-the-art

on one horizon sometimes transfer poorly to another, and evaluations that aggregate results across horizons without separating them run the risk of overstating generality [72].

3.1. From Statistical Models to Deep Learning

Persistence and the Box–Jenkins ARIMA family remain deceptively strong baselines at very short horizons: complex nonlinear models struggle to convincingly defeat their linear, low-variance structure once forecasting skill is properly cross-validated [73], and Kalman-filter state-space formulations extend this baseline to slowly time-varying statistics. Classical machine learning, including support vector regression, random forests, and gradient boosting, outperformed these baselines mainly by exploiting a richer combination of NWP and SCADA variables rather than through modeling power [61]. Deep learning was then used to alter the field along two axes. LSTM and GRU sequence models directly capture temporal dependence from raw or lightly processed series; convolutional and temporal convolutional (TCN) feature extractors exchange recurrence for parallelizable dilated convolutions [64]; and, most recently, Transformer architectures use self-attention to directly weigh the relevance of distant time steps without the vanishing-gradient constraints of recurrence [74]. Second, empirically, a thorough benchmarking of Transformer-based renewable forecasting reveals systematic accuracy gains over traditional deep learning across solar and wind datasets, while specifically identifying deployment complexity, data hunger, and quadratic attention cost as adoption barriers that point-error metrics alone fail to capture [23]. At the short-to-medium horizon, ensemble techniques that integrate NWP output with statistically post-processed station data remain a robust and reasonably priced substitute for end-to-end deep structures [71]. Systematic reviews of deep forecasting attest to both the field's explosive architectural growth and its fragmentation of benchmarks. Site, horizon, and metric selections differ sufficiently between studies that head-to-head architecture comparisons derived from the literature should be interpreted with caution rather than as a ranking [30].

3.2. Hybrid Decomposition–Optimization–Learning Pipelines

In recent years, the most published architecture is the hybrid decomposition–optimization–learning pipeline. A raw wind speed or power series is first decomposed into components of different frequency content, most commonly by wavelet transforms, empirical mode decomposition (EMD) and its noise-robust variant CEEMDAN, or variational mode decomposition (VMD). This recasts decomposition as a non-recursive variational optimization that adaptively estimates each mode's center frequency and bandwidth rather than relying on recursive sifting [75]. Each resulting component is then modeled separately, typically by a GRU, LSTM, or Transformer, and the component forecasts are recombined into the final prediction, with metaheuristics commonly used to optimize decomposition parameters (decomposition level, mode count) and network hyperparameters jointly [23,24]. A representative recent example combines dual-stage decomposition with metaheuristic-tuned GRU and Transformer sub-models, and it reports lower errors than eleven single and hybrid baselines on four wind speed series from Germany and China [24]. The magnitude of such gains depends mainly on the baseline strength: against the strongest hybrid competitor in that study, the reported improvement ranges from 0.5% to 5.5% depending on the dataset, whereas considerably larger reductions can appear when the baseline is a single untuned model or persistence. Reported percentages should therefore be read together with the baseline, the horizon, and the tuning budget, and three methodological caveats further temper them. First, and most seriously, decomposition is frequently applied to the full series before the train–test split. This leaks future values into the training components through the decomposition's global frequency estimation. Because methods such as EMD and VMD

are not causal, this is a genuine information leak rather than a benign preprocessing step, and it inflates reported accuracy in a way that does not survive deployment where only past data are available at decomposition time [75]. Second, hyperparameter optimization budgets between the proposed hybrid and non-hybrid baselines are rarely matched fairly. So, part of the reported gain may reflect unequal tuning effort rather than architectural superiority. Third, gains obtained on one site or one season often shrink substantially when the same pipeline is retrained on a different location [30]. Figure 2 shows a hybrid decomposition–optimization–learning pipeline for wind speed and power forecasting.

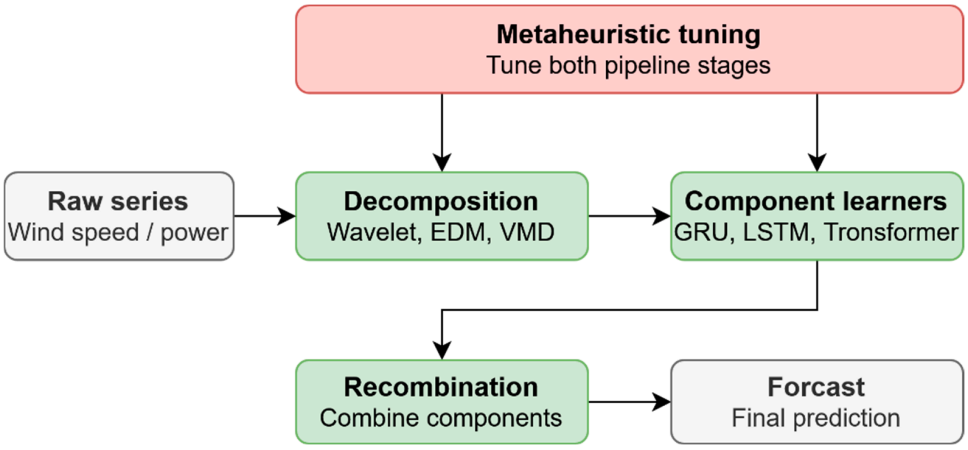


Figure 2. Hybrid decomposition–optimization–learning pipeline for wind speed and power forecasting.

3.3. Probabilistic and Spatiotemporal Forecasting

Grid operation, reserve sizing, and market participation require calibrated distributions, not single point values. Evaluating those distributions requires metrics designed for that purpose. Proper scoring rules such as the pinball (quantile) loss and the continuous ranked probability score (CRPS) reward forecasts that are both sharp and honestly calibrated, unlike RMSE or MAE, which say nothing about whether a model's stated uncertainty is trustworthy [76]. Several probabilistic families have been applied to wind power. Quantile regression directly estimates conditional quantiles rather than a conditional mean. Conformal prediction wraps an arbitrary point forecaster with a distribution-free coverage guarantee, and a recent context-aware variant adapts its intervals to the local, time-varying volatility of wind power rather than assuming stationary error statistics [77]. Generative approaches, most notably denoising-diffusion probabilistic models, learn to sample directly from the full predictive distribution by reversing a learned noising process, with reported calibration improvements over quantile and ensemble baselines under the strongly non-stationary statistics typical of wind power [78]. At the spatial scale, graph neural networks represent a wind farm as a graph of turbines connected by wake or proximity edges and exploit inter-turbine correlation that single-turbine models discard. Applied to offshore SCADA data, a spatio-temporal graph model reduced turbine potential-power estimation error by about 30% relative to the IEC power-curve method, although probabilistic extensions evaluated for joint calibration across a farm remain scarce [65,66]. Despite this progress, the very-short-term horizon, where turbine and storage control operate and where probabilistic information is most actionable, remains comparatively underserved by the probabilistic forecasting literature relative to the day-ahead and medium-term horizons [72].

To maximize sharpness subject to calibration, a probabilistic forecast should be evaluated on two properties: calibration, the statistical consistency between predicted probabilities and observed frequencies, and sharpness, the concentration of the predictive distribution [79]. For prediction intervals, calibration is measured by the prediction interval coverage probability (PICP), the fraction of observations within the interval, which should match the nominal level. Sharpness is measured by the prediction interval average width (PIAW), often normalized by the range of the target [80]. Because either metric can be improved at the expense of the other, a combined score is needed. The Winkler score adds to the interval width a penalty proportional to the distance of each observation falling outside it [81], while the CRPS and the pinball loss evaluate the full distribution or a set of quantiles [79]. Reliability diagrams complement these scores by showing calibration at all nominal levels [82]. Probabilistic forecasts matter because they feed downstream decisions: reserve requirements can be sized from forecast quantiles, stochastic unit commitment uses forecast scenarios to schedule generation under uncertainty [83], and in a two-price imbalance market the optimal bid of a wind producer is a quantile of the predictive distribution determined by the imbalance prices [84]. The link with scheduling is discussed further in Section 5.3. A detailed treatment of power-system operation is beyond the scope of this review.

The recent architectures discussed above (conformal wrappers, diffusion models, graph networks, and Transformers) are usually introduced with comparisons that are not matched. The proposed model and its baselines differ in input features, horizon, data split, or hyperparameter-tuning budget, so the reported gain cannot be attributed to the architecture alone [30]. Matched comparisons, in which only the component under study changes, can reverse the conclusion. In a pipeline in which only the temporal module was swapped, a GRU outperformed the Transformer, which was judged unsuitable given its computational cost [74]. Fair benchmarking requires the same data and split, the same horizon, an equal tuning budget for all candidates, and, for probabilistic models, a report of both calibration and sharpness rather than a single accuracy metric. A sample of forecasting studies is introduced in Table 3.

Table 3. Representative forecasting studies spanning method families, horizons, and evaluation rigor.

Ref	Method	Horizon	Dataset	Results	Output
Ullah et al. [24]	Dual-stage decomposition + optimized GRU + Transformer	Wind speed: 10 min–1 h	4 wind speed series (Germany, China)	R^2: 3.39, 3.93, 5.53, and 0.50% improvement for each dataset, compared to the TVFEMD-MPE-QPSO-GRU-DWT-Autoformer	Point
Jiang et al. [74]	VMD + GNN (GIN) + temporal model (Transformer/GRU/LSTM/MLP)	Wind speed: 1 h–12 h	Shenzhen meteorological tower	Best variant (GRU): MAE 0.1085 m/s (1-h)	Point
Du [71]	Ensemble ML + NWP post-processing	Wind power short-term	2-year data collected at Electricity Reliability Council of Texas	Mean absolute percentage error (MAPE) reduction from 4.57 to 3.74% at 3 h	Point
Liu et al. [78]	Denoising-diffusion generative model	Wind power short-term	NREL WIND Toolkit, 16 MW farm	PIAW: 1.901, PICP: 88.5%, Winkler score: −2.246 (α = 90%); WS: −5.32% vs. best SOTA	Probabilistic
Lahouar & Ben Hadj Slama [61]	RF (up to 18 inputs; no tuning)	Wind power: 1 h	Sidi Daoud wind farm, Tunisia (53.5 MW, 2011, 10-min SCADA)	NA	Point

Table 3. *Cont.*

Ref	Method	Horizon	Dataset	Results	Output
Zhu, Liao & Wang [64]	TCN	Short-term	Wind power series	TCN beats SVM, MLP, LSTM, GRU on both RMSE and MAE, train and test	Point
Daenens et al. [65]	Spatio-temporal graph neural network	Short-term	Offshore farm, >40 WTs	Normalized MAE 0.0333/MAPE 8.65% (test); −30.3% MAE vs. power-curve baseline; undetected power-loss events 0.02% vs. 12.6%	Point

MAE: mean absolute error; RMSE: root mean square error; MAPE: mean absolute percentage error; PICP: prediction interval coverage probability (fraction of observations inside the interval; target = nominal level); PIAW: prediction interval average width (sharpness; lower is better at equal coverage); Winkler score: interval width plus a penalty proportional to the distance of observations falling outside the interval (lower is better); CRPS: continuous ranked probability score.

4. Metaheuristic and AI-Based Control of Wind Energy Conversion Systems

As illustrated in Figure 3, the wind turbine control system is designed around two operating regions: region 2 (R2) and region 3 (R3). In R2, when the wind speed exceeds the cut-in speed, the control system maximizes power output. In R3, the control system maintains power production at the rated value when the wind speed exceeds the rated value. If the wind speed is above the cut-in value but below the rated value (R1), the blades turn very slowly and produce no power. Therefore, the wind turbine turns off. If the wind speed exceeds the cut-off value, the control system should stop the rotation speed at R4.

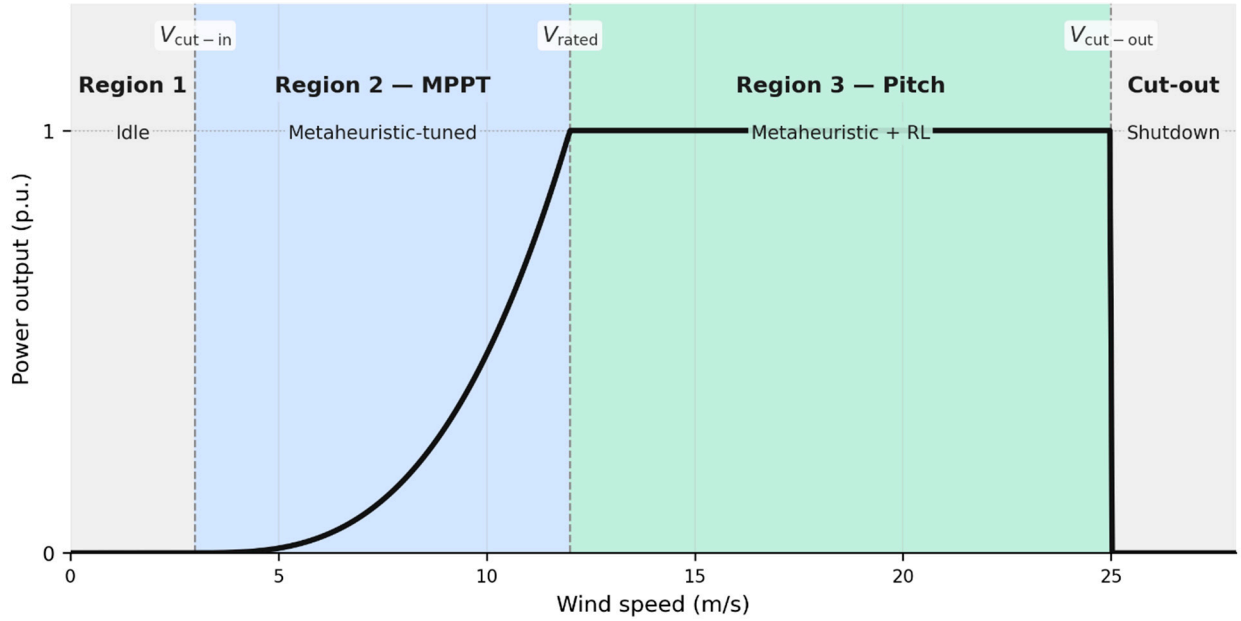


Figure 3. Wind turbine operating regions mapped to the dominant metaheuristic and AI control approach in each.

4.1. MPPT

The control goal is to extract as much power as possible below the rated wind speed. Classical methods, such as tip-speed-ratio (TSR) control, optimal torque, power-signal feedback, and perturb-and-observe (P&O), face a fundamental trade-off between tracking speed and oscillation, despite their conceptual simplicity. Fast adaptation improves tracking

under gusty inflow, but it also increases chattering around the optimum. Conversely, sluggish adaptation limits energy capture during transients but smooths the output [25]. To address this trade-off, metaheuristic approaches use two methods: online metaheuristic searches that re-optimize the operating point and offline optimization of fixed controller gains (such as PI/PID coefficients or fuzzy membership function shapes), which are tuned once against a representative wind profile. One recent example of an offline technique is a fuzzy-logic maximum power point tracking (MPPT) algorithm with membership functions optimized by GA and PSO for permanent magnet synchronous generator (PMSG)-based systems. Unlike traditional P&O, this method has been demonstrated to increase energy yield across simulated wind profiles [44]. An example of an online approach is an MPPT strategy based on the greater cane rat algorithm, which reached a mean tracking efficiency of about 99% under step, ramp, and turbulent wind profiles in simulation, compared with 94–96% for PSO and GWO and 90–91% for P&O [58]. Reviews of MPPT schemes categorize them as direct, indirect, hybrid, or intelligent and assess them according to their robustness to sensor noise and parameter uncertainty, convergence speed, and steady-state efficiency [25].

4.2. Pitch, Torque, and Generator Control

When wind speeds exceed the rated level, pitch control becomes a pivotal mechanism for regulating power output and constraining structural loads. The multi-objective nature of the problem—namely, the conflicting objectives of regulating quality, actuator wear, and fatigue loading—is precisely the type of problem for which population-based multi-objective optimization was designed. This is because a single scalar cost cannot represent these conflicting objectives without an arbitrary weighting choice [85]. Metaheuristic tuning of generator and rotor-side converter controllers has been widely reported for DFIG-based systems. A recent experimental study that employed cuckoo-search and WOA tuning of the direct voltage control of a standalone DFIG under unbalanced loads demonstrated a reduction in overshoot of up to 88%, an acceleration in rise time of 99%, and a suppression of stator-voltage THD of 82%, compared to conventional PI tuning in simulation (overshoot reduced from 36.8% to 4.2%); in the laboratory test, the overshoot fell from 16.4% to 1.5%. Because these percentages refer to one standalone configuration and a conventional PI baseline, they indicate what careful tuning can achieve in that setting rather than a general improvement [86]. This study and the energy-valley-optimized controller discussed below are the only two controller studies in this section validated on laboratory hardware; the other results were obtained in simulation only. A fractional-order sliding mode controller with a perturbation observer (RPO-FOSMC), tuned by a multi-objective grasshopper optimizer, has been proposed to improve the MPPT and fault ride-through capabilities of a DFIG-based turbine [87]. Adaptive fuzzy controllers tuned by energy-valley optimization have been validated experimentally for the power extraction and dynamic response of a grid-tied DFIG [59]. A new line of inquiry is emerging that combines cost weights carefully crafted using metaheuristic algorithms with model predictive control (MPC). This method makes use of MPC's ability to handle explicit constraints and a population-based search over weighting factors, which are often adjusted through trial and error [88].

4.3. Reinforcement Learning and Data-Driven Control

MPPT, pitch scheduling, yaw control, and farm-level wake steering are just a few of the applications of deep reinforcement learning. Instead of using a hand-specified control law, this method learns policies directly from interaction or high-fidelity simulation. For real-time wake steering across a simulated multi-turbine farm row, a representative closed-loop example combines policy-gradient RL with recursive least squares, reporting

mean power improvements in the low double digits under partial-wake conditions [89]. This is an uncommon instance of RL being assessed dynamically rather than using static, pre-calculated yaw tables. In the context of farm-level coordination, where wake interactions between turbines might render analytic gradients from a wake model unavailable or inaccurate, this approach shows particular promise. Conversely, classical and metaheuristic-tuned controllers work well for single-turbine pitch or MPPT control. Four issues separate these simulation results from field deployment. First, sample efficiency: deep RL typically requires far more interactions than a turbine can safely provide, so policies are trained in simulation, and the number of environment steps should be reported alongside the performance. Second, safety during exploration: random exploratory actions on a physical turbine can violate load or grid-code limits, which motivates constrained or shielded formulations in which a certified controller retains authority [90]. Third, sim-to-real transfer: policies trained on engineering wake or aeroelastic models inherit their modeling errors, and techniques such as domain randomization or fine-tuning on measured data have not yet been demonstrated on operating wind turbines [91]. Fourth, benchmarking: RL controllers are rarely compared with classical and metaheuristic-tuned controllers under matched conditions (same plant model, wind fields, constraints, and tuning effort), so their added value over well-tuned conventional designs remains unquantified. Certification requirements, discussed in Section 8.3, add a further obstacle.

Evaluation caveats: The control results presented in Table 4 are not directly comparable, despite similar percentage improvements. First, validation approaches differ: most studies employ MATLAB/Simulink or OpenFAST simulations, whereas two DFIG controllers were validated on laboratory test benches using a dSPACE DS1104 real-time board [59,86]. The reinforcement-learning wake-steering result was derived from dynamic simulation of a four-turbine row [89]. Second, reported improvements are measured against baselines of varying robustness, typically a conventional PI or P&O controller, with tuning procedures often undocumented. As a result, substantial relative gains may reflect weak reference controllers. Third, operating conditions differ, including standalone versus grid-connected machines, step versus turbulent wind profiles, and balanced versus unbalanced loads. Results obtained under one set of conditions are rarely re-evaluated under alternative scenarios. Finally, although metaheuristic tuning is inherently stochastic, variability across repeated runs is infrequently reported. Among the studies in Table 4, only the wake-steering study provides both a mean and a standard deviation [89]. These criteria are systematically compared across domains in Section 8.5.

Table 4. Controller-optimization studies spanning plant type, algorithm, and validation rigor.

Ref	Machine	Controller	Algorithm	Objectives	Key Results
Bouchakour et al. [44]	PMSG	Fuzzy-logic MPPT	GA & PSO (offline gain tuning)	Energy yield vs. P&O	The proposed algorithm's efficiency is higher by 53.6% compared to the other algorithms
AboRas et al. [58]	Generic WECS	MPPT	Greater cane rat algorithm	Tracking efficiency, IAE	Mean tracking efficiency: 99.1–99.2% vs. 94.2–95.8% (PSO, GWO) and 89.6–91.2% (P&O)
Soued et al. [86]	DFIG	Direct voltage control (rotor-side PI)	Cuckoo search + WOA	Transient response, THD	Overshoot: −88%, rise time: −99%, THD: −82% vs. conventional PI (simulation); overshoot: 16.4% vs. 1.5% (laboratory test)
Elnaghi et al. [59]	DFIG	Adaptive fuzzy logic	Energy valley optimizer	Power extraction, dynamic response	Lowest IAE among all controllers compared

Table 4. *Cont.*

Ref	Machine	Controller	Algorithm	Objectives	Key Results
Reddy et al. [88]	DFIG	Model predictive control	PSO, AVOA, GTO (cost-function weighting factor)	Multi-objective MPC design	I_d and I_q steady-state errors of 0.03% and 0.02% with a 0.7% reduction in the switching frequency
Liew et al. [89]	Multi-turbine farm (4-turbine row)	Farm-level wake steering	Policy-gradient RL + recursive least squares	Power gain via wake redirection	Mean power gain: 11.6 ± 3.0%

5. Design and Planning Optimization

5.1. Wind Farm Layout

Layout optimization under wake interaction is the canonical wind metaheuristic issue. The objective function is nonconvex, high-dimensional, and costly to evaluate because turbine positions interact through wake effects. This is because before a fitness value is known, each candidate layout necessitates a complete wake-field computation [4]. GA and PSO variations dominate the literature, with recent additions from novel swarm algorithms and hybrid discrete–continuous encodings that manage both turbine count and location. A layout problem is fully specified only when the wind rose (direction and speed bins, often Weibull-distributed), the wake model and its superposition rule, the objective (annual energy production, wake loss, or levelized cost of energy), the constraints (site boundary, minimum turbine spacing, number of turbines), and the evaluation budget are stated.

In practice, the final layout depends not only on the optimizer but also on two upstream modeling choices: wind-rose fidelity, typically represented by Weibull-distributed speed and direction bins, which determine how much wake loss the objective actually penalizes, and wake-model fidelity. The latter ranges from the simple top-hat Jensen/Park model to the Gaussian analytical model, which agrees better with wind-tunnel and LES data. A round-robin comparison organized within IEA Wind Task 37 illustrates both effects. With a common wake model, different optimizers reached different annual energy production values, and when participants also chose their own wake models, the resulting layouts and energy estimates differed further [42,92]. The relative weight of modeling and optimizer choices, therefore, depends on the case and cannot be stated in general. AI primarily enters as surrogate modeling. Neural networks trained on CFD or engineering wake model outputs accelerate fitness evaluation by orders of magnitude compared with running the wake model directly in the optimization loop. This enables richer layout search spaces and joint layout–cable–foundation optimization, which would otherwise be computationally prohibitive [2].

Figure 4 schematically illustrates the difference between a naive grid layout and a wake-optimized layout. Figure 5 shows the plan view of a single-turbine wake with measurement plane and the corresponding lateral velocity-deficit profile, comparing the top-hat (Jensen/Park) and Gaussian (Bastankhah–Porté-Agel) wake models. The Gaussian profile matches wind-tunnel and LES data more closely than the sharp-edged top-hat model [93].

Several factors complicate layout result comparisons. First, the reported gain depends on the reference layout chosen (e.g., an aligned grid, a staggered grid, or a previously optimized layout), the wake model, and the wind-rose discretization. As a result, identical percentage improvements may correspond to substantially different outcomes. Second, the optimized layout is permanently evaluated using the same engineering wake model adopted in the optimization process. Consequently, the reported gain is often internal to the model, and validation against measured production data or higher-fidelity simulations is uncommon. Third, because each fitness evaluation requires a wake-field computation, the

number of evaluations is the primary cost driver. Nevertheless, competing algorithms are seldom compared using equivalent evaluation budgets (see Section 8.1). Therefore, a rigorous layout study should specify at least four fundamental elements: the reference layout, the wake model and its calibration, the wind-rose resolution, and the evaluation budget.

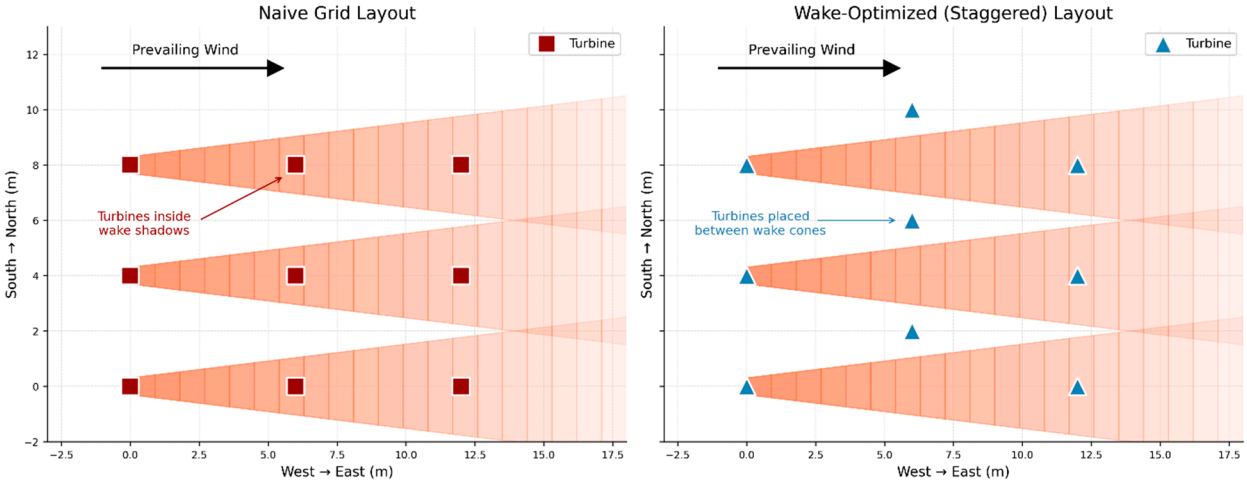


Figure 4. Schematic comparison of a naive grid layout and a wake-optimized layout under a single prevailing wind direction (illustrative, not computed).

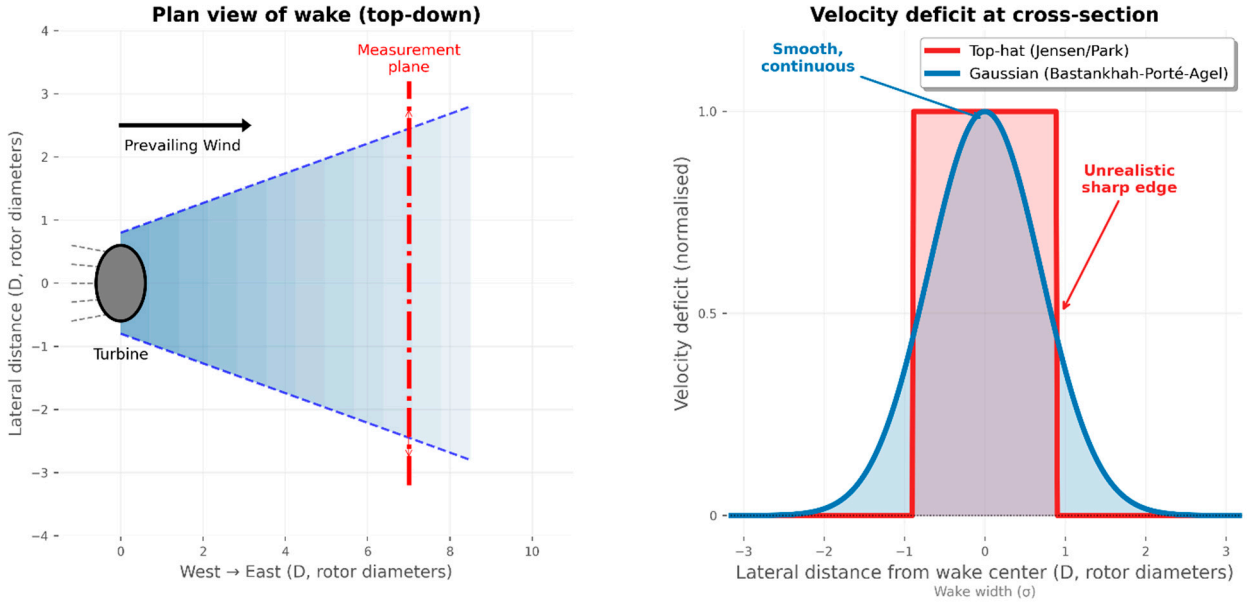


Figure 5. Plan view of a single-turbine wake with measurement plane (**left**) and the corresponding lateral velocity-deficit profile (**right**), comparing the top-hat Jensen model [94] and the Gaussian model of Bastankhah and Porté-Agel. Both profiles are normalized to the same peak deficit and are qualitative (not to scale).

5.2. Turbine and Blade Design

Aerodynamic blade optimization, involving chord and twist-distribution evaluation and airfoil selection, is an inherently multi-objective problem. It aims to resolve conflicts by maximizing the aerodynamic performance while minimizing blade mass and the structural loads that scale with it [95]. This trade-off has been observed in a variety of independent studies, across various specific objective pairings. When applied to the NREL 5 MW reference blade with maximum power coefficient and minimum blade mass as objectives, an enhanced NSGA-II, incorporating controlled elitism and dynamic crowding distance,

yields a Pareto-optimal front rather than a single best design [96]. Similarly, a separate study maximizing annual energy production against minimum blade mass on the same reference turbine by deploying standard NSGA-II with explicit strain constraints reports a qualitative structure analogous to the front of non-dominated designs trading energy yield against mass [97]. The designer then selects from this front based on project priorities, rather than an algorithmically determined "best" answer. Obtaining a front is a property of Pareto-based algorithms and is not evidence of a trade-off. The evidence is that the reported fronts do not collapse to a single design across different objectives and constraint-handling strategies. As the formulation changes and gets more complicated, the same structural pattern remains. The front becomes a trade-off surface instead of a curve with three objectives (adding root bending moment or blade cost) [98]. With explicit material and manufacturing constraints, the admissible front shifts but does not collapse to a point [99]. Optimizing the blade and tower jointly at the system level leads to the emergence of coupling effects that sequential component-wise design cannot capture [100]. Most of these studies report a single algorithm, which prevents conclusions about relative optimizer performance. Vianna Neto et al. [101] is one of the few studies that compare NSGA-II, QMEA, MOEA/D, and MODE on the same blade design problem.

Several caveats must be considered when evaluating the blade-design studies presented in Table 5. Although most studies utilize the NREL 5 MW reference blade as a benchmark, they do not share a common evaluation protocol. In five out of six cases, a single algorithm is applied, making it difficult to attribute the reported Pareto front to the optimizer rather than the problem formulation. The studies also pursue different objectives, including power coefficient, annual energy production, blade mass, root bending moment, and cost, which precludes direct comparison of reported improvements. The comparison of four multi-objective algorithms on the same blade problem represents an exception [101]. Additional like-for-like comparisons, including front-quality indicators such as hypervolume and their dispersion across independent runs, are necessary before any algorithm is recommended for blade design.

Table 5. Design optimization studies.

Ref	Benchmark	Algorithm	Objectives	Constraints	Key Finding
Wang, Wang & Luo [96]	NREL 5 MW	Improved NSGA-II (elitism + dynamic crowding)	Max Cp, min blade mass	Not reported	Pareto-optimal front, not a single optimum
Yang et al. [97]	NREL 5 MW	NSGA-II	Max AEP, min blade mass	Strain constraints	Front of non-dominated, structurally admissible designs
Sessarego et al. [98]	NREL 5 MW	Hybrid NSGA-II (+ gradient local search)	Max AEP, min flapwise root-bending moment, min blade mass	Not reported	3D Pareto surface; hybrid achieved +1.8% AEP, −4.7% root moment; local gradients did not improve convergence
Zhu et al. [99]	1.5 MW-class HAWT	NSGA-II (BEM + FEM in MATLAB/ANSYS loop)	Max AEP, min blade mass, min blade cost	Strain, deflection, vibration, buckling limits	CFRP/GFRP material mix on Pareto front; cost as explicit third objective
Vianna Neto et al. [101]	Original benchmark (Brazilian wind conditions)	NSGA-II vs. QMEA vs. MOEA/D vs. MODE (comparative)	Max energy production, min blade mass	Not reported	Four MOEAs compared on the same problem—rare like-for-like algorithm comparison in blade design
Zhu et al. [100]	Coupled blade–tower system	NSGA-II (system-level)	Max AEP, min total turbine mass	Stress, strain, deflection, vibration, buckling	Extends the trade-off from component to system level; captures blade–tower coupling effects sequential design misses

5.3. Link Between Probabilistic Forecasting and Stochastic Scheduling

Metaheuristics are also applied to macro-siting under multi-criteria constraints, reactive power and voltage optimization for grid-code compliance, and the sizing of hybrid wind–solar–storage plants. These topics fall outside the scope defined in Section 1 and are only mentioned here. Prediction-informed stochastic optimization is the interface where Sections 3 and 5 directly intersect. Instead of using a single-point prediction, unit commitment and storage scheduling decisions are optimized against a probability distribution over future wind power. Therefore, the quality and calibration of the forecast uncertainty representation—rather than just the scheduling method itself—determines the optimization's value [102]. This dependence underscores the significance of the probabilistic forecasting gap discussed in Section 3.3, which extends beyond predictive accuracy alone. As the metrics defined there (coverage, sharpness, and proper scores) determine a forecast's value in scheduling, forecasting studies would be more useful to operators if they reported statistical scores and the downstream cost of their forecasts in a representative scheduling problem. A simpler schedule based on an uncertain forecast distribution may outperform a well-optimized dispatch schedule based on a poorly calibrated one. Furthermore, sustainability benefits should be assessed through downstream outcomes, such as renewable energy utilization and operating cost, rather than inferred from forecast accuracy alone.

6. SCADA-Based Condition Monitoring and Fault Diagnosis

6.1. Data Landscape

SCADA systems are a highly cost-effective data source for monitoring because they capture statistics of hundreds of channels, including power, rotor speed, pitch, temperatures, and vibrations, at 10-min intervals (increasingly at 1 Hz). The reason is that nearly all commercial turbines collect them automatically rather than requiring specialist equipment [7,27]. Public datasets (Penmanshiel, EDP, ENGIE La Haute Borne, Kelmarsh) have enabled replicable study in a field formerly dominated by proprietary operator data. Recent studies have gone beyond single-farm, single-season evaluation by utilizing multi-year, multi-million-record archives combined with maintenance event logs [103]. Because maintenance records reflect what technicians choose to log rather than a controlled fault taxonomy, this scale has two drawbacks: it allows deep architectures that would starve on smaller datasets, and the event logs used as ground truth are themselves weak labels. This limitation is directly addressed in Section 6.3.

6.2. Methods

The coexistence of three methodological generations does not imply that the older ones are out of date. Normal behavior modeling (NBM) avoids the fault-label scarcity that restricts supervised alternatives by using only healthy-condition data to train a regressor (ANN, Gaussian process, gradient boosting) of a target channel, such as gearbox bearing temperature, and monitoring the residual between predicted and observed values for the deviation that signals incipient failure. Comparative studies of kNN-bagging, XGBoost, and ANN coupled to PCC/PCA feature selection and grid, random, or Bayesian hyperparameter search map out this design space in detail [62,104]. Deep anomaly detection extends NBM's residual logic to a learned, nonlinear notion of "normal" across multiple channels simultaneously (rather than one target channel at a time). By using autoencoders, LSTM-AE, and denoising architectures, it reconstructs multivariate operating signatures and detects deviations via reconstruction errors [105]. The structure of a typical first-generation NBM-based gearbox failure detection pipeline is shown in Figure 6. Using healthy SCADA channels, an ANN predicts target temperatures and triggers an alarm if the persistent residual deviation exceeds an empirically tuned threshold. This configuration suppressed

false alarms across an eight-turbine validation fleet, and it detected a planetary bearing failure approximately nine months in advance in the study from which the workflow is derived [106]. This demonstrates why NBM generation remains the industrial workhorse despite its straightforward architecture.

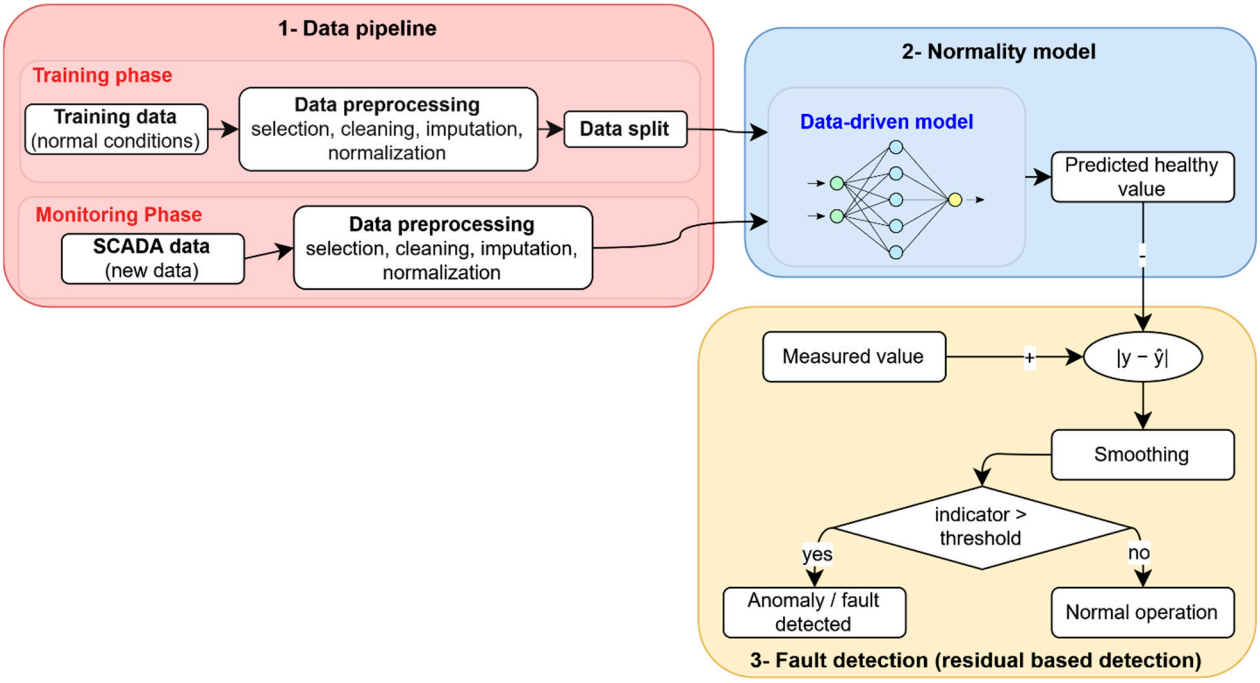


Figure 6. Overview of fault detection workflow model based on SCADA system, ANN-based temperature prediction, and residual/threshold-based anomaly detection. Adapted from [106].

In the study underlying Figure 6 [106], the anomaly indicator is the absolute residual between the measured and predicted gear bearing temperature, smoothed with a one-week moving average. An alarm is raised when this indicator exceeds $\mu + \kappa\sigma$, where μ and σ are computed for each turbine on its healthy training residuals. The factor κ was selected empirically: $\kappa = 6$ triggered short-lived false alarms on two healthy turbines, whereas $\kappa = 9$ triggered an alarm only on the failed turbine, about nine months before the documented failure. Because κ was calibrated on the same test period that contains the only documented failure, this lead time is an in-sample result. Confirming it would require further failure events or other wind farms.

Unlike standalone anomaly flagging, end-to-end deep diagnosis uses CNN, temporal convolutional, graph, and attention-based architectures to directly classify and localize faults. Recent state-of-the-art methods include self-attention-based dynamic graph representation learning and variable-level normalizing flows, which explicitly model relationships between SCADA channels instead of treating them as an unstructured feature vector [107]. A comprehensive study on the monitoring of SCADA systems based on ANNs provides a summary of this rapidly expanding body of literature across all three generations [27].

6.3. Persistent Challenges

Class imbalance is almost structural in this domain because faults are rare by construction (a normal turbine generates far more data than a failing one). Label scarcity and noise aggravate this problem, as operator maintenance logs are inconsistent and frequently record the repair action rather than the root cause. Concept drift caused by seasonal and aging effects implies that a model trained on one year's healthy baseline can misclassify normal behavior one year later. Furthermore, cross-turbine and cross-farm transfers

remain challenging because nominally identical turbines display measurably different SCADA signatures due to micro-siting and individual component variation [7,27]. The most straightforward solution to the last challenge is transfer learning between turbines. It applies a model trained on data-rich turbines to data-poor ones on the same or a separate farm [108]. Additional active responses include weakly supervised formulations that treat maintenance logs as noisy rather than ground-truth labels and physics-informed features derived from known failure mechanisms rather than learned end-to-end. In practice, weak supervision means aligning work orders with alarm and status logs, treating the period that precedes a documented repair as uncertain rather than faulty, and learning from positive and unlabeled data. The studies reviewed here reflect the two ends of this methodological spectrum: a study that uses an event-log-labeled multi-year archive [103], and a study in which the authors had to infer maintenance dates from the monitored signal itself due to missing access to maintenance reports, preventing independent verification of alarm lead times [62]. Metaheuristics contribute to this pipeline mainly in two ways: selecting features and tuning hyperparameters for diagnostic models (e.g., using genetic-algorithm-based selection of ensemble members [109]) and optimizing maintenance downstream scheduling of diagnosis. Their benefit is rarely isolated. Feature selection is frequently performed on the same data used later for evaluation, leading to information leakage and inflated accuracy. Moreover, the gain over simple filter methods (correlation or mutual information) or over budgeted search strategies (grid, random, and Bayesian) is seldom reported. The comparative study in [62] is one of the few that explicitly compares search strategies. Once a fault or degradation trend is identified, determining the time to intervene is a classical condition-based-maintenance optimization problem, extensively studied in prior literature preceding the deep learning methods reviewed above [110].

In SCADA-based diagnosis, the choice of validation protocol can influence reported performance as much as the model architecture itself. Random train–test splits of SCADA records may place temporally adjacent, highly correlated samples on both sides of the split, artificially inflating accuracy. In contrast, chronological splits and held-out turbines yield more realistic performance estimates. However, only a subset of studies in Table 6 use these approaches [106,108]. Detection outcomes are further influenced by the alarm threshold, which is often tuned empirically, and by the ground truth, typically derived from maintenance logs of variable quality (Section 6.3). Metrics most relevant to operators, such as the false-alarm rate per turbine per year and detection lead time, are reported far less often than classification accuracy. While public datasets (Penmanshiel, Kelmarsh, EDP, La Haute Borne) now facilitate reproducibility, many studies still rely on proprietary operator data without releasing code.

A systematic evaluation framework for SCADA-based fault diagnosis should report at least the following. Detection performance must include precision, recall, and F1 score at the operating threshold, alongside the area under the precision–recall curve rather than only the ROC AUC, as faults are rare and the ROC-AUC is optimistic under strong class imbalance. Operational value should list detection lead time before documented failure and the number of false alarms per turbine-year, since an alarm rate that is acceptable on a single turbine becomes unmanageable across a fleet. Label quality requires knowing the source of the ground truth (work orders, alarm logs, or expert annotation) and how ambiguous periods were handled. Finally, robustness must be assessed through temporal performance to reveal concept drift from seasonal and aging effects, and on turbines or farms not used for training. Community guidelines for sharing data and evaluating fault-detection models [111] and recent public datasets with labeled anomaly events and a dedicated scoring procedure [112] are first steps toward such a framework. Few studies in Table 6 report all of these elements.

Table 6. SCADA-based fault diagnosis studies.

Ref	Dataset	Target	Method	Validation Protocol	Metrics/Key Result
Ng and Lim [62]	La Haute Borne (ENGIE, public), 4 turbines, 2014–2016	Gearbox	kNN-bagging, XGBoost, ANN + PCC/PCA feature selection	Held-out turbine (3 train, 1 test); alarms checked against inferred maintenance (no maintenance log)	R^2 0.95, RMSE 0.84 (kNN); alarms about 2–4 weeks before maintenance dates inferred by the authors
Daenens et al. [65]	Offshore farm, >40 turbines >8 MW, 2021–2022 SCADA (confidential)	Turbine potential power	Spatio-temporal graph neural network (LSTM encoder + GENConv)	Chronological split (2021 train; 2022 validation/test); MAE −30.3% vs. IEC power curve; undetected power-loss events 0.02% vs. 12.6%	MAE −30.3% vs. IEC power curve; undetected power-loss events 0.02% vs. 12.6%
Chesterman et al. [104]	5 onshore wind farms, 2 MW turbines (confidential)	Generator bearing, generator fan, rotor brush	Normal behavior modeling (ANN, GP, gradient boosting)	Chronological: first ~6 months of healthy data per turbine for training; 5-fold CV for tuning	Healthy-test RMSE; ΔPE (median prediction error, unhealthy–healthy)
Chen et al. [105]	SCADA multivariate signals	Turbine performance	LSTM-AE deep anomaly detection	Not reported	Adaptive threshold reported to reduce false alarms; values not reported in the abstract
Zhu et al. [107]	SCADA multivariate signals	Fault classification + localization	Self-attention dynamic graph + normalizing flows	Not reported	Not reported
Bronskiy et al. [103]	Penmanshiel (3.6M records) + event logs	Multiple, event-log-labeled	Integrated SCADA analysis + ML	Multi-year, multi-turbine	Not reported
Xie et al. [108]	Cross-turbine SCADA	Multiple (transfer setting)	Deep transfer learning	Cross-turbine generalization test	Not reported
Puruncajas et al. [106]	9-turbine wind farm, 2018–2022 SCADA	Gearbox (planetary bearing)	ANN-based NBM on gear bearing temperature	Chronological train/validate/test split; healthy-only training; validated across 8 turbines	Alarm ~9 months before the documented failure; 0 false alarms at $\kappa = 9$ (2 at $\kappa = 6$); one documented failure

7. Floating Offshore Wind: An Emerging Frontier

Integrating floating offshore wind turbines (FOWTs) into existing offshore wind farm infrastructure further complicates the already complex control problem, introducing new challenges in platform hydrodynamics, mooring dynamics, and wave–wind interactions. By linking platform pitch motion back into the blade-pitch control loop, negative aerodynamic damping, a prominent feature of FOWTs, can destabilize conventional pitch control under above-rated conditions [31,113]. This impact, specific to FOWTs and absent in fixed-bottom turbines, underscores the need for thorough investigation and analysis to ensure the dependable and secure operation of these hybrid energy systems. This phenomenon engenders distinctive opportunities for the methods reviewed here.

Multi-objective metaheuristic tuning of structural-load-aware controllers has already been demonstrated. Specifically, a gray wolf-optimized gain-scheduled controller has been shown to substantially reduce platform pitch motion while preserving rated power output on a barge-type FOWT [31]. This is achieved without the extensive mathematical approximations that analytical gain-scheduling otherwise requires. This line of work has since been extended toward disturbance-rejection and model-free control formulations. An optimized active disturbance rejection controller has been reported to reduce generator-speed fluctuations by 82.9%, torque error by 56.1%, and tower damage-equivalent load by approximately 19% relative to standard ADRC and GSPI baselines [32]. Additionally, a resilient model-free controller has been proposed specifically to couple power regulation with fatigue-load reduction, addressing both objectives that Section 5.2's blade-design discussion identified

as structurally in tension [33]. A separate blade-pitch strategy targeting power regulation and structural-motion alleviation has been reported and experimentally validated [34]. Experimental validation is rare in this field. Along with the laboratory-scale DFIG studies discussed in Section 4 [59,86], it is one of the few control results in this review that were not obtained from simulation alone. Another demonstrated approach with reported load-mitigation benefits is wave-feedforward control, which uses real-time wave forecasts to preemptively adjust blade pitch ahead of an impending disturbance, rather than reacting after platform motion has already begun [114]. On the other hand, the application of artificial intelligence (AI) to monitoring dynamic cables and moorings can be considered a reasonable, though largely unexplored, extension of the SCADA-based diagnostic techniques discussed in Section 6. The inspection of these components is intrinsically costly and does not constitute a well-established body of knowledge, highlighting the importance of this developing topic as a viable avenue for future research rather than a definitive result. Reviews of optimization have shown a clear inclination to focus on offshore systems, including design, architecture, and control. From 2021 to 2025, this trend is more noticeable. The rapid expansion of this literature is notable, but it is also marked by the lack of a cohesive analytical framework across various related fields [2]. Due to the lack of operational FOWT data, high-fidelity simulation (OpenFAST-class aero-hydro-servo-elastic programs) remains the predominant benchmark environment. This approach has been used to validate most of the controllers presented here, and it raises a sim-to-real transfer question that impacts both learned controllers and monitors.

Several of the FOWT control studies cited above come from the authors' group, and the broader literature extends beyond them. The integrated design of the platform, tower, mooring system, and blade-pitch controller has been formulated as a single optimization problem, showing that controller parameters and structural dimensions interact and should not be fixed sequentially [115]. Model-based wave-feedforward pitch control has been designed and evaluated independently of the study cited above [116]. For dynamic power cables, deep learning models have been trained to predict the mechanical response of the cable cross-section as part of digital-twin monitoring, and the authors note that machine-learning models for this task were previously lacking [117]. Because fault data from floating turbines are scarce, few-shot approaches have been proposed for offshore turbine fault diagnosis, including ordinal prototypical networks [118] and meta-learning for generator faults trained on simulated and field data [119]. Three constraints limit the transfer of these methods to operation. First, certification: floating turbines are designed and certified against dedicated standards [120], which require the controller to be verified over a prescribed set of design load cases; learned or online-adapted controllers do not yet fit this process. Second, simulation-to-real transfer: nearly all controllers above are validated in aero-hydro-servo-elastic codes, whose hydrodynamic and mooring models carry their own uncertainty. Third, data scarcity: few floating turbines are in operation, and their data are rarely public, which limits both the training of data-driven monitors and independent validation. This implies that floating offshore sustainability requires the assessment of the coupled turbine, platform, mooring, and cable system. Accordingly, load-aware control, co-design, and monitoring can address production, structural requirements, reliability, and maintenance access. However, simulated motion or fatigue reductions are intermediate outcomes; implications for service life, material demand, and environmental impact require validated lifetime models and operational evidence.

8. Critical Discussion

8.1. Novel-Algorithm Inflation and the No-Free-Lunch Reality

A single survey covering only wind power prediction applications already lists dozens of such algorithms introduced within a few years, directly illustrating the pace of this proliferation [16]. The wind literature, like other energy domains, is home to a proliferation of newly named metaphor-based algorithms, each typically demonstrated on a favorable case study. According to the no-free-lunch theorem, universal superiority claims are meaningless [121]. Reputable contributions must be statistically tested over numerous runs and benchmarked against well-tuned baselines (PSO, DE, CMA-ES, and, for costly issues, Bayesian optimization) under equal evaluation budgets. Only a minority of the controller-tuning and layout studies reviewed here meet this standard. A direct critical assessment of the metaphor-based literature argues that many "novel" algorithms differ from established methods only in narrative framing rather than mechanism [122]; this argument deserves wider citation and enforcement by reviewers in the wind community specifically, since the pattern documented for optimization in general is, if anything, more visible in wind applications given how many recent algorithms use an energy-system case study as their very first validation.

8.2. Benchmark Fragmentation and Data Leakage

Forecasting studies rarely share sites, horizons, and metrics, making the literature nearly impossible to synthesize quantitatively. A systematic review of deep-learning wind forecasting models explicitly flags this fragmentation as a barrier to cumulative progress in the field [30]. Two technical flaws recur across the studies reviewed here, both already identified in Section 3: decomposition-based leakage, where signal decomposition computed on the full series before the train–test split lets future information contaminate training components (Section 3.3), and evaluation restricted to a single site or season, which cannot distinguish a genuinely robust method from one that happens to fit local conditions. A corresponding issue in layout studies exacerbates this: two articles reporting "20% wake-loss reduction" may not be comparable at all because the baseline each improves upon was calculated differently due to incompatible wake models and wind-rose representations. Diagnostic studies introduce the same failure mode that drives the cross-turbine transfer-learning literature covered in Section 6.3: evaluating models on the same turbines used for training. This practice inflates reported accuracy in a way that does not transfer to a new turbine or farm. Community benchmarks can reorganize incentives to support similar, cumulative outcomes. These benchmarks include fixed public datasets, prescribed splits, and required probabilistic indicators, such as pinball loss and CRPS, for forecasting. This strategy aligns with broader machine learning precedents, where shared benchmarks like ImageNet have redefined what constitutes a legitimate contribution.

8.3. From Simulation to Certification

Most metaheuristic control experiments conducted have been verified only using OpenFAST or MATLAB/Simulink simulations. Experimental hardware validation remains rare in the literature. Exceptions include a cuckoo-search/WOA-tuned DFIG controller that reports overshoot, rise-time, and THD improvements confirmed on physical hardware rather than in simulation alone is one exception that has been experimentally validated in the literature [86], and an energy-valley-optimized adaptive fuzzy controller for a DFIG, also validated with a DS1104 board [59]. This laboratory-scale validation shows that hardware validation is feasible and significantly reinforces the study's conclusions, as it is the exception rather than the rule. Nevertheless, other obstacles to industrial adoption, including control software certification, grid code compliance, real-time turbine controller

execution, and fail-safe behavior verification, need to be addressed. Notably, population-based online optimization techniques do not yet sufficiently address these issues. We believe that offline metaheuristic optimization of certifiable controller architectures is the most practical implementation approach. At the same time, online learning is pushed to supervisory levels, where a qualified controller has the final decisions.

8.4. Interpretability, Uncertainty, and Trust

Although deep predictors and diagnostic models are rapidly outperforming transparent baselines, traders and grid operators need calibrated uncertainty rather than a single point estimate, and operators need justifications for maintenance decisions based on alarms. Converging solutions to this gap include conformal and diffusion-based probabilistic forecasting [77,78], attention-based attribution techniques for diagnostic classifiers, and physics-informed hybrid structures that limit learning models toward physically consistent behavior [69]. It should be standard practice to report calibration, not just sharpness or point inaccuracy. It may be argued that a forecaster who is honestly uncertain is less risky than one who is generally accurate but badly calibrated in the tail situations that matter most for reserve sizing.

8.5. Evaluation Practice Across Application Domains

Table 7 compares the four application domains and floating offshore wind against the same evaluation criteria. It summarizes patterns observed in the studies discussed in Sections 3–7 and provides a qualitative synthesis rather than a meta-analysis. Three observations follow. First, the weakest link differs by domain: in forecasting it is the validation protocol (decomposition leakage and single-site evaluation), in control it is the validation level (mostly simulation), in layout and blade design it is comparability (different wake models, wind roses, reference designs, and objectives), and in diagnostics it is the ground truth (weak labels from maintenance logs and same-turbine evaluation). Second, all domains share two weaknesses: the dispersion of stochastic methods over repeated runs is rarely reported, and competing methods are rarely compared under equal computational budgets. Third, reproducibility is highest where public data or reference models exist (open wind datasets, public SCADA archives, NREL reference turbines) and lowest in control and floating offshore wind, where results depend on specific simulation set-ups or test benches. Consequently, the gains reported in Tables 3–6 should be read as evidence that a method can work under the stated conditions, not as a ranking of methods.

Table 7. Evaluation practice across application domains, synthesized from the studies reviewed in Sections 3–7.

Domain	Typical Data/ Benchmark	Typical Baseline	Dominant Validation Protocol	Main Evaluation Weakness	Validation Level Reached
Forecasting	Site-specific wind speed or power series; few public datasets (e.g., NREL WIND Toolkit)	Other deep or hybrid models; persistence and statistical baselines less often	Hold-out split; decomposition sometimes computed before the split	Leakage, single site or season, point metrics without calibration	Offline evaluation on historical data
Control	Simulated turbines (MATLAB/Simulink, OpenFAST); laboratory DFIG benches	Conventional PI or P&O, tuning often undocumented	Step or synthetic wind profiles; single operating condition set	Baseline strength; incomparable operating conditions; single runs	Simulation; laboratory scale in a few studies [59,86]
Layout	Reference or synthetic wind farms; engineering wake models	Aligned or staggered grid; earlier optimized layout	Evaluation with the wake model used for optimization	Dependence on wake model and wind rose; unequal evaluation budgets	Model-based

Table 7. *Cont.*

Domain	Typical Data/ Benchmark	Typical Baseline	Dominant Validation Protocol	Main Evaluation Weakness	Validation Level Reached
Blade design	NREL 5 MW reference blade; BEM and FEM models	Reference blade geometry	Single algorithm; visual comparison of Pareto fronts	Different objectives; no front-quality indicators or repeated runs	Model-based
Diagnostics	Proprietary SCADA; public archives (Penmanshiel, Kelmarsh, EDP, La Haute Borne)	Power-curve or NBM residuals; standard classifiers	Random or chronological split; held-out turbines in a few studies	Weak labels; empirically tuned thresholds; same-turbine evaluation	Historical field data; prospective deployment rare
Floating offshore wind	Aero-hydro-servo-elastic simulation of reference FOWTs	Gain-scheduled PI; standard ADRC	Simulated load cases	Few studies; mostly simulation	Simulation; experimental validation reported in one study [34]

8.6. Problem–Decision–Method Framework

Table 8 links the main problems addressed in this review to their decision level and horizon, their characteristics, and the method families suited to them. The table does not identify a universally superior method; it indicates which families are appropriate under which conditions. Three patterns emerge. Strategic design problems (layout, blade, and floating platform design) are expensive, multi-objective, and evaluated offline. These characteristics make them well-suited to population-based multi-objective search when the model is a black box and gradient-based or surrogate-assisted optimization when the evaluation cost dominates. Tactical problems (controller tuning, day-ahead forecasting, maintenance scheduling) tolerate offline computation but require robustness across operating conditions. Operational problems (real-time control, very-short-term forecasting, monitoring) impose latency and safety constraints that favor simple, verifiable models at the core, with learning and optimization in supervisory or offline roles.

Table 8. Problem–decision–method framework: decision level and horizon, problem characteristics, suitable method families, and conditions of use.

Problem	Decision Level/Horizon	Key Characteristics	Suitable Method Families	Conditions and Limits
Blade design	Strategic (design stage)	Expensive, multi-objective, continuous variables; BEM/FEM models	Multi-objective EAs (NSGA-II, MOEA/D); gradient-based MDO when gradients are available; surrogate-assisted search [96–101]	MOEAs when the model is a black box and a trade-off front is needed; report budget and repeated runs
Wind farm layout	Strategic	Nonconvex, multimodal, many variables, costly wake evaluations	GA/PSO variants; gradient-based methods with continuation; surrogate or Bayesian optimization for expensive cases [2,4,39,42,92]	Results depend on wake model and wind rose; compare optimizers at equal evaluation budgets
FOWT platform–mooring–controller co-design	Strategic	Coupled multidisciplinary problem; costly aero-hydro-servo-elastic simulation	Gradient-based MDO on linearized models; surrogate-assisted metaheuristics [35,113,115]	Final verification over the design load cases required for certification
Controller tuning (MPPT, pitch, generator)	Tactical (commissioning, offline)	Few parameters; simulation-based, often conflicting objectives	Offline single- or multi-objective metaheuristic tuning of a fixed controller structure [44,58,59,86–88]	Robustness across operating conditions must be checked; keep a certifiable structure

Table 8. *Cont.*

Problem	Decision Level/Horizon	Key Characteristics	Suitable Method Families	Conditions and Limits
Real-time turbine control	Operational (ms–s)	Hard real-time, safety-critical	Classical and model predictive control; learning or online optimization only at supervisory level [25,88]	Population-based online search rarely suitable
Farm-level wake steering	Operational (minutes)	Coupled turbines; wake-model uncertainty	Model-based optimization with engineering wake models; RL with model adaptation [89]	Added value of RL over optimization not yet quantified under matched conditions
Very-short- and short-term forecasting	Operational (seconds–6 h)	Nonstationary, autocorrelated, site-specific	Persistence/ARIMA baselines; ML and deep models; decomposition hybrids; probabilistic methods [24,70,72,74]	Must beat persistence; causal decomposition; calibrated intervals
Day-ahead forecasting	Tactical (6–72 h)	Driven by synoptic weather	NWP with ML post-processing; ensembles; probabilistic forecasts [70]	Physical NWP dominates beyond a few hours
Condition monitoring and fault diagnosis	Operational to tactical (days–months)	Rare faults, weak labels, concept drift	Normal behavior models; deep anomaly detection; transfer learning [7,27,104,106,108]	Chronological or held-out evaluation; report false alarms per turbine-year and lead time
Maintenance scheduling	Tactical (weeks–months)	Combinatorial, uncertain, weather-window constraints	Mathematical programming; metaheuristics; condition-based maintenance models [110]	Value depends on the reliability of the diagnostic input

8.7. Integrated and Hybrid Approaches

Hybrid approaches are discussed in several sections of this review. Table 9 brings them together, indicating, for each, the problem addressed, the techniques combined, the main benefits and limitations, representative implementation tools, and the validation level achieved in the reviewed studies. Across all hybrids, the literature-reported benefits are observed mainly offline or in simulation, with the limitations identified in Section 8.5: evaluation leakage, unequal tuning budgets, and the absence of field validation. The tools listed are representative open-source or commercial packages for each approach and are not necessarily those used in the cited studies.

Table 9. Synthesis of integrated and hybrid approaches.

Hybrid Approach	Problem Addressed	Combined Techniques	Main Benefits	Main Limitations	Representative Tools	Validation Level
Decomposition–optimization–learning	Wind speed and power forecasting	EMD/CEEMDAN/VMD + metaheuristic hyperparameter tuning + LSTM/GRU/Transformer [23,24,74]	Handles nonstationarity; high reported accuracy	Leakage with non-causal decomposition; unequal tuning budgets; single-site evaluation	PyEMD, vmdpy, PyTorch, TensorFlow, Optuna	Offline historical data
Metaheuristic-tuned controllers	MPPT, pitch, generator, and FOWT control	PSO/GA/GWO/WOA and variants + PI, fuzzy, sliding-mode, or MPC controllers [31,34,44,58,59,86]	Systematic offline tuning; explicit multi-objective trade-offs	Single operating scenario; robustness and repeated runs rarely reported	MATLAB/Simulink, OpenFAST, ROSCO	Simulation; laboratory scale in two studies

Table 9. *Cont.*

Hybrid Approach	Problem Addressed	Combined Techniques	Main Benefits	Main Limitations	Representative Tools	Validation Level
Surrogate-assisted optimization	Layout, blade, and FOWT design	Neural-network, Gaussian-process, or kriging surrogates + metaheuristic or Bayesian optimization [2,35]	Far fewer expensive evaluations; larger design spaces	Surrogate error and extrapolation; retraining cost	FLORIS, PyWake, WISDEM, OpenMDAO, SMT	Model-based
Physics-informed learning	Forecasting, normal behavior models, wake and structural modeling	Governing equations or physical features embedded in neural-network training [69]	Better extrapolation; lower data requirements	Training difficulty; choice of physical priors	DeepXDE, PyTorch	Mostly simulation or offline data
RL with model-based adaptation	Farm-level wake steering	Policy-gradient RL + recursive least squares [89]	Adapts without precomputed yaw tables	Sample efficiency, safe exploration, sim-to-real transfer	FLORIS-based environments, Stable-Baselines3	Dynamic simulation
Metaheuristics for model design	Diagnostics and forecasting	GA/PSO feature selection and hyperparameter search [62,109]	Automates model configuration	Leakage when selection uses evaluation data; rarely compared with simple baselines	scikit-learn, pymoo, DEAP, Optuna	Offline historical data

8.8. Limitations of This Review

This review has four limitations. First, it is a narrative review: the studies were selected at the authors' discretion, the search results were not systematically screened, and some sub-topics may be under-represented. Second, the quantitative results reported in Tables 3–6 are taken from the original studies and were not reproduced; they inherit the evaluation choices criticized in this review. Third, for several studies whose full text was not accessible to the authors, the table entries were based on the published abstracts, as indicated in the tables. Fourth, several floating offshore wind control studies discussed in Section 7 are from the authors' group, introducing a potential selection bias in that section; to mitigate this, we have also cited independent work on the same topics.

9. Future Research Directions

The fragmentation described in Section 8.2 (incompatible sites, horizons, and metrics in forecasting [30], incomparable wake models and wind roses in layout studies, and same-turbine evaluation in diagnosis) is a structural barrier to cumulative improvement rather than a collection of discrete defects. The field requires transformative benchmarks that reshape machine learning, including fixed public datasets with specified train–test splits, required probabilistic metrics (pinball loss, CRPS [76]) for forecasting, shared wake-model and wind-rose specifications for layout, and held-out turbines for diagnosis. While components exist, as demonstrated by the multi-method layout comparisons already shown on common problems [42] and the public SCADA archives [103], community collaboration to freeze them into standards is still lacking. Where wind engineering most needs them (extreme loads, uncommon faults, and undiscovered sites), purely data-driven models extrapolate poorly. In addition to providing better extrapolation and lower data requirements, physics-informed formulations that incorporate aerodynamic, electrical, or structural priors as training constraints also support the certification arguments that end-to-end black boxes cannot [69]. Practical rather than conceptual open questions remain: which priors are

important for which tasks, how much accuracy is sacrificed for physical consistency, and whether the resulting models remain competitive when training data is plentiful.

Farm-level coordinated optimization. Wake steering, induction control, and storage co-scheduling remain mostly studied in isolation. The closed-loop reinforcement-learning wake-steering demonstration of [89] (dynamic simulation, double-digit power gains under partial wake) is the current reference point, but it has not yet been benchmarked against classical optimization-based wake steering under matched wake models and conditions. Hence, the genuine added value of learning over optimization at farm level remains unquantified. Extending such comparisons to include probabilistic forecast inputs [77,78] would connect the forecasting and control literature that this review has shown to be methodologically adjacent but practically separate.

Floating offshore co-design. Metaheuristic-tuned FOWT control [31] and wave-feedforward load mitigation [114] are demonstrated in simulation, but platform, moorings, and controller are still designed largely sequentially. Joint co-design with multi-fidelity surrogates (cheap engineering models exploring broadly, aero-hydro-servo-elastic codes [113] verifying finalists) is the natural next step, and the surrogate-assisted optimization tooling already validated in other energy applications [35] transfers directly. AI-based monitoring of moorings and dynamic cables remains a gap, driven by extreme inspection costs rather than an active literature; the SCADA-based methods in Section 6 are the obvious starting point, but operational FOWT data scarcity means transfer and few-shot approaches [108] will be prerequisites rather than refinements.

Sample-efficient optimization. Population-based metaheuristics are profligate with function evaluations, which is tolerable for cheap ECM-style objectives but prohibitive for CFD and aeroelastic simulation. Bayesian and surrogate-assisted methods [35,40] should be systematically compared with metaheuristics under equal evaluation budgets on expensive wind problems. The current literature rarely runs this comparison, so the community does not actually know how much performance the standard swarm-based defaults leave on the table.

Hybrid plant integration and grid formation. When wind farms take on grid-service obligations, design, forecasting, and control concerns become intertwined: dispatch depends on both, storage sizing depends on forecast calibration, and hydrogen coupling adds a layer [102]. Rather than being an internal forecasting metric in this context, probabilistic forecast quality directly yields optimization value [76]. Consequently, the multi-objective and uncertainty-aware machinery outlined here must operate in tandem rather than across distinct papers.

Table 10 links each direction to a research question, a benchmark, the metrics that should be reported, and the validation required before a claim of practical benefit.

Table 10. Future research directions with associated research questions, benchmarks, metrics, and validation requirements.

Direction	Research Question	Benchmark/Data	Metrics	Validation Requirement
Community benchmarks	Do reported gains survive fixed data, splits, and horizons?	Public wind and SCADA datasets with prescribed splits	MAE/RMSE per horizon; CRPS, pinball loss; F1, lead time, false alarms per turbine-year	Held-out sites, seasons, and turbines
Physics-informed learning	Which physical priors improve extrapolation, and at what accuracy cost?	Rare-event and out-of-distribution test sets	Error on extreme and unseen conditions; physical-constraint violations	Comparison with purely data-driven models at equal data size
Farm-level coordinated control	Does learning add value over optimization-based wake steering?	Common wake model and wind fields	Farm power gain; load changes; sample count	Matched comparison; then high-fidelity or field test

Table 10. *Cont.*

Direction	Research Question	Benchmark/Data	Metrics	Validation Requirement
FOWT co-design	How much do integrated platform–mooring–controller designs gain over sequential design?	Reference FOWT models	Cost proxy; platform motion; fatigue loads	Design load cases required for certification
Sample-efficient optimization	When do Bayesian or surrogate methods outperform metaheuristics on expensive problems?	Expensive layout, blade, and aeroelastic cases	Best objective vs. number of evaluations; dispersion over runs	Equal evaluation budgets; repeated runs
Hybrid plant integration	How does forecast calibration translate into dispatch and storage value?	Representative scheduling problems	Operating cost; imbalance cost; reliability indices	Out-of-sample scheduling with probabilistic forecasts

10. Conclusions

Artificial intelligence (AI) and metaheuristic optimization algorithms (MOAs) are currently used in wind energy engineering, including forecasting, control, farm design, and predictive maintenance to ensure sustainable energy generation. This review has discussed published studies reporting the benefits of these methods. These case studies include surrogate-accelerated layout optimization; SCADA-based normal behavior models that detected gearbox failure several months in advance in individual cases; hybrid deep forecasting pipelines whose reported gains depend strongly on the baseline; and two metaheuristic-tuned generator controllers validated on laboratory hardware. But the field's expansion has surpassed its methodological discipline. Most published results have limited transferability due to fragmented benchmarks, the proliferation of metaphor-driven algorithms, leakage-prone pipelines, and simulation-only validation. Moreover, in every application domain examined here, hardware validation is still the exception rather than the rule. Building complex networks or introducing new metaphors will have less of an impact on the next stage of evolution. Instead, it will be enhanced by the rigor and integration of different components. Key focus areas will include co-design techniques for floating offshore systems, physics-informed hybrid models, uncertainty-aware forecasting combined with optimization, and community benchmarks for easier comparison of claims. With these foundational elements in place, computational intelligence can transition from a prolific research theme to a reliable engineering substrate for the terawatt-scale wind fleets of the impending decades.

In terms of method choice, the evidence reviewed here supports three conclusions (Table 8). For strategic design problems, multi-objective metaheuristics and surrogate-assisted optimization are appropriate when the models are black boxes and trade-off fronts are needed; however, results depend on modeling choices and remain model-based. For tactical problems, offline metaheuristic tuning of certifiable controller structures and NWP-based machine-learning forecasts are the most mature options, validated mostly in simulation and on historical data. For operational problems, normal behavior models for monitoring and well-benchmarked statistical or deep forecasters are the closest to practice; reinforcement learning and online optimization remain at the simulation stage. The sustainability implications follow the same pattern and should be read with this validation level in mind: layout optimization, MPPT, and wake steering target energy yield; load-aware pitch and floating-platform control and early fault detection target reliability, maintenance cost, and asset life; calibrated probabilistic forecasts reduce reserve and imbalance costs and ease grid integration; and integrated co-design addresses lifecycle cost from the design stage. Quantifying these benefits in the field, rather than in simulation, is the key step in demonstrating their contributions to sustainable wind energy.

Author Contributions: Conceptualization, A.M.N. and W.M.; methodology, A.M.N. and M.A.E.; software, A.M.N. and W.M.; validation, A.M.N., W.M. and M.A.E.; formal analysis, A.M.N. and W.M.; investigation, A.M.N. and M.A.E.; resources, A.M.N.; data curation, A.M.N., W.M. and M.A.E.; writing—original draft preparation, A.M.N. and W.M.; writing—review and editing, A.M.N., W.M. and M.A.E.; visualization, A.M.N. and M.A.E.; supervision, A.M.N.; project administration, A.M.N.; funding acquisition, A.M.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Prince Sattam bin Abdulaziz University, project number PSAU/2025/RV/11.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing does not apply to this article.

Acknowledgments: The authors extended their appreciation to Prince Sattam bin Abdulaziz University for funding this research work through the project number PSAU/2025/RV/11.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. World Wind Energy Association. *Global Statistics—Wind Power Capacity Worldwide*; World Wind Energy Association (WWEA): Bonn, Germany, 2026.
2. Zhang, F.; Wang, H.; Zhang, Y.; Miao, W.; Hou, C. A Review of Mathematical Optimization Methods in Offshore Wind Energy Systems: Design, Layout, and Control. *J. Mar. Sci. Eng.* **2026**, *14*, 1033. [CrossRef]
3. Ren, Z.; Verma, A.S.; Li, Y.; Teuwen, J.J.E.; Jiang, Z. Offshore Wind Turbine Operations and Maintenance: A State-of-the-Art Review. *Renew. Sustain. Energy Rev.* **2021**, *144*, 110886. [CrossRef]
4. Herbert-Acero, J.; Probst, O.; Réthoré, P.-E.; Larsen, G.; Castillo-Villar, K. A Review of Methodological Approaches for the Design and Optimization of Wind Farms. *Energies* **2014**, *7*, 6930–7016. [CrossRef]
5. Rezk, H.; Olabi, A.G.; Wilberforce, T.; Sayed, E.T. A Multi-Faceted Review of Wind Turbine Optimization Techniques: Metaheuristics and Related Issues. *Int. J. Thermofluids* **2025**, *26*, 101077. [CrossRef]
6. Wang, Y.; Zou, R.; Liu, F.; Zhang, L.; Liu, Q. A Review of Wind Speed and Wind Power Forecasting with Deep Neural Networks. *Appl. Energy* **2021**, *304*, 117766. [CrossRef]
7. Pandit, R.; Astolfi, D.; Hong, J.; Infield, D.; Santos, M. SCADA Data for Wind Turbine Data-Driven Condition/Performance Monitoring: A Review on State-of-Art, Challenges and Future Trends. *Wind Eng.* **2023**, *47*, 422–441. [CrossRef]
8. Dantzig, G.B. *Linear Programming and Extensions*; Princeton Landmarks in Mathematics and Physics; Princeton University Press: Princeton, NJ, USA, 2016; ISBN 978-0-691-05913-6.
9. Bazaraa, M.S.; Sherali, H.D.; Shetty, C.M. *Nonlinear Programming: Theory and Algorithms*, 3rd ed.; Wiley-Interscience: Hoboken, NJ, USA, 2006; ISBN 978-0-471-48600-8.
10. Bellman, R. *Dynamic Programming*; Princeton Landmarks in Mathematics; First Princeton landmarks in mathematics edition; Princeton University Press: Princeton, NJ, USA, 2010; ISBN 978-0-691-14668-3.
11. Holland, J.H. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*; The MIT Press: Cambridge, MA, USA, 1992; ISBN 978-0-262-27555-2.
12. Kennedy, J.; Eberhart, R. Particle Swarm Optimization. In *Proceedings of the ICNN'95-International Conference on Neural Networks, Perth, WA, Australia, 27 November–1 December 1995*; IEEE: Piscataway, NJ, USA, 1995; Volume 4, pp. 1942–1948.
13. Storn, R.; Price, K. Differential Evolutio–A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *J. Glob. Optim.* **1997**, *11*, 341–359. [CrossRef]
14. Mirjalili, S.; Mirjalili, S.M.; Lewis, A. Grey Wolf Optimizer. *Adv. Eng. Softw.* **2014**, *69*, 46–61. [CrossRef]
15. Mirjalili, S.; Lewis, A. The Whale Optimization Algorithm. *Adv. Eng. Softw.* **2016**, *95*, 51–67. [CrossRef]
16. Sharma, G.; Lal, M.; Attwal, K.P.S. An Extensive Review Of Metaheuristic Algorithms For Wind Power Prediction. *IOSR J. Comput. Eng.* **2025**, *27*, 31–38. [CrossRef]
17. Ferahtia, S.; Houari, A.; Rezk, H.; Djerioui, A.; Machmoum, M.; Motahhir, S.; Ait-Ahmed, M. Red-Tailed Hawk Algorithm for Numerical Optimization and Real-World Problems. *Sci. Rep.* **2023**, *13*, 12950. [CrossRef] [PubMed]
18. Cortes, C.; Vapnik, V. Support-Vector Networks. *Mach. Learn.* **1995**, *20*, 273–297. [CrossRef]
19. Breiman, L. Random Forests. *Mach. Learn.* **2001**, *45*, 5–32. [CrossRef]
20. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]

21. Bai, S.; Kolter, J.Z.; Koltun, V. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv* **2018**, arXiv:1803.01271.
22. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention Is All You Need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 6000–6010.
23. Hussan, U.; Wang, H.; Peng, J.; Jiang, H.; Rasheed, H. Transformer-Based Renewable Energy Forecasting: A Comprehensive Review. *Renew. Sustain. Energy Rev.* **2026**, *226*, 116356. [CrossRef]
24. Ullah, S.; Chen, X.; Han, H.; Wu, J.; Dong, J.; Liu, R.; Ding, W.; Liu, M.; Li, Q.; Qi, H.; et al. A Novel Hybrid Ensemble Approach for Wind Speed Forecasting with Dual-Stage Decomposition Strategy Using Optimized GRU and Transformer Models. *Energy* **2025**, *329*, 136739. [CrossRef]
25. Abdullah, M.A.; Yatim, A.H.M.; Tan, C.W.; Saidur, R. A Review of Maximum Power Point Tracking Algorithms for Wind Energy Systems. *Renew. Sustain. Energy Rev.* **2012**, *16*, 3220–3227. [CrossRef]
26. Haq, I.U.; Kumar, A.; Rathore, P.S. Machine Learning Approaches for Wind Power Forecasting: A Comprehensive Review. *Discov. Appl. Sci.* **2025**, *7*, 1139. [CrossRef]
27. Sheng, L.; Li, C.; Gao, M.; Xi, X.; Zhou, D. A Review of SCADA-Based Condition Monitoring for Wind Turbines via Artificial Neural Networks. *Neurocomputing* **2025**, *633*, 129830. [CrossRef]
28. Moshtaghi, P.; Hajialigol, N.; Rafiei, B. A Comprehensive Review of Artificial Intelligence Applications in Wind Energy Power Generation. *Sustain. Futures* **2025**, *9*, 100638. [CrossRef]
29. Wan, A.; Liao, X.; Iqbal, A.; Al-Bukhaiti, K.; Zhu, P.; Cheng, X.; Ji, X. Machine Learning in Offshore Wind Energy Systems: A Systematic Review (2022–2025). *J. Mech. Sci. Technol.* **2026**, *40*, 6937–6945. [CrossRef]
30. Manzano, E.A.; Nogales, R.E.; Rios, A. A Systematic Review of Wind Energy Forecasting Models Based on Deep Neural Networks. *Wind* **2025**, *5*, 29. [CrossRef]
31. Ferahtia, S.; Houari, A.; Machmoum, M.; Ait-Ahmed, M.; Saim, A. Optimal Control Strategy for Floating Offshore Wind Turbines Based on Grey Wolf Optimizer. *Appl. Sci.* **2023**, *13*, 11595. [CrossRef]
32. Ferahtia, S.; Houari, A.; Machmoum, M.; Ait-Ahmed, M. Robust Pitch-Angle Control of Floating Offshore Wind Turbines Using Optimized Active Disturbance Rejection Control. *Smart Grids Sustain. Energy* **2026**, *11*, 7. [CrossRef]
33. Ferahtia, S.; Houari, A.; Hamida, M.A.; Machmoum, M.; Ait-Ahmed, M. A Resilient Model-Free Controller for Power Regulation and Fatigue Load Reduction in Floating Offshore Wind Turbines. *Mar. Struct.* **2026**, *109*, 104081. [CrossRef]
34. Ferahtia, S.; Houari, A.; Machmoum, M.; Mojallizadeh, M.R.; Ait-Ahmed, M.; Bonnefoy, F. Floating Offshore Wind Turbine Optimized Control for Power Regulation with Experimental Validation. *IEEE J. Ocean. Eng.* **2025**, *50*, 1231–1243. [CrossRef]
35. Thebelt, A.; Tsay, C.; Lee, R.M.; Sudermann-Merx, N.; Walz, D.; Tranter, T.; Misener, R. Multi-Objective Constrained Optimization for Energy Applications via Tree Ensembles. *Appl. Energy* **2022**, *306*, 118061. [CrossRef]
36. Malisani, P.; Bartement, T.; Bozonnet, P. Offshore Wind Farm Layout Optimization with Alignment Constraints. *Wind Energy Sci.* **2025**, *10*, 1611–1623. [CrossRef]
37. Talbi, E. *Metaheuristics: From Design to Implementation*, 1st ed.; Wiley: Hoboken, NJ, USA, 2009; ISBN 978-0-470-27858-1.
38. Yang, X.-S. *Nature-Inspired Metaheuristic Algorithms*, 2nd ed.; Luniver Press: Frome, UK, 2010; ISBN 978-1-905986-28-6.
39. Grady, S.A.; Hussaini, M.Y.; Abdullah, M.M. Placement of Wind Turbines Using Genetic Algorithms. *Renew. Energy* **2005**, *30*, 259–270. [CrossRef]
40. Hansen, N. The CMA Evolution Strategy: A Tutorial. *arXiv* **2016**, arXiv:1604.00772.
41. Hansen, N.; Müller, S.D.; Koumoutsakos, P. Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evol. Comput.* **2003**, *11*, 1–18. [CrossRef] [PubMed]
42. Thomas, J.J.; Baker, N.F.; Malisani, P.; Quaeghebeur, E.; Perez-Moreno, S.S.; Jasa, J.; Bay, C.; Tilli, F.; Bieniek, D.; Robinson, N.; et al. A Comparison of Eight Optimization Methods Applied to a Wind Farm Layout Optimization Problem. *Wind Energy Sci.* **2023**, *8*, 865–891. [CrossRef]
43. Yin Kwong, W.; Yun Zhang, P.; Romero, D.; Moran, J.; Morgenroth, M.; Amon, C. Multi-Objective Wind Farm Layout Optimization Considering Energy Generation and Noise Propagation With NSGA-II. *J. Mech. Des.* **2014**, *136*, 091010. [CrossRef]
44. Bouchakour, A.; Zarour, L.; Bessous, N.; Bechouat, M.; Borni, A.; Zaghba, L.; Rabehi, A.; Alwabli, A.; El-Abd, M.; Ghoneim, S.S.M. MPPT Algorithm Based on Metaheuristic Techniques (PSO & GA) Dedicated to Improve Wind Energy Water Pumping System Performance. *Sci. Rep.* **2024**, *14*, 17891. [CrossRef] [PubMed]
45. Blum, C. Ant Colony Optimization: Introduction and Recent Trends. *Phys. Life Rev.* **2005**, *2*, 353–373. [CrossRef]
46. Karaboga, D.; Basturk, B. A Powerful and Efficient Algorithm for Numerical Function Optimization: Artificial Bee Colony (ABC) Algorithm. *J. Glob. Optim.* **2007**, *39*, 459–471. [CrossRef]
47. Qais, M.H.; Hasanien, H.M.; Alghuwainem, S. A Grey Wolf Optimizer for Optimum Parameters of Multiple PI Controllers of a Grid-Connected PMSG Driven by Variable Speed Wind Turbine. *IEEE Access* **2018**, *6*, 44120–44128. [CrossRef]

48. Mohamed, A.-A.A.; Haridy, A.L.; Hemeida, A.M. The Whale Optimization Algorithm Based Controller for PMSG Wind Energy Generation System. In *Proceedings of the 2019 International Conference on Innovative Trends in Computer Engineering (ITCE), Aswan, Egypt, 2–4 February 2019*; IEEE: Piscataway, NJ, USA, 2019; pp. 438–443.
49. Mahmoud, M.M.; Ratib, M.K.; Aly, M.M.; Abdel–Rahim, A.-M.M. Application of Whale Optimization Technique for Evaluating the Performance of Wind-Driven PMSG Under Harsh Operating Events. *Process Integr. Optim. Sustain.* **2022**, *6*, 447–470. [CrossRef]
50. Mirjalili, S.; Gandomi, A.H.; Mirjalili, S.Z.; Saremi, S.; Faris, H.; Mirjalili, S.M. Salp Swarm Algorithm: A Bio-Inspired Optimizer for Engineering Design Problems. *Adv. Eng. Softw.* **2017**, *114*, 163–191. [CrossRef]
51. Qais, M.; Hasanien, H.M.; Alghuwainem, S. Salp Swarm Algorithm-Based TS-FLCs for MPPT and Fault Ride-through Capability Enhancement of Wind Generators. *ISA Trans.* **2020**, *101*, 211–224. [CrossRef] [PubMed]
52. Kirkpatrick, S.; Gelatt, C.D.; Vecchi, M.P. Optimization by Simulated Annealing. *Science* **1983**, *220*, 671–680. [CrossRef] [PubMed]
53. Rashedi, E.; Nezamabadi-pour, H.; Saryazdi, S. GSA: A Gravitational Search Algorithm. *Inf. Sci.* **2009**, *179*, 2232–2248. [CrossRef]
54. Mirjalili, S. SCA: A Sine Cosine Algorithm for Solving Optimization Problems. *Knowl.-Based Syst.* **2016**, *96*, 120–133. [CrossRef]
55. Abualigah, L.; Diabat, A.; Mirjalili, S.; Abd Elaziz, M.; Gandomi, A.H. The Arithmetic Optimization Algorithm. *Comput. Methods Appl. Mech. Eng.* **2021**, *376*, 113609. [CrossRef]
56. Askarzadeh, A. A Novel Metaheuristic Method for Solving Constrained Engineering Optimization Problems: Crow Search Algorithm. *Comput. Struct.* **2016**, *169*, 1–12. [CrossRef]
57. Mohammadi-Balani, A.; Nayeri, M.D.; Azar, A.; Taghizadeh-Yazdi, M. Golden Eagle Optimizer: A Nature-Inspired Metaheuristic Algorithm. *Comput. Ind. Eng.* **2021**, *152*, 107050. [CrossRef]
58. AboRas, K.M.; EL-Banna, M.H.; Megahed, A.I.; Hammad, M.R. Optimal Maximum Power Point Tracking Strategy Based on Greater Cane Rat Algorithm for Wind Energy Conversion System. *Sci. Rep.* **2025**, *15*, 32474. [CrossRef] [PubMed]
59. Elnaghi, B.E.; Ismaiel, A.M.; El Sayed Abdel-Kader, F.; Abelwhab, M.N.; Mohammed, R.H. Validation of Energy Valley Optimization for Adaptive Fuzzy Logic Controller of DFIG-Based Wind Turbines. *Sci. Rep.* **2025**, *15*, 711. [CrossRef] [PubMed]
60. Yakoub, G.; Mathew, S.; Leal, J. Power Production Forecast for Distributed Wind Energy Systems Using Support Vector Regression. *Energy Sci. Eng.* **2022**, *10*, 4662–4673. [CrossRef]
61. Lahouar, A.; Slama, J.B.H. Hour-Ahead Wind Power Forecast Based on Random Forests. *Renew. Energy* **2017**, *109*, 529–541. [CrossRef]
62. Ng, E.Y.-K.; Lim, J.T. Machine Learning on Fault Diagnosis in Wind Turbines. *Fluids* **2022**, *7*, 371. [CrossRef]
63. Wei, J.; Wu, X.; Yang, T.; Jiao, R. Ultra-Short-Term Forecasting of Wind Power Based on Multi-Task Learning and LSTM. *Int. J. Electr. Power Energy Syst.* **2023**, *149*, 109073. [CrossRef]
64. Zhu, R.; Liao, W.; Wang, Y. Short-Term Prediction for Wind Power Based on Temporal Convolutional Network. *Energy Rep.* **2020**, *6*, 424–429. [CrossRef]
65. Daenens, S.; Verstraeten, T.; Daems, P.-J.; Nowé, A.; Helsen, J. Spatio-Temporal Graph Neural Networks for Power Prediction in Offshore Wind Farms Using SCADA Data. *Wind Energy Sci.* **2025**, *10*, 1137–1152. [CrossRef]
66. Scarselli, F.; Gori, M.; Tsoi, A.C.; Hagenbuchner, M.; Monfardini, G. The Graph Neural Network Model. *IEEE Trans. Neural Netw.* **2008**, *20*, 61–80. [CrossRef] [PubMed]
67. Sutton, R.S.; Barto, A. *Reinforcement Learning: An Introduction*, 2nd ed.; Adaptive Computation and Machine Learning; The MIT Press: Cambridge, MA, USA; London, UK, 2020; ISBN 978-0-262-03924-6.
68. Band, S.S.; Lin, T.J.; Qasem, S.N.; Ameri, R.; Shahmirzadi, D.; Aslam, M.S.; Pai, H.-T.; Salwana, E.; Mosavi, A. A Deep Reinforcement Learning Approach for Wind Speed Forecasting. *Eng. Appl. Comput. Fluid Mech.* **2025**, *19*, 2498355. [CrossRef]
69. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations. *J. Comput. Phys.* **2019**, *378*, 686–707. [CrossRef]
70. Soman, S.S.; Zareipour, H.; Malik, O.; Mandal, P. A Review of Wind Power and Wind Speed Forecasting Methods with Different Time Horizons. In *Proceedings of the North American Power Symposium 2010, Arlington, TX, USA, 26–28 September 2010*; IEEE: Piscataway, NJ, USA, 2010; pp. 1–8.
71. Du, P. Ensemble Machine Learning-Based Wind Forecasting to Combine NWP Output with Data from Weather Station. *IEEE Trans. Sustain. Energy* **2018**, *10*, 2133–2141. [CrossRef]
72. Tawn, R.; Browell, J. A Review of Very Short-Term Wind and Solar Power Forecasting. *Renew. Sustain. Energy Rev.* **2022**, *153*, 111758. [CrossRef]
73. Wilson, G.T. Time Series Analysis: Forecasting and Control, 5th Edition. *J. Time Ser. Anal.* **2016**, *37*, 709–711. [CrossRef]
74. Jiang, W.; Liu, B.; Liang, Y.; Gao, H.; Lin, P.; Zhang, D.; Hu, G. Applicability Analysis of Transformer to Wind Speed Forecasting by a Novel Deep Learning Framework with Multiple Atmospheric Variables. *Appl. Energy* **2024**, *353*, 122155. [CrossRef]
75. Dragomiretskiy, K.; Zosso, D. Variational Mode Decomposition. *IEEE Trans. Signal Process.* **2013**, *62*, 531–544. [CrossRef]
76. Gneiting, T.; Raftery, A.E. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Am. Stat. Assoc.* **2007**, *102*, 359–378. [CrossRef]

77. Moradi, A.; Tanneau, M.; Zandehshahvar, R.; van Hentenryck, P. Enhanced Renewable Energy Forecasting Using Context-Aware Conformal Prediction. *arXiv* **2025**, arXiv:2510.15780.
78. Liu, J.; Zang, H.; Cheng, L.; Ding, T.; Wei, Z.; Sun, G. Generative Probabilistic Forecasting of Wind Power: A Denoising-Diffusion-Based Nonstationary Signal Modeling Approach. *Energy* **2025**, *317*, 134576. [CrossRef]
79. Gneiting, T.; Balabdaoui, F.; Raftery, A.E. Probabilistic Forecasts, Calibration and Sharpness. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2007**, *69*, 243–268. [CrossRef]
80. Khosravi, A.; Nahavandi, S.; Creighton, D.; Atiya, A.F. Lower Upper Bound Estimation Method for Construction of Neural Network-Based Prediction Intervals. *IEEE Trans. Neural Netw.* **2010**, *22*, 337–346. [CrossRef] [PubMed]
81. Winkler, R.L. A Decision-Theoretic Approach to Interval Estimation. *J. Am. Stat. Assoc.* **1972**, *67*, 187–191. [CrossRef]
82. Pinson, P.; Nielsen, H.A.; Møller, J.K.; Madsen, H.; Kariniotakis, G.N. Non-parametric Probabilistic Forecasts of Wind Power: Required Properties and Evaluation. *Wind Energy* **2007**, *10*, 497–516. [CrossRef]
83. Bouffard, F.; Galiana, F.D. Stochastic Security for Operations Planning with Significant Wind Power Generation. *IEEE Trans. Power Syst.* **2008**, *23*, 306–316. [CrossRef]
84. Bitar, E.Y.; Rajagopal, R.; Khargonekar, P.P.; Poolla, K.; Varaiya, P. Bringing Wind Energy to Market. *IEEE Trans. Power Syst.* **2012**, *27*, 1225–1235. [CrossRef]
85. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **2002**, *6*, 182–197. [CrossRef]
86. Soued, S.; Boureguig, K.; Chabani, M.S.; Himair Swhli, K.M. Experimental Validation of Metaheuristic-Optimized Control for Standalone DFIG Dynamic Performance Enhancement. *Sci. Rep.* **2026**, *16*, 10432. [CrossRef] [PubMed]
87. Falehi, A.D. An Innovative Optimal RPO-FOSMC Based on Multi-Objective Grasshopper Optimization Algorithm for DFIG-Based Wind Turbine to Augment MPPT and FRT Capabilities. *Chaos Solitons Fractals* **2020**, *130*, 109407. [CrossRef]
88. Reddy, K.; Sarma, R.; Guha, D. Performance Analysis of Advanced Metaheuristics for Optimal Design of Multi-Objective Model Predictive Control of Doubly Fed Induction Generator. *Processes* **2025**, *13*, 221. [CrossRef]
89. Liew, J.; Göçmen, T.; Lio, W.H.; Larsen, G.C. Model-free Closed-loop Wind Farm Control Using Reinforcement Learning with Recursive Least Squares. *Wind Energy* **2024**, *27*, 1173–1187. [CrossRef]
90. García, J.; Fernández, F. A Comprehensive Survey on Safe Reinforcement Learning. *J. Mach. Learn. Res.* **2015**, *16*, 1437–1480.
91. Zhao, W.; Queralta, J.P.; Westerlund, T. Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: A Survey. In *Proceedings of the 2020 IEEE Symposium Series on Computational Intelligence (SSCI), Canberra, Australia, 1–4 December 2020*; IEEE: Piscataway, NJ, USA, 2020; pp. 737–744.
92. Baker, N.F.; Stanley, A.P.; Thomas, J.J.; Ning, A.; Dykes, K. Best Practices for Wake Model and Optimization Algorithm Selection in Wind Farm Layout Optimization. In *Proceedings of the AIAA Scitech 2019 Forum*; American Institute of Aeronautics and Astronautics: San Diego, CA, USA, 2019.
93. Bastankhah, M.; Porté-Agel, F. A New Analytical Model for Wind-Turbine Wakes. *Renew. Energy* **2014**, *70*, 116–123. [CrossRef]
94. Jensen, N.O. *A Note on Wind Generator Interaction*; Risø-M: Roskilde, Denmark, 1983; ISBN 978-87-550-0971-4.
95. Hansen, M. *Aerodynamics of Wind Turbines*, 3rd ed.; Routledge: London, UK, 2015; ISBN 978-1-317-67103-9.
96. Wang, L.; Wang, T.; Luo, Y. Improved Non-Dominated Sorting Genetic Algorithm (NSGA)-II in Multi-Objective Optimization Studies of Wind Turbine Blades. *Appl. Math. Mech.* **2011**, *32*, 739–748. [CrossRef]
97. Yang, Y.; Li, C.; Zhang, W.; Yang, J.; Ye, Z.; Miao, W.; Ye, K. A Multi-Objective Optimization for HAWT Blades Design by Considering Structural Strength. *J. Mech. Sci. Technol.* **2016**, *30*, 3693–3703. [CrossRef]
98. Sessarego, M.; Dixon, K.R.; Rival, D.E.; Wood, D.H. A Hybrid Multi-Objective Evolutionary Algorithm for Wind-Turbine Blade Optimization. *Eng. Optim.* **2015**, *47*, 1043–1062. [CrossRef]
99. Zhu, J.; Cai, X.; Gu, R. Multi-Objective Aerodynamic and Structural Optimization of Horizontal-Axis Wind Turbine Blades. *Energies* **2017**, *10*, 101. [CrossRef]
100. Zhu, J.; Zhou, Z.; Cai, X. Multi-Objective Aerodynamic and Structural Integrated Optimization Design of Wind Turbines at the System Level through a Coupled Blade-Tower Model. *Renew. Energy* **2020**, *150*, 523–537. [CrossRef]
101. Neto, J.X.V.; Guerra, E.J., Jr.; Moreno, S.R.; Ayala, H.V.H.; Mariani, V.C.; dos Santos Coelho, L. Wind Turbine Blade Geometry Design Based on Multi-Objective Optimization Using Metaheuristics. *Energy* **2018**, *162*, 645–658. [CrossRef]
102. Wu, Q.; Peng, L.; Han, G.; Shu, J.; Yuan, M.; Wang, B. Deep-Learning-Based Scheduling Optimization of Wind-Hydrogen-Energy Storage System on Energy Islands. *Energy* **2025**, *320*, 135107. [CrossRef]
103. Bronskiy, S.; Cheng, J.C.P.; Cheung, C.K.C.; Jiang, T. Enhancing Wind Turbine Fault Prediction through Integrated SCADA Data Analysis and Machine Learning Techniques. *IET Conf. Proc.* **2025**, *2025*, 259–265. [CrossRef]
104. Chesterman, X.; Verstraeten, T.; Daems, P.-J.; Nowé, A.; Helsen, J. Overview of Normal Behavior Modeling Approaches for SCADA-Based Wind Turbine Condition Monitoring Demonstrated on Data from Operational Wind Farms. *Wind Energy Sci.* **2023**, *8*, 893–924. [CrossRef]

105. Chen, H.; Liu, H.; Chu, X.; Liu, Q.; Xue, D. Anomaly Detection and Critical SCADA Parameters Identification for Wind Turbines Based on LSTM-AE Neural Network. *Renew. Energy* **2021**, *172*, 829–840. [CrossRef]
106. Puruncajas, B.; Castellani, F.; Vidal, Y.; Tutivén, C. Use of Artificial Neural Networks and SCADA Data for Early Detection of Wind Turbine Gearbox Failures. *Machines* **2025**, *13*, 746. [CrossRef]
107. Zhu, Y.; Xie, B.; Wang, A.; Qian, Z. Wind Turbine Fault Detection and Identification via Self-Attention-Based Dynamic Graph Representation Learning and Variable-Level Normalizing Flow. *Reliab. Eng. Syst. Saf.* **2025**, *253*, 110554. [CrossRef]
108. Xie, P.; Zhang, X.; Jiang, G.; Cui, J.; He, Q. Investigation of Deep Transfer Learning for Cross-Turbine Diagnosis of Wind Turbine Faults. *Meas. Sci. Technol.* **2023**, *34*, 044009. [CrossRef]
109. Khan, P.W.; Yeun, C.Y.; Byun, Y.C. Fault Detection of Wind Turbines Using SCADA Data and Genetic Algorithm-Based Ensemble Learning. *Eng. Fail. Anal.* **2023**, *148*, 107209. [CrossRef]
110. Jardine, A.K.S.; Lin, D.; Banjevic, D. A Review on Machinery Diagnostics and Prognostics Implementing Condition-Based Maintenance. *Mech. Syst. Signal Process.* **2006**, *20*, 1483–1510. [CrossRef]
111. Barber, S.; Izagirre, U.; Serradilla, O.; Olaizola, J.; Zugasti, E.; Aizpurua, J.I.; Milani, A.E.; Sehnke, F.; Sakagami, Y.; Henderson, C. Best Practice Data Sharing Guidelines for Wind Turbine Fault Detection Model Evaluation. *Energies* **2023**, *16*, 3567. [CrossRef]
112. Gück, C.; Roelofs, C.M.A.; Faulstich, S. CARE to Compare: A Real-World Benchmark Dataset for Early Fault Detection in Wind Turbine Data. *Data* **2024**, *9*, 138. [CrossRef]
113. Jonkman, J.M. *Dynamics Modeling and Loads Analysis of an Offshore Floating Wind Turbine*; NREL: Golden, CO, USA, 2007; p. 921803.
114. Ma, Y.; Sclavounos, P.D.; Cross-Whiter, J.; Arora, D. Wave Forecast and Its Application to the Optimal Control of Offshore Floating Wind Turbine for Load Mitigation. *Renew. Energy* **2018**, *128*, 163–176. [CrossRef]
115. Hegseth, J.M.; Bachynski, E.E.; Martins, J.R.R.A. Integrated Design Optimization of Spar Floating Wind Turbines. *Mar. Struct.* **2020**, *72*, 102771. [CrossRef]
116. Fontanella, A.; Al, M.; Van Wingerden, J.-W.; Belloli, M. Model-Based Design of a Wave-Feedforward Control Strategy in Floating Wind Turbines. *Wind Energy Sci.* **2021**, *6*, 885–901. [CrossRef]
117. Liu, J.; Li, B. A Deep Learning Model for Predicting Mechanical Behaviors of Dynamic Power Cable of Offshore Floating Wind Turbine. *Mar. Struct.* **2025**, *99*, 103705. [CrossRef]
118. Jin, Z.; Xu, Q.; Jiang, C.; Wang, X.; Chen, H. Ordinal Few-Shot Learning with Applications to Fault Diagnosis of Offshore Wind Turbines. *Renew. Energy* **2023**, *206*, 1158–1169. [CrossRef]
119. Qiao, L.; Zhang, Y.; Wang, Q.; Li, D.; Peng, S. Fault Diagnosis for Wind Turbine Generators Based on Model-Agnostic Meta-Learning: A Few-Shot Learning Method. *Expert Syst. Appl.* **2025**, *267*, 126171. [CrossRef]
120. Internationale Elektrotechnische Kommission (Ed.) *Design Requirements for Floating Offshore Wind Turbines*, 1st ed.; Wind Energy Generation Systems/International Electrotechnical Commission Part 3; International Electrotechnical Commission: Geneva, Switzerland, 2019; ISBN 978-2-8322-5986-3.
121. Wolpert, D.H.; Macready, W.G. No Free Lunch Theorems for Optimization. *IEEE Trans. Evol. Comput.* **1997**, *1*, 67–82. [CrossRef]
122. Sörensen, K. Metaheuristics—The Metaphor Exposed. *Int. Trans. Oper. Res.* **2015**, *22*, 3–18. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.