International Journal of
*Turbomachinery
Propulsion and Power*

MDPI

Review

The Reliability Paradox: Machine Learning Applications in Industrial Fans and the Perspectives of Industry Experts

Lorenzo Tieghi [1,*], Giovanni Delibra [2] and Lorenzo Battisti [1]

1 Fluid Machinery Laboratory, Department of Civil, Environmental and Mechanical Engineering, University of Trento, 38123 Trento, Italy; lorenzo.battisti@unitn.it
2 Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, 00184 Rome, Italy; giovanni.delibra@uniroma1.it
* Correspondence: lorenzo.tieghi@unitn.it

Abstract

The integration of Artificial Intelligence (AI) in turbomachinery and fan systems is transforming traditional design, diagnostics, and operational strategies. Artificial Intelligence allows for the efficient exploration of wide design space, easy and fast prediction of fan performance and improving existing system operation and maintenance. Nevertheless, this AI-driven revolution still raises concerns and diffidence in the community, as highlighted by the results of a survey delivered to over 100 fan experts and discussed in this paper. This manuscript aims to provide an overview of Fan-AI applications through a comprehensive literature review of notable use cases. The applications target different stages of the life cycle of fans, from ML-assisted three-dimensional design/optimization to data-driven performance prediction, AI-driven fan control and fault analysis/prognosis. For each of these categories, the relevant application are discussed, highlighting trends, adopted algorithms and strategies, as well as limiting factors. This study also shares the views of experts on both fan design, optimization and operations and AI methods in the upcoming challenges for fan industry. Starting from the need of high-quality data, the improvement of model generalization and the embedding of Fan-AI in the standard engineering practices. This paper concludes with a discussion on the future role of AI in fans, suggesting pathways for research and industrial adoption that balance technological innovation with domain-specific constraints.

Keywords: Industrial-AI; fan engineering; machine learning; expert survey; review; fault prognosis; digital transformation

Check for updates

Academic Editor: Marcello Manna

Received: 28 January 2026
Revised: 13 March 2026
Accepted: 25 May 2026
Published: 5 June 2026

1. Introduction

Data-driven methods, Machine Learning (ML) and Artificial Intelligence (AI) are affecting many engineering sectors, including turbomachinery (TM). AI applications in TM already cover different scopes, including design and optimization, validation [1], aerodynamics and acoustics performance prediction [2], maintenance, prognosis [3], control and manufacturing [4].

Despite the wide range of end uses, the ML applications share a common characteristic: being specifically designed, trained and optimized for a single task. Following the taxonomy of Gabsi et al. [5], it can be generally referred to as Narrow-AI, and underlines the highly-functional and segregated operations of the AI. Narrow-AI models are extremely proficient in solving specific tasks, with limited extrapolation and (frequent) lack of self-learning capabilities. Narrow-AI can be seen as the current stage of development of turbomachinery applications.

However, the growing demand of AI methods coupled with the increased maturity of the topic is pointing to a novel paradigm of the so called Industrial-AI [6], based on self-learning, interconnection and adaptability of the models. However, at the moment only a limited number of Industrial-AI applications have been successfully implemented, like the IoT-based diagnostic of turbines of [7].

The View of Fan Experts on AI

The fan industry is also affected by the changes in technology driven by AI. The audience of the Fan 2025. conference in Antibes [8]—one of the most relevant event in the fan community—was invited to a panel session on the topic, and was asked to compile a survey to disclose its views on the subject. Questions to the participants focused on basics of machine learning, their views on the potential applications, expectations and limiting factors. Over one hundred experts replied, mostly belonging to the industry sector (≈75% of the participants). The most relevant issues are discussed in this section. A first interesting result regards the general expectations on the role of AI in the in the next 5–10 years within the fan industry. To this question, 57.3% of the surveyed audience, considering all levels of expertise, anticipated a significant impact on their workflow, while 34.4% expected a moderate effect on fan-related research (Figure 1).

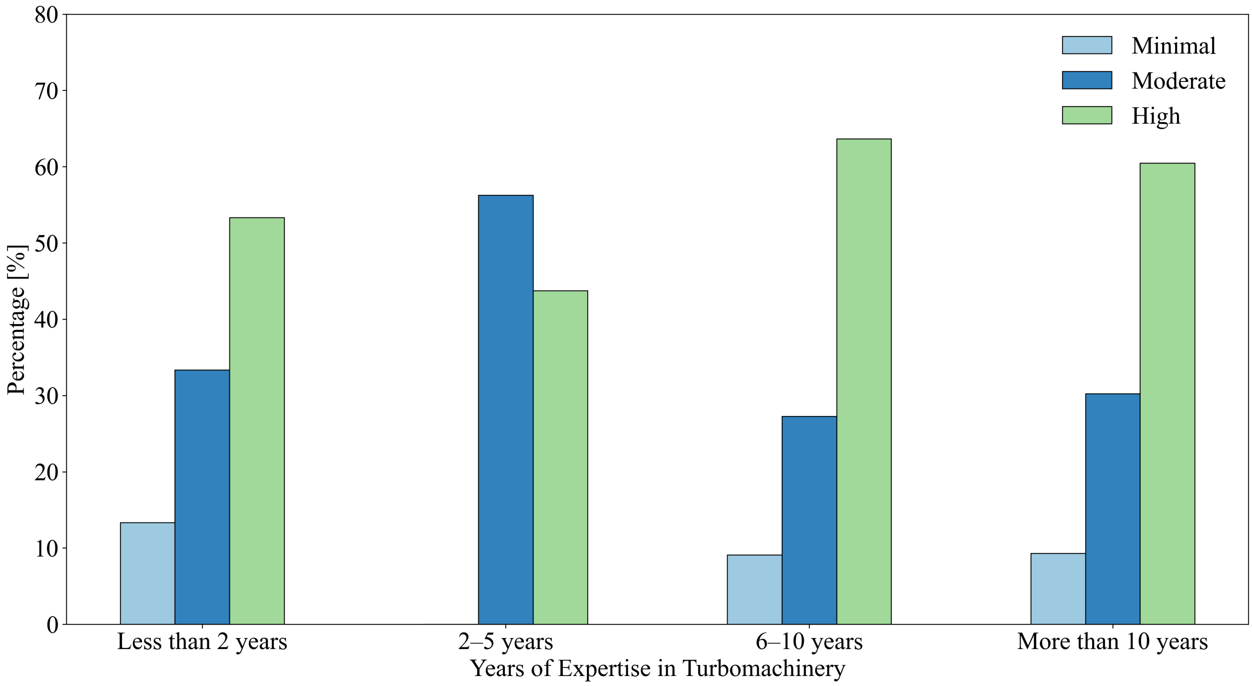

Figure 1. Fan 2025 survey: *"What impact do you expect from AI/ML applications in the turbomachinery field in the next 5–10 years?"*.

Among the possible applications of ML, the most anticipated are improvements in optimization techniques for fan aerodynamics and support for O&M, as prognostic or failure detection tools. However, this positive outlook on ML/AI also conceals a general lack of confidence in ML methods, as highlighted by the results reported in Figure 2.

The general impression is that experts surveyed believe that ML methods are less reliable than traditional approaches regardless of their professional level. We also must point out that this lack of confidence seems stronger among younger engineers. Also, some of the veterans in the field do not trust ML at all. This outcome cannot be ascribed to

https://doi.org/10.3390/ijtpp11020028

the sole lack of knowledge, as more than 80% of the audience considered themselves as moderately familiar with machine learning methods. Therefore an apparent paradox arises: on one hand, the fan community expects ML to play a central role in the near future; on the other, they think these methods may produce unreliable results. A compelling question thus arises: *"Are we really experiencing an artificial intelligence-driven revolution for fans? If so, is it possible that this will lead to undesired reliability issues?"*

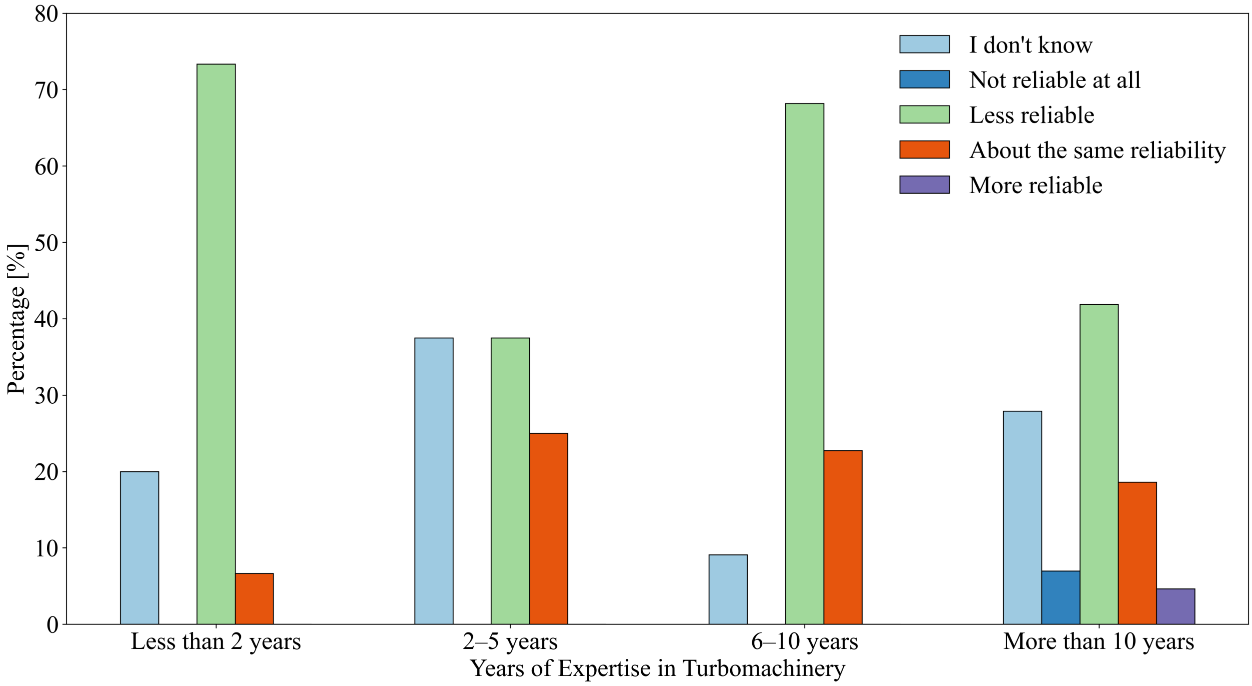

Figure 2. Fan 2025 survey: *"How would you rate the reliability of AI/ML-driven design and validation methods compared to traditional CFD and experimental approaches?"*.

This review aims thus to provide a breakdown of existing and effective applications of fan-AI, highlighting capabilities, recent trends, and also the limitations. The manuscript is organized as follows.

The next sections present a comprehensive literature review of relevant works that explore the application of artificial intelligence and machine learning methods to fans. This literature survey excludes research conducted on turbomachinery that may be relevant to compression processes (i.e., pumps and compressors), although some considerations regarding the applicability of these methods to fans are discussed. For a more in-depth understanding, readers may refer to related literature surveys (e.g., [9–13]).

Here, the survey has led to a classification of the works in four main themes, based on the final purposes of machine learning, following the scheme reported in Figure 3. Then, a section dedicated to the critical analysis of discuss follows. The last part of the paper is dedicated to a discussion on the authors' view of open challenges and possible upcoming developments of the discipline. Conclusions then summarize the content of this article.

Applications of Artificial Intelligence to Industrial Fans

3D Design and Optimization

- Surrogate-based and data-driven optimization of fan geometries using neural networks, ensemble learning and evolutionary algorithms, applied to axial, centrifugal and mixed-flow fans [14–19]
- Multi-objective blade shape optimization targeting efficiency, pressure rise, operating range and noise mitigation across diverse operating conditions [18–21]
- Machine-learning-assisted optimization of blade features and passive flow control strategies for enhanced aerodynamic performance and stability margins [16,20,21]

Performance Prediction

- Data-driven prediction of pressure–flow characteristics, efficiency maps, operating envelopes and off-design behavior using regression and deep learning models [16,22–24]
- Reduced-order, similarity-based and hybrid physics–ML models for scalable performance estimation, map reconstruction and generalization across fan sizes [17,22,24]
- Prediction of flow fields, wakes, aerodynamic loading and aeroacoustic quantities using data-driven and hybrid modeling approaches [25–28]

AI-driven Control and Operation

- Identification and classification of operating regimes, including stall and near-instability conditions, using supervised and unsupervised learning techniques [29–31]
- Intelligent control and monitoring strategies for efficiency maximization, energy savings and operational robustness in ventilation and industrial fan systems [30,32,33]
- Adaptive and learning-based control frameworks, including reinforcement learning, for dynamic fan operation under variable boundary and load conditions [29,30,33]

Fault Analysis and Prognosis

- Machine-learning-based fault detection, diagnosis and classification using vibration, acoustic, electrical and operational data [28,31,34,35]
- Condition monitoring and early fault diagnosis employing classical ML methods and deep learning architectures for industrial fan systems [28,34,35]
- Prognostic assessment of performance degradation, blade fouling and abnormal operating conditions using non-intrusive data-driven techniques [31,35]

Figure 3. Overview of the revised works, grouped by application [14–35].

2. Three-Dimensional Design and Optimization

Aerodynamic optimization and the development of novel semi-empirical design correlations through data-driven methods are by far the most popular ML applications in fan industry. Blade shape optimization, in particular, has been frequently carried out coupling evolutionary algorithms with ML regressors used as cost/fitness functions. In fact, ML methods have been proved to be particularly effective in the context of multi-objective optimization, thanks to their high accuracy and reduced computational costs for algorithm inference, surpassing traditional approaches like semi-empirical models or Computational Fluid Dynamics (CFD). Within this context, a part from few exceptions, supervised learning methods based on neural networks and their modifications are predominant. When complex blade topology is involved, unsupervised learning dimensionality reduction algorithms may find application, e.g., Principal Component Analysis, Projection to Latent Basis, Active Subspaces. In fact, these methods map the highly-dimensional original feature space to a latent and reduced space. This operation preserves—almost perfectly—the statistical characterization of the original dataset, allowing the optimization algorithms to operate on the reduced dimensional space. Consequently, computational costs of the training operations are reduced and ML model with lower complexity can be adopted.

https://doi.org/10.3390/ijtpp11020028

2.1. Axial Fans

In the work of Sun et al. [15], 11 geometrical and operational parameters are used as input for a Neural Network (NN) to target static pressure rise in a axial flow fan. The training dataset was generated through lab tests of 19 small axial fans for electrical equipment cooling. The fans were characterized by different geometrical and operating characteristics, including solidity, blade geometry, rotational speed and volumetric flow rates. Based on this data collection, two separate Deep Neural Networks (DNNs) were trained to predict flow rate and static pressure rise respectively. The two regression models were then coupled with a Genetic Algorithm (GA) to optimize new rotor designs. The authors also highlighted how this approach, even if time-consuming at the beginning, may be extremely beneficial in industrial applications, where the original dataset is constantly expanded due to the development of new products.

Bamberger et al. [14] proposed and validated a surrogate-based optimization methodology for axial fans, to reduce design time while maintaining high prediction accuracy of aerodynamic performance. The approach was based on a combination of:

- A Design of Experiments (DOE) approach to generate the training dataset of 4000+ RANS computations;
- A multi-objective optimization algorithm (NSGA-II);
- A NN surrogate model, trained on a dataset of 3D RANS simulations.

The authors defined the fan geometry using a set of 10 design parameters, including hub-to-tip ratio, blade number, chord length, stagger angle and blade thickness and performed computations of the aerodynamic performance of the fan at design and off-design conditions. These data were used to train a NN surrogate model, which was then coupled with a multi-objective optimization algorithm (NSGA-II) to derive the optimal fan design. The authors found that the surrogate model was able to predict the fan performance with a good accuracy, reducing the computational cost of the optimization process by up to 90% compared to the traditional RANS approach.

In contrast to the predominant use of multi-layer perceptrons in machine learning applications, a recent study by Cheng et al. [21] demonstrates the capability of a boosting algorithms. In this work, a dynamic XGBoost model was trained to predict the aerodynamic performance of an aeroengine fan. The authors adopted a quite complex framework to achieve improvements in the isentropic efficiency, pressure ratio and mass flow rate of the fan. The blade geometries were first parametrized to perform a Latin Hypercube Sampling and then solved using standard CFD. Based on the numerical results, XGBoost was optimized to create a map between fan performance and the original parameters. They further applied the Shapley method to visualize and explain the model behavior and to quantify its uncertainty. Finally, an optimization algorithm was applied to improve the fan aerodynamic robustness. At different stages of the framework, the authors have adopted solutions to reduce model overfitting, particular detrimental in boosting algorithms. The reported approach, although complex, greatly supports the design phase thanks to its high interpretability—superior to standard deep learning methods.

An advanced neural network architecture, Dendrite Networks (DNNs), was employed within an NSGA-II optimization framework to optimize the radial distribution of axial fans [19]. While DNNs share several features with conventional neural networks, they more closely mimic natural biological processes, which enables enhanced predictive performance with lower model complexity. The industrial fan here in analysis is equipped with outlet guide vanes and simulated through CFD means to build a training dataset for the DNN tuning, utilizing maximum efficiency and total pressure rise as target variables. Then, the blade section geometrical parameters and radial distribution of the twisting law are fed to a NSGA-II algorithm to perform the fan optimization, exploiting the pre-trained DNN for the individual fitness evaluation. The optimized fan, measured experimentally, achieves a notably increase by 2.47% and 5.55% in total efficiency and pressure rise, respectively, albeit with increased tip secondary flows. Nevertheless, the potential benefits of DNNs over traditional neural networks in fan optimization require further investigation, due to the lack of studies adopting this family of algorithms.

ML methods have also been applied for the derivation of optimal configurations of enhanced fan geometries that implement passive flow strategies. For example, in [20], the authors aim to optimize the design of an axial fan with leading edge serrations to improve the aerodynamic and aeroacoustics performance. The dataset is generated through an experimental campaign, deriving 65 time- and frequency-domain metrics as targets. In particular, design parameters are set to serration amplitude and wavelength, flow coefficient and inlet turbulent intensity. The four target variables come in terms of aerodynamics performance—fan pressure coefficient and efficiency—and acoustic emissions—overall sound pressure levels and overall sound power level. The relationship between input and target parameters is built using a neural network that is embedded in a multi-objective optimization based on particle-swarm to determine the best serration geometries. The tests carried out at 1/3-octave band spectra seems to confirm the generalization capability of this approach, making this methodology suitable for complex blade topologies.

Angelini et al. [18,36] applied a machine-learning model for the aerodynamic optimization of reversible blades for axial fans commonly found in tunnel and metro ventilation systems.

In [18], the authors formulate a multi-objective optimization problem using the NSGA-II genetic algorithm, targeting aerodynamic efficiency and stall margin. The reversible geometry of the airfoils allows the parameterization of the profile using a sixth-degree B-spline of the suction side. The aerodynamic performance of the airfoils is here evaluated using XFoil solver. The link between profile parameters and aerodynamics is also generated by training a two surrogate models based on LSM and ANN respectively. These surrogate models are adopted within the NSGA-II as alternative fitness evaluation methods. Moreover, two distinct integration strategies are applied: a simple-level direct fitness replacement approach in which optimization relies exclusively on surrogate predictions, and a bi-level indirect fitness replacement, where surrogate-based optimization is periodically corrected using XFoil. A comprehensive test matrix covering a wide range of Reynolds numbers and specified lift coefficients representative of reversible fan operating characteristics is generated. The results show that the single-level approach may lead to unreliable and false Pareto fronts, primarily due to poor exploration of the design space and progressive divergence between the predictions of the surrogate models and true values. In contrast, the bi-level strategy yields Pareto fronts that closely match or outperform those obtained using direct XFoil-based optimization, while significantly reducing the number of expensive fitness-function evaluations, resulting in up to 50% reduced computational time. To this extent, both NN- and LSM-based surrogates are quite effective within the bi-level framework, with the former generally providing higher variance in the solution space, as shown in Figure 4.

The same authors have numerically investigated through a two-dimensional RANS campaign the [36] the influence of four key parameters—Reynolds number, solidity, pitch (stagger) angle and angle of cascade (AoC)—on lift, drag, deflection capability and stall margins. A Central Composite Design (CCD) within a Response Surface Methodology (RSM) framework is adopted to minimize the number of required simulations and allowing the derivation of second-order polynomial surrogate models.

Meta-models for lift and drag coefficients are derived using least-squares regression and further refined through Analysis of Variance (ANOVA) to retain only statistically significant terms. The improved models achieve high predictive accuracy for lift and moderate accuracy for drag while substantially reducing computational cost. The metamodels are then integrated into a quasi-3D in-house fan performance solver, replacing conventional XFoil-based polar data. Comparisons with full 3D CFD simulations show notable improved spanwise prediction of velocities, flow deflection and load coefficient, and a reduction of specific-work overprediction from approximately 50% to about 15%.

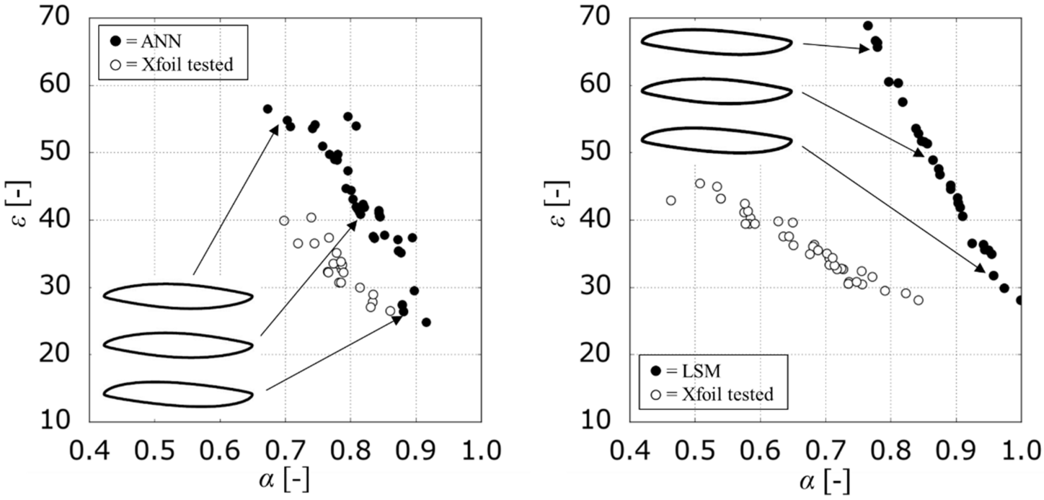

Figure 4. Comparison between algorithm performance in a multi-objective optimization problem of a reversible fan airfoil (angle of incidence α versus stall margin ε), from [18].

Few attempts to model losses in fans can be found in the literature. In [37], unsupervised clustering methods are applied to two-dimensional fan blades with sinusoidal leading edges. The authors have first modified NACA 4 digit profiles with a sinusoidal leading edge, with variable amplitude and wavelength. These profiles are then simulated through RANS simulations at variable angle of incidence, leading to a training dataset with 76 samples. The results of the numerical campaign are used first to evaluate cascade losses adopting the Lieblein's theory. The flow field is then partitioned using a combination of PCA and Gaussian Mixture clustering, leading to a decomposition of the flow field in various clusters that reflects the physical mechanisms of boundary layer separation and vorticity generation. Finally, a XGBoost regressor was trained to correlate the inflow conditions and geometrical parametrization of the leading edges to the cascade deflection. Considering the promising results, it may applied to three-dimensional flow fields for the derivation of novel correlations.

The work of Pan et al. [38] investigated a novel ML model for the determination of aerodynamic losses in boundary layer ingesting fan rotors. The fan under investigation is the NASA Rotor 67, operating in the air compressibility range, although this methodology may be extended to low-speed industrial fans. Training data are provided through means of RANS CFD simulations of a single vane subjected to various intake distortion intensities. The number of samples is 635. The authors have then applied a region-segmentation combinational loss model in which the blade was divided into sub-regions spanwise, based on distinct loss sources, represented through the entropy loss that is chosen as target variable. Local Radial Basis Function (RBF)-based loss models are then trained for each region and eventually aggregated to produce a full-span loss prediction. A region-dependent feature selection is applied, following the physical understanding of the loss mechanism. For example, at the blade tip, the relative Mach number and the blade load coefficient can be exploited to model tip leakage and shock losses. The model results showed a better capability of the ML approach with respect to traditional loss models, demonstrating the feasibility of this zonal approach.

2.2. Centrifugal / Mixed Flow Fans

In the work of Khalkhali et al. (2010) [17], neural networks were adopted to carry out the optimization of a forward-curved blade of a centrifugal fan. No details on the fan characteristics or dimensions are provided. The camber line of the fan blades was approximated using a Bezier curve, reducing the number of design parameters to four: inlet, outlet and stagger blade angles and number of blades. First, the authors carried out 132 RANS simulations to characterize the influence of the design parameters on the head rise and head loss of the fan. The Group Method of Data Handling (GMDH) algorithm is trained on the CFD results, leading to parametric polynomial correlations between design parameters and head rise and loss. GMDH algorithms are self-organizing neural networks that adopts complex polynomial functions as neuron activation and are particularly effective in solving search and selection problems. An evolutionary algorithm is eventually applied exploiting the GMDH-derived fitness functions to derive a Pareto front of the optimal designs. The main advantage of this approach lies in the easy interpretability and fast implementation of the empirical correlation from the GMDH algorithm, although data generation may be computationally intensive for complex geometries (e.g., three-dimensional development of the blade) or large design spaces.

Zhang et al. [39] discretized the backward-blade geometry of a centrifugal fan. After the problem specification in terms of 1D design the feasible blade geometries are parametrized using Bezier curves. To reduce the number of tests, the authors adopt a design of experiment to build a design space that is later investigated through means of Spalart-Allmaras RANS simulations. A NN is then trained to map between input geometrical parameters and blade aerodynamic response. The NN is then embedded in a optimization algorithm to solve a multi-optimization problem targeting total pressure rise and efficiency. A second optimization—this time not based on ML—is carried out to find the optimal width-to-tip ratio. Eventually, the optimized blade shows a 9.5% and 5.5% increase in pressure and efficiency at the highest flow rate.

In [40], the authors focus their analysis on the G4-73 centrifugal fan. The fan is here optimized in terms of blade number, exist stagger angle and impeller outlet width. An orthogonal design is applied to limit the number of experimental tests (49) to generate the training dataset. Then, a NN is trained to predict values of total pressure and efficiency, exploited as surrogate fitness function within an evolutionary algorithm. Experimental measurements of the optimized model show improved aerodynamic and aeroacoustic performance in the range of stable operating flow rates of the fan.

Bamberger et al. [41] extended the previous work carried out on axial fans to impellers, testing also a Local Model Network to reduce overfitting and provide faster inference and better generalization. Approaches like this are particularly useful in engineering problems with a limited amount of data. An important issue in fan-AI applications is in fact the availability of high-quality data, which is often limited by the cost and time, and the limited amount of poor solutions among these data. Most of the fan geometries to improve are in fact corresponding to designs with already good performance and statistical methods like these are prone to overfitting when the number of solutions is limited to a small number of good enough designs. The typical DoE approach is in fact operating according to range and steps to explore the design space and in so doing represents the design space as a grid rather than a continuous space.

An alternative solution to speed up the optimization process was proposed by Meng et al. [42], based on Extreme Learning Machines (ELM), a single-layer feedforward NN that does not require backpropagation for training. Although these algorithms do not lift the burden of the high number of computations required for data generation, they significantly accelerate ML model training and optimization. ELM, coupled with Particle Swarm Optimization (PSO), was applied to a centrifugal fan to determine optimal blade position and curvature.

A high-dimensional optimization process was carried out by Lopez et al. [43], who parametrized a transonic fan applying Active Design Subspaces (ADSs). Design parameters included blade sweep, skew, camber, etc., distributed over five radial sections, leading to a total of 35 degrees of freedom. The latent space obtained from ADS is then adopted as the input vector to a NN, trained on medium-fidelity CFD data and implemented within an optimization algorithm. A comparison with a standard adjoint optimization method proved the higher accuracy of the ADS-based approach.

The work in [16] is among the few dedicated to the analysis of mixed-flow fans. In particular, the authors here aim to improve the fan aerodynamics and efficiency by optimizing the shape of the outlet guide vanes (OGV). According to this scope, they adopt as design variables are the axial spacing between impeller and outlet guide vanes and the radius, length and coordinates of the camber line. The design space is generated through Latin hypercube sampling of the OGV features, that have been preliminary constrained to assure the feasibility of the solutions. It is then investigated through RANS numerical simulations, and the results are treated using a Least Squares–Support Vector Machine (LS-SVM) algorithm that is capable of predicting system performance with a tolerance of 2% error at a specified flow rate. The authors then combine a Genetic Algorithm with a Particle Swarm Optimization (PSO) method to create a robust optimization framework that queries the LS-SVM model to predict the fitness of each individual. The numerical and experimental validation of the approach proves the feasibility of the proposed methodology, by achieving a impressive increase of 106 Pa in total pressure and a 16.3% improvement in total pressure. Although it was not explicitly considered in the optimization loop, a noise reduction of 3.6 dB is also reported, obtained by a suppression of pressure fluctuations induced by wall vortices.

3. Data-Driven Performance Prediction

A different application of ML methods is the prediction of aerodynamic and aeroacoustic performance of fan and fan systems.

In this case, the followed methodologies, as well as the type of learning method adopted, may vary significantly from case to case. Nevertheless, the underlying idea is that most fans may exhibit similar behavior, even with different geometrical and operational parameters. This is not entirely novel, especially considering the invaluable previous works of Balje and Cordier in deriving hidden correlations using similitude theory.

Some researchers have already investigated large collections of fan designs along with their corresponding aerodynamic performance, with the aim of predicting the characteristics of novel designs. Although similar processes have also been carried out in some of the works reported in the previous section, the aim here is not the optimization of fan geometry, but rather the prediction of fan characteristics once all design parameters are known.

A comprehensive review of the state of the art in industrial fans as of the year 2020 can be attributed to Masi et al. [22], who conducted a survey of more than 500 fans based on catalogue data provided by major manufacturers. Despite not being based on unsupervised algorithms, this work demonstrates an excellent application of data-driven and statistical methods capable of revealing the state of the art, design trends and possible future developments in industrial fans.

Two pioneering works by Angelini et al. first introduced a methodology to derive three novel composite parameters for axial fans through the application of dimensionality reduction algorithms (PLS, PCA), analogous to non-dimensional speed and flow rate [44]. These non-dimensional groups combined additional fan design parameters like hub ratio, solidity, blade count, blade twist and aspect ratio. This application effectively enhanced the existing Balje chart, ultimately leading to a powerful preliminary design tool [24].

A more fan-specific approach has been followed by Moradihaji et al. [45], who presented a novel methodology that combines transfer and ensemble learning to predict a radial fan performance starting from a reduced dataset. The underlying algorithm was random forest, coupled with a Gaussian Process Regressor (GPR) to target the probablistic error of the prediction. The main advantage of this approach was that, starting from a reduced number of simulations exploited as training data, it was possible predict the performance of new fans within the original design space.

The same authors [23] tried to model the off-design performance of an industrial fan. The dataset is generated through RANS simulations of the fan at various volumetric flow rates and rotational speeds, leading to 70 samples, with 50 eventually exploited for the algorithm training. Three different ML algorithms are tested based on neural networks, support vectors and random forests respectively. The comparison of their accuracy showed a higher interference capability of the NN-based approach, especially in terms of efficiency. The authors have also tested the extrapolated capability of the NN model by inferring off-design performance of the same fan with an increased blade number, finding a deterioration in the algorithm accuracy as the blade count progressively move away from the original value. Nevertheless, this work suggests the possibility of implementing novel transfer learning methodologies to significantly speed-up design processes.

Fasquet et al. [27] have applied an advanced ML method, based on U-Net NN to predict two-dimensional entropy distribution in the fan wake, as well as efficiency and pressure ratio, providing blade geometry and operating conditions. The training dataset was generated through numerical simulations of 75 different fan geometries. A comparison with a standard surrogate modelling technique, i.e., Proper Orthogonal Decomposition (POD)-Kringing, highlighted the superior performance of the ML model with respect to the latter. This methodology could also be applied for the optimization of the blade design, similarly to the optimization pipelines described in the previous section.

The work of Huang et al. [25] also targeted the prediction of the wake of the fans, which could lead to fast and reliable predictions on the wake-induced noise. With this in mind, the authors have built a training dataset of CFD simulations of four similar fan geometries evaluated at seven rotational speeds and variable inlet mass flow rates, with 268 results. The input data entail fan speed, mass flow rate, trailing edge shape, chord length, local stagger angle and boundary layer thickness on pressure/suction side, counting a total of 13 input features. Model output was built by interpolating the fan wake over a $30 \times 28 \times 100$ sampling points regularly distributed at different axial, radial and circumferential locations respectively. The authors emphasize the role of a proper selection of training/test samples, as high-density-dataset may lead to the introduction of biases in the model training, artificially boosting the model accuracy. They tested different combinations of train/test points, eventually utilizing an entire geometry for the algorithm test. Two different ML frameworks are then evaluated. First, each point of the grid is treated as independent, and a XGboost algorithm and a NN are used to predict mean velocity vectors, turbulent intensity and length scale. The second approach is instead based on a Convolutional Neural Network (CNN) to model the whole wake field in a single step. Results show that XGboost and NN algorithms have a mixed and competing accuracy, with the authors eventually proposing a combination of the two models to achieve reasonable results. The CNN exhibits a similar level of accuracy, with a significant simplification of the framework, comparable computational costs and a more physical representation of the problem. Even if further testing is required, the proposed methodology may entirely bypass the need of numerical simulations once a generalization of the approach is achieved.

Fans are frequently found in larger systems that constraint their operating characteristics and performance. Within this context, a fan-AI agent based on multi-source sensors—not directly related to the machine itself—can identify the actual operating conditions of the fan. The opposite is also true, leading to the identification of the overall system performance based on the analysis of the fan status. A perfect example of the latter is the work of Ahmed et al. [46] where a part of the operational parameters of the fan are used to predict the overall system performance. In this case, the targets included a turbofan engine specific thrust, propulsive efficiency, specific and total fuel consumption. A more comprehensive review of works related to the performance estimation of ventilation systems through ML methods can be found in [47].

Noise emissions in fans has also been investigated using ML approaches. Li et al. [26] were able to predict aerodynamic and aeroacoustic performance of low-pressure fans with and without leading-edge serrations. The numerical dataset, generated using a factorial design, defined 130 target values also expressed in the time and frequency domains. The data include multi-source observations, experiments and RANS simulations, and spanned over 545 cases of 8 fan geometries. NNs were trained to predict each target. The described approach is also feasible to be embedded in multi-objective optimization of fans.

In electrical traction systems, the fan can contribute up to the 75% of the overall noise emissions. In [48] the authors have measured the A-weighted sound pressure level of a centrifugal fan characterized by different operating (rotational speed and flow rates) and geometrical (number of blades) parameters. Unfortunately, due to the complexity of aeroacoustics measurements, the authors have investigated a limited combination of these factors, with a final training size of 21 samples. In the methodology, multiple regression equations (MGAs) are originally built to map the fan parameters to the noise emissions at various frequencies. The derived semi-empirical correlations are further statistically analyzed to select the relevant features, eventually limited to the blade number, static pressure rise, flow rate and rotating speed. The proposed multi-regression analysis is combined with a two-layer fuzzy neural network to predict the noise spectrum. In the first stage, noise emissions at different frequency bands are used as input regressors in multiple linear regression models to extract intermediate variables that describe local spectral relationships. A selection mechanism is then applied to retain the most informative regression outputs, reducing redundancy among frequency components. The selected regressors are provided to the fuzzy neural network, where fuzzy hidden layers model nonlinear interactions and uncertainties across frequency bands. Finally, the output layer performs multi-output regression to reconstruct the complete noise spectrum, enabling accurate frequency-domain noise prediction. The framework, compared with other ML models, shows a fairly high accuracy—more than 93%—even if the small considered dataset may lead to a poor model generalization capability. Nevertheless, the MGA-FNN approach can be a power tool to support optimization and fault diagnosis, thanks to the high non-linearity and embedded uncertainty treatment.

Tieghi et al. [49] presented one of the few applications of unsupervised methods for the identification and investigation of noise sources in a benchmark axial fan. This paper investigates the use of unsupervised machine-learning clustering techniques to identify and characterize noise generation mechanisms in a low-speed axial fan. High-fidelity Large Eddy Simulations (LESs) are performed and coupled with a perturbed convective wave equation (PCWE) formulation to compute dipole noise sources in the flow field. The resulting dataset includes instantaneous flow variables, pressure and velocity gradients, turbulence quantities and PCWE-based acoustic source terms. Given the extremely large size of CFD-generated datasets, the authors propose a data-reduction strategy based on sampling flow features on two cylindrical surfaces located at 50% (midspan) and 95% (near-tip) blade span, representing reference flow behavior and tip-leakage-driven noise mechanisms, respectively. Prior to clustering, the data are normalized and analyzed using correlation analysis, Principal Component Analysis (PCA) and t-SNE, showing that global dimensionality reduction is ineffective when applied to the full 3D dataset but becomes informative when restricted to spanwise surfaces Three families of clustering algorithms are evaluated: k-means (partitional), Gaussian Mixture Models (distribution-based) and HDBSCAN (density-based). Their ability to reconstruct the spatial distribution of PCWE noise sources is quantitatively assessed using an image similarity metric. Among the tested methods, Gaussian Mixture clustering demonstrates the highest and most consistent similarity (>80%) with the computed noise source distributions at both midspan and tip regions, outperforming density-based approaches, which struggle with data agglomeration and scalability issues.

The statistical analysis of the resulting clusters reveals that noise source intensity is strongly linked to the statistical distributions of pressure and velocity gradients, with distinct mechanisms dominating at midspan and near the blade tip. In particular, the results highlight the role of spanwise work distribution and tip-leakage-related gradients in driving noise generation. The study concludes that clustering-based unsupervised learning is a promising tool for uncovering hidden aeroacoustic patterns in turbomachinery flows, although computational cost and scalability remain key challenges for industrial applications.

4. AI-Driven Control

This section reports a number of case studies on system control and off-design operations using ML. These applications are the only ones that demonstrate practical utilization of reinforcement learning, which is particularly effective for the creation of smart controllers.

In [29], ML is applied for the prediction of operating regimes in contra-rotating axial fans. This work adopts a novel coupling of unsupervised learning through k-Means clustering and NNs. Here, the underlying assumption is that pressure fields, analyzed through Fourier representation, exhibit particular patterns as the fans are approaching stall. To support the algorithm training, the authors have performed experimental measurements of the fan with high-frequency data acquisition, gradually reducing the mass flow rate down to the stall region. Following a statistical analysis of data, the input features are eventually constituted by the variance of Fourier amplitude, Shannon's entropies and a coefficient of determination for a polynomial fit of fixed degree. A K-means algorithm is then used to classify the originally unlabeled dataset, leading to three different clusters in the collected time series of measurements. A NN is finally trained to classify, based on the same input features as the clustering model, what is the current state of operation of the fan. Although this methodology—validated using five independent subsets — is here applied to a single fan setup, it may lead to significant progress in fan stall control, for example to drive the decisions of controllers based on semi-reinforced learning.

Peng et al. [30] targeted the monitoring of fans used for the cooling of server cabinets. To achieve an accurate evaluation of the fan thermal performance and power absorption it is in fact compulsory to decouple during acquisition the real rotational speed signal to the associated noise. To address this problem, the authors have implemented and modified a Unscented Kalman Filter (UKF) that considers state-dependent noise in velocity observations and a mathematical discretization of the associated uncertainties. The experimental validation of the framework shows the denoising capabilities of the proposed algorithm. This pioneering work proves that it is possible to include model, operational and installation uncertainties—commonly found in standard fan operations—within the ML framework, constituting a use case of real-time AI-assisted monitoring and control systems.

The work of Iranfar et al. [33] both target the thermal efficiency of fan-driven systems. In the former, dynamic thermal management of a processor unit is carried out by controlling fan speed using a reinforcement learning-based method, resulting in a significant reduction in cooling power consumption. During operations, the ML agent is capable of selecting the number of cores, operating frequency and fan speed—and their temporal evolution—to maximize thermal efficiency and minimize power absorption of the CPU unit while satisfying the processor workload. A notable reduction in fan power absorption, between 27 and 40 %, is achieved.

Corsini et al. [32] treated the interaction between the ventilation system and a complex industrial environment—a gas turbine enclosure. They built a physical transfer learning framework that, starting from a set of simplified set of methane purge simulations in simplified geometries, is able to detect poorly ventilated areas in the actual enclosure. Their approach has been validated in [50], showing a good agreement between computationally expensive URANS simulations and the fast ML predictions. Despite this being a design for safety tool, it may be coupled with a fan controller to optimize the ventilation efficiency in presence of fuel leaks.

5. ML-Based Fault Analysis & Prognosis

Fault diagnostics through data-driven methods is not entirely novel in the turbomachinery sector. For example, SCADA data from wind turbines is often analyzed using ML methods, with the aim of predicting remaining useful life or component damage [51,52]. This class of applications is probably the least developed, as it requires a large collection of time-series data and relies on online monitoring systems, which are not frequently found in fan installations.

Tao et al. [31] carried out fault diagnosis of axial flow induced draft fans. The training dataset is constituted by automatic acquisitions of the control system of a power plant, with one minute sampling interval over a span of three months. Temporal sequences of vibrations of the drive system, bearing temperatures, flow rate measurements and fan pressure rise are used as input features. During this period, five major fault occurred caused by different components. Due to the relative low number of observations of faults, data are augmented using SMOTE algorithm. Three different boosting algorithms, decisions trees, Random Forests (RFs) and Adaboost, are then trained to categorize the faulting components. The RF-based classifier is eventually found as the most effective, with above 90% of precision and recall.

A real-time anomaly detection methodology is presented in [28]. The acquired data consist of vibration and current signals collected from accelerometers and electrical sensors installed on five small electric fans used for equipment cooling. In addition to nominal operating conditions, the fans are tested under multiple faulty states, including broken or stuck blades, as well as unstable and highly vibrating operating regimes. The input signals are first processed using a moving average to extract time-domain features and subsequently analyzed through Fast Fourier Transformation (FFT) techniques to compute frequency-domain quantities such as spectral power, RMS, skewness, and kurtosis, resulting in 69 time-domain and 66 frequency-domain features. The resulting feature space is then analyzed using a combination of CNNs and deep neural networks to classify the fan operating condition into one of six predefined states. The proposed framework achieves a final precision and accuracy exceeding 95%. Owing to the high processing capability

https://doi.org/10.3390/ijtpp11020028

of the CNN-based architecture, the methodology can be extended to other types of electrical equipment and fan systems and represents a promising solution for operation and maintenance activities in industrial environments.

The work of Zhang et al. [34] focuses on the blade damage accumulation. The authors have performed measurements at various rotation speeds of a centrifugal fan with blades exhibiting 10 and 20 mm cracks penetrating the whole blade thickness. Acoustic and vibrational signals are stacked to generate the input features. A one-dimensional CNN eventually classifies the state of the blade between normal and cracked. The model shows an excellent accuracy, with the most notable errors—although still below a 2%—occurring between the 10 mm cracks and undamaged blades. This works not only demonstrates the capability of ML models to treat multi-source data but also proves how heterogeneous samples may also lead to an enhanced algorithm predicting capability.

In [35], the problem of blade fouling in air-movement domestic fans is treated by building a ML algorithm, based on pre-trained CNN models, to identify the operating conditions of the fan. Dust accumulation, in fact, leads to a peculiar acoustic signatures that non-linear algorithms may learn to identify, activating the adequate maintenance operations. To this extent the authors performed experimental measurements of the acoustic emissions of the fan in an anechoic camber. The acoustic signals were later mapped to a three-dimensional space, i.e., spectogram, with the frequency, time and intensity constituting the Cartesian dimensions. With this projection, acoustic signals can be interpreted as images with the intensity encoded as its brightness. Data augmentation is then performed, following the standard image transformations (Figure 5), e.g., cropping, rotation, distortion, etc., reducing the burden of large experimental campaigns. A CNN classifier is eventually built on this dataset, predicting the fan operating conditions—clean or dusted blades.

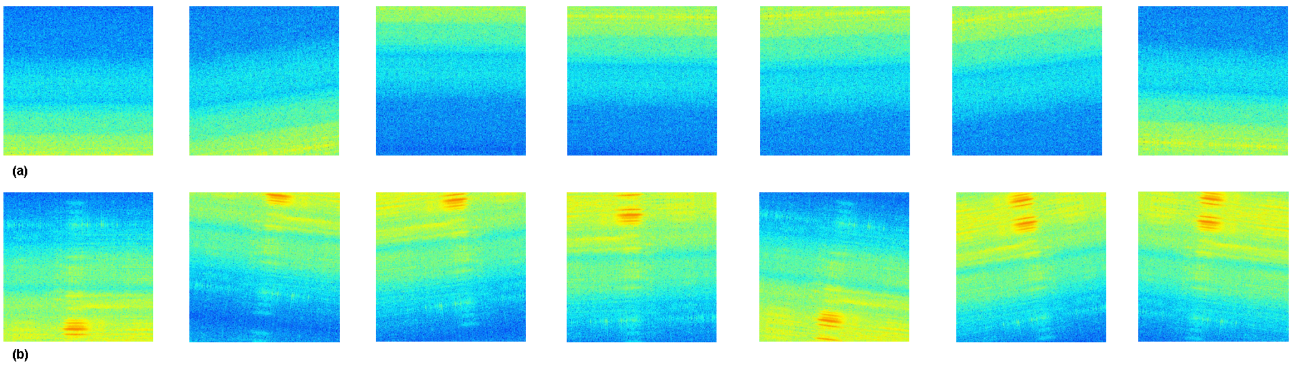

Figure 5. Data augmentation through cropping, rotation, distortion, performed on the spectogram projection of acoustic signals of a domestic fan in normal (**a**) and fault (**b**) conditions, from [35].

The goal was to determine whether fouling had occurred on the blades, achieving a significant 95% accuracy. This pioneering work is of particular interest, as the methodology is based on a non-intrusive and simplified measurement system that could facilitate maintenance operations in larger plants.

6. Critical Analysis and Current Trends

The structured comparison presented in Table 1 allows several cross-cutting observations on the maturity and generalization capability of Fan-AI applications.

Table 1. Summary of datasets and models from the reviewed Fan-AI literature. Combined column shows number of input and output features and database population. A dash (–) denotes information not explicitly reported.

Ref.	Fan Type	Data Source	$N_{in}/N_{out}/N_s$	ML Model (s)	Reported Performance	Perf. Assessment
[15]	Axial	Exp.	11/2/19	DNN	–	Train/test split
[14]	Axial	RANS CFD	10/2/4000+	NN	90% cost reduction vs. direct CFD	CFD validation
[21]	Aeroengine	RANS CFD (LHS)	–/3/–	XGBoost	Shapley-based uncertainty	Train/test split; Shapley
[19]	Axial	RANS CFD	–/2/–	DNN (Dendrite Net)	+2.47% eff.; +5.55% Δp_t	Exp. validation
[20]	Axial	Exp.	4/4/65	NN	Good generalization	Exp. validation
[18]	Axial (rev.)	XFoil	6/2/–	LSM; ANN	50% time reduction	Pareto comparison
[36]	Axial cascade	RANS CFD (CCD)	4/4/–	LSM + ANOVA	Overprediction reduced to ∼15%	3D CFD comparison
[37]	Axial cascade	RANS CFD	–/1/76	PCA + GMM	–	Visual consistency
[38]	Aeroengine (R67)	RANS CFD	Region-dep./1/635	RBF	Superior to loss models	Loss-model comparison
[17]	Centrifugal	RANS CFD	4/2/132	GMDH	–	Train/test split
[39]	Centrifugal	RANS CFD (DOE)	–/2/–	NN	+9.5% Δp	CFD validation
[40]	Centrifugal	Exp.	3/2/49	NN	Improved aero-acoustics	Exp. validation
[41]	Centrifugal	RANS CFD	–/–/–	Local Model Net	Reduced overfitting	CFD validation
[42]	Centrifugal	RANS CFD	–/–/–	ELM	Faster training	–
[43]	Transonic	Med.-CFD	35/–/–	ADS + NN	Better than adjoint	Adjoint comparison
[16]	Mixed-flow	RANS CFD	4/–/–	LS-SVM	+106 Pa; +16.3% eff.	Exp. + CFD validation
[22]	Axial	Catalogue	–/–/500+	Statistical	Mapping state-of-art	Statistical analysis
[24,44]	Axial	Catalogue	7/3/500+	PLS; PCA	Enhanced Balje chart	Cross-validation
[45]	Radial	RANS CFD	–/–/Reduced	RF + GPR	Good extrapolation	Train/test split
[23]	Radial	RANS CFD	2/–/70	NN; SVM; RF	NN best for eff.	Extrapolation test
[27]	Axial	RANS CFD	–/3/75	U-Net	Superior to POD-Kriging	Method comparison
[25]	Axial (AE)	RANS CFD	13/Wake/268	CNN; XGB	CNN most physical	Leave-one-out
[46]	Turbofan	–	–/4/–	NN	–	–
[26]	Axial (LP)	Exp. + RANS	–/130/545	NN	–	Train/test split
[48]	Centrifugal	Exp.	4/Noise/21	MGA-FNN	>93% accuracy	Train/test split
[49]	Axial	LES + PCWE	Flow/Noise/1	GMM; k-Means	>80% similarity	Image similarity
[29]	Axial	Exp.	3/3 classes/–	k-Means + NN	–	5-fold CV
[30]	Cooling fan	Exp.	–/1/–	Modified UKF	Effective denoising	Exp. validation
[33]	CPU fan	Simulation	States/Actions/–	RL	27–40% power reduction	Simulation benchmark
[32,50]	Ventilation	URANS	CFD/Eff./–	Transfer Learning	Good URANS agreement	CFD validation
[31]	Axial	SCADA	4+/5 classes/–	RF; DT	>90% precision	Precision/Recall
[28]	Small fan	Exp.	135/6 classes/–	CNN + DNN	>95% accuracy	Accuracy
[34]	Centrifugal	Exp.	–/–/–	1D-CNN	<2% error	Confusion matrix
[35]	Domestic fan	Exp.	Spectrogram/–/–	Pre-trained CNN	95% accuracy	Accuracy

Data sources and sample sizes. A first evident trend is the strong reliance on numerically-generated data, with the vast majority of the reviewed studies employing RANS-CFD simulations to populate the training set. Only a minority of works rely on experimental measurements or catalogue data. Sample sizes vary by more than two orders of magnitude, from as few as 19 [15] or 21 [48] to over 4000 [14]. This heterogeneity has direct consequences on model robustness: small training sets inevitably restrict the generalization envelope of the resulting models, while large-scale CFD campaigns—although computationally expensive—enable a broader exploration of the design space.

Input/output dimensionality. The number of input features spans a wide range, from 2 to 4 parameters in low-dimensional optimization studies to 35 degrees of freedom in the active-subspace framework of Lopez et al. [43] and the 135 features extracted in the time–frequency domain by Cong et al. [28]. Target variables are also diverse: scalar performance metrics (efficiency, pressure rise) dominate in the design and optimization category, whereas spatially-resolved fields (e.g., wake velocity distributions, noise source maps) appear in the more recent performance-prediction works. Notably, several studies do not explicitly report the number of inputs and outputs, which complicates a fair comparison and points to a lack of standardization in reporting practices.

Model selection. Neural networks, in their various architectures (feedforward, convolutional, dendrite-based, U-Net), remain the dominant modelling choice across all categories. Ensemble and boosting methods (RF, XGBoost, AdaBoost) are gaining traction, particularly in fault diagnostics and in recent optimization studies that emphasize interpretability. Unsupervised techniques (PCA, GMM, k-Means, HDBSCAN) are confined to a supporting role—either for dimensionality reduction or for the identification of flow patterns—and are rarely used as the primary predictive tool. The emerging adoption of reinforcement learning and transfer learning in the control and ventilation domains signals a shift from purely regression-oriented frameworks toward more adaptive paradigms, although evidence is still limited to isolated case studies.

Performance assessment. The methods used to evaluate model accuracy vary significantly. A non-negligible fraction of the reviewed studies relies on simple train/test splits without reporting metrics such as R^2, RMSE, or classification scores (precision, recall, F1). Where quantitative figures are available, they are expressed in heterogeneous and application-specific terms—percentage improvements in efficiency, Pareto-front comparisons, image similarity scores, confusion matrices—making a direct cross-study benchmarking virtually impossible. Only a few works employ rigorous protocols such as k-fold cross-validation [29] or leave-one-geometry-out testing [25]. The absence of a shared benchmark dataset further exacerbates this limitation and constitutes one of the main barriers to assessing the true generalization capability of the proposed models.

Maturity across categories. The design and optimization category is clearly the most developed, featuring the largest number of studies, the most diverse algorithmic landscape, and several instances of experimental validation of the optimized designs. In contrast, the performance-prediction category, while growing, still faces challenges related to the transferability of models across different fan families. The AI-driven control and fault-diagnosis categories remain at an early stage: control applications are limited to specific thermal-management or stall-detection scenarios, and fault-prognosis studies are constrained by the scarcity of labeled failure data. In both cases, the practical deployment in industrial environments—where robustness, real-time inference and interpretability are paramount—is yet to be demonstrated at scale.

Methodological Gaps and Limitations

A closer inspection of the reviewed literature also allows us to address a set of fundamental questions on the nature of the machine-learning problems encountered in fan engineering, which are discussed below.

Smoothness of input–output mappings. In many of the reviewed design and optimization studies, the input–output relationships are indeed expected to be relatively smooth, particularly in the vicinity of the design point. Aerodynamic performance metrics such as total pressure rise and efficiency are governed by the Navier–Stokes equations, which—for attached, subsonic flows typical of industrial fans—produce continuous and differentiable responses to moderate variations of geometrical and operational parameters. This inherent regularity partly explains the widespread success of low-complexity surrogate models (e.g., polynomial response surfaces, radial basis functions) reported in works such as [14,17,36]. However, smoothness cannot be taken for granted across the entire operating envelope. Near stall, in regions of massive flow separation, or in the presence of strong tip-leakage vortices, the mapping becomes highly nonlinear and may exhibit steep gradients or even discontinuities in the performance space. Similar considerations apply to aeroacoustic targets, where broadband noise sources are intrinsically linked to turbulent, unsteady phenomena that are poorly captured by steady-state descriptors. In these regimes, simple surrogate models are prone to large interpolation errors, and more expressive architectures—deep neural networks, convolutional networks, or ensemble methods—become necessary, as demonstrated by [21,25,49]. Consequently, the degree of smoothness is not a universal property of Fan-AI problems but depends strongly on the operating regime, the choice of target variables, and the fidelity of the underlying data.

Structure of the datasets. The training datasets employed in the reviewed works are, in most cases, highly structured by construction. The vast majority of studies rely on CFD-generated data produced through formal sampling strategies—Design of Experiments (DOE), Latin Hypercube Sampling (LHS), Central Composite Design (CCD), or orthogonal arrays—that impose a regular and controlled coverage of the parameter space [14,16,21,40]. This structured nature is further reinforced by the physics itself: non-dimensional similarity laws constrain fan performance to well-defined manifolds in the feature space (e.g., the Cordier diagram), which effectively reduces the intrinsic dimensionality of the problem. As a result, the machine-learning models reviewed here typically operate on low-dimensional, physics-informed latent spaces rather than on raw high-dimensional data. While this structural regularity generally facilitates model training and contributes to the favourable accuracy figures reported, it also introduces a caveat: the apparent success of a model may partly reflect the regularity of the sampling grid rather than its true ability to capture the underlying physics. Works that adopt catalogue or experimental data—inherently less structured and potentially affected by noise and heterogeneous measurement conditions—constitute notable exceptions [22,28,31,44] and tend to require more robust algorithmic choices (ensemble methods, data augmentation) to achieve comparable accuracy.

Interpolation-dominated validation. A critical observation emerging from Table 1 is that model validation is overwhelmingly performed within the training design space, i.e., in an interpolation setting. The most common protocol is a random train/test split drawn from the same parameter space—and, frequently, from the same fan geometry. Only a handful of studies explicitly assess extrapolation capability: Moradihaji et al. [23] tested the prediction of off-design performance for a fan with increased blade count, observing a progressive deterioration of accuracy; Huang et al. [25] adopted a leave-one-geometry-out protocol, which constitutes a more rigorous test of generalization across configurations; and the bi-level framework of Angelini et al. [18] periodically corrected surrogate predictions with high-fidelity evaluations, implicitly acknowledging the risks of pure surrogate-based extrapolation. The general absence of systematic extrapolation tests represents a significant gap in the current literature and limits the confidence with which the reviewed models can be transferred to novel fan designs outside their original training domain. This issue is particularly relevant for industrial applications, where a model trained on a specific fan family is expected to provide reliable predictions also for new, untested configurations.

Overfitting mitigation strategies. The approaches to overfitting mitigation reported in the reviewed works are heterogeneous and, in several cases, not explicitly discussed. Among the more rigorous practices, the following can be identified:

- *Architectural choices.* Bamberger et al. [41] adopted Local Model Networks specifically to reduce overfitting in data-scarce centrifugal fan optimization, while Meng et al. [42] employed Extreme Learning Machines, whose single-layer structure inherently limits model complexity. The dynamic XGBoost framework of Cheng et al. [21] included regularization at multiple stages of the pipeline.
- *Bi-level and hybrid strategies.* Angelini et al. [18] demonstrated that single-level surrogate optimization can produce unreliable Pareto fronts due to progressive divergence between surrogate predictions and true values, and proposed a bi-level correction scheme that periodically re-evaluates the surrogate against high-fidelity solvers. This constitutes one of the most explicit treatments of overfitting in the surveyed literature.
- *Feature selection and dimensionality reduction.* Several works employ PCA, ANOVA-based variable selection, or Active Design Subspaces to reduce the effective dimensionality of the input space prior to model training [36,37,43], thereby limiting the degrees of freedom available for overfitting.
- *Data augmentation.* In the fault-diagnosis category, where labeled data are inherently scarce, SMOTE-based oversampling [31] and image-domain augmentation (rotation, cropping, distortion) [35] have been adopted to artificially expand the training set.
- *Cross-validation.* Rigorous validation protocols such as k-fold cross-validation [29] or leave-one-geometry-out testing [25] are employed in only a minority of the reviewed studies.

Overall, the treatment of overfitting remains largely ad hoc and study-specific. A systematic comparison of regularization strategies across different Fan-AI tasks—and, more broadly, the definition of best practices for model complexity selection in turbomachinery-ML—is still missing and would represent a valuable contribution to the field.

7. Open Challenges

This literature survey clearly highlights how powerful and helpful tools are gradually being developed by industrial and academic stakeholders, with several successful results on different aspects of the life cycle of fans. Nevertheless, very few applications can be considered as technological mature or ready to be transferred into current workflows. This section aims to discuss some of the limiting factors to the integration of ML in the fan practice, with a critical view on possible solutions. The starting point for the discussion is the results shown in Figure 6, that surveys the audience perspective toward the biggest challenges.

Results are not surprising. Data gathering is a crucial step in every ML application and the statistical background of machine learning approaches inherently relies on a wide number of results to leverage the potential of most methods. The amount of data openly available is almost non-existent and manufacturers or other stakeholders with access to experiments, numerical simulations and design parameters often are not willing to spend time on R&D projects on ML applications. This can be partially explained by a lack of expertise from engineers (both at senior and junior level), but often entails the lack of an organized and structured database with results from previous works.

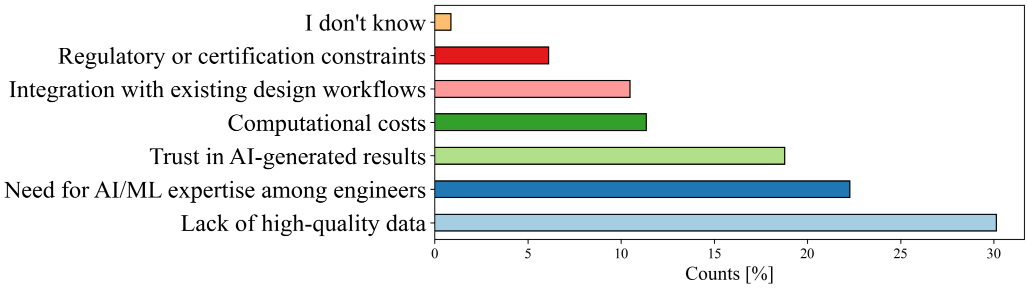

Figure 6. Survey answer to *"What are the biggest challenges in adopting AI/ML for fan applications?"*—multiple answers were allowed.

The reliability of results from AI is also an issue, with several effective proof of concepts and applications, but still far from industrial technological readiness. The statistical methodology behind ML is inherently prone to some errors and the preliminary applications developed so far still lack serious stress tests. The turbomachinery industry, at all levels, is in fact reliant of a strong quality assessment that has been stratified and certified for decades. The world of software is still well behind this level of certification and often consider as mere "bugs" errors that in a fan product would impact the credibility (and revenues) of the manufacturer. Data scientists and AI experts often do not understand this inherent difference between the software and hardware market, and training of young engineers to deal with the application of ML to product development is not an easy task, especially given the salary differences between the two jobs.

Computational costs are no longer seen as a strongly limiting factor, considering the relative affordability of GPU units and the availability of novel HPC and cloud GPU infrastructures. In addition, the development of novel machine learning models is increasingly oriented toward more efficient, pre-trained architectures that better capitalize on the computational costs incurred during the training phase. In this context, one may question how effectively a general pre-trained model can perform when applied to a specific engineering problem; however, this discussion lies beyond the scope of the present work. As an example, Claude Sonnet [53] is seen at the time of writing as the best LLM-driven code-assistant on the market. However, unlike other general scope, LLMs, it was specifically tuned for this scope and a large number of potential users dealing with code development. It is highly unlikely that a software company would do the same to a specific engineering task like turbomachinery development.

Normative and regulative aspects can also constitute a future issue, especially when dealing with industrial applications. If shared and open AI models are developed, then the individual contribution of the stakeholder is reduced, and thus is revenue, possibly leading to conflicts or coincident fan designs. The normative aspect is however entirely subject to speculations, given that novelty of the topic and challenges.

(I) Creation of Shared High-fidelity Dataset

One of the most influential factors that impact the quality of ML models is surely the training data. Considering that the maximum reachable accuracy from current ML methods is the same as that of the training data, improving their reliability immediately results in a general increase in the robustness of AI-driven applications. In a two-fold view, it also means that improving model accuracy does not automatically result in better modeling of the physical solutions under analysis, but rather in a closer approximation of the original data. For example, if the training data include errors—as may occur in numerical computations—the turbomachinery-ML models will also be trained to reproduce these errors.

From this point of view, turbomachinery-ML is partially in contrast with other big-data disciplines, e.g., the analysis of social media, where local errors may be mitigated by the large amount of available data. So far, in all the analyzed works, the training data were generated in-house, following best practices for simulations/experiments. Frequently the source of data comes from historical collection of results, carried out in a mixed-reliable approach, like it may occur through the collection of years of product development in a company. These aspect alone raises several concerns, since unless novel datasets are generated each time, the mixture of results may strongly impair model performance. Even if best practices and guidelines guarantee a certain level of reliability in the methods adopted for data generation, it is still impossible to perform a fair comparison between different algorithms and approaches to the problem, since each model operates in the data space it was trained on.

These aspects advocate the creation of common repositories and database to train, compare and develop algorithms specifically designed for fans applications, similarly to what is already occurring in other disciplines, as already happened for example for SCADA data in wind turbines or CFD benchmarks. This open-source collection of data should exists as a shared effort from academic and industrial partners, to create a common ground to move forward the new generation of AI algorithms. It must also be underlined that this is surely not an easy task, due to the evident industrial relevance of designs and fan performance potentially slowing down this process.

(II) Standardization Procedures and Methods

Data are characterized not only by their quantity but also by attributes such as accuracy, observables, frequency, source, preprocessing methodology and more. For example, in the applications previously discussed, source data are generated using various sources, from catalogue data to experimental/numerical campaigns, each operating at different levels of detail or with varying observed quantities.

Although data science provides rules and methodologies that are well suited for developing machine learning (ML) algorithms, the data typically encountered by experts in this domain present some unique characteristics. In particular, certain physical information must also be retained in turbomachinery-AI models, similar to what is already occurring in physics-informed models in fluid mechanics.

Methods commonly used in the turbomachinery field—such as projection into cylindrical coordinates or non-dimensional analysis—can be integrated into ML practices, helping bridge the gap between the two disciplines. Moreover, outliers may not be as detrimental here as they are in other fields, as they often correspond to relevant physical phenomena. For instance, in the modeling of boundary layer losses, the region with high velocity gradients would appear as a statistical outlier compared to the core flow, yet it carries essential physical meaning.

Guided by these underlying data patterns, turbomachinery-ML experts follow a process partially informed by heuristics inherited from data science but enriched by engineering experience and domain-specific creativity. Nevertheless, there is a need for standardization in terms of optimal models, preprocessing methodologies, and algorithm performance. Such standardization would not only establish a shared state-of-the-art but also enhance confidence in the reliability and robustness of these methods.

(III) Training of hybrid AI/Turbomachinery experts

As previously highlighted, ML-TM requires a set of specific knowledge that spans over a wide range of topics, from statistics, programming, data science to turbomachinery design, validation and verification. This mix of skills and knowledge falls far from the traditional approach to engineering and turbomachinery field, requiring instead specialized training

https://doi.org/10.3390/ijtpp11020028

and studies. Although we are experiencing an increasing number of courses aimed to higher levels of formation as specialized courses, workshops and seminar, also undergraduate and master courses should be targeted, with the aim of creating novel skilled engineers specifically focused on ML-TM. This could also greatly incentive the technological and knowledge transfer between computer/data scientists and turbomachinery experts in large-scale industrial applications of the methods, as advocated from the industry experts (Figure 6).

8. Conclusions

In this manuscript, we have investigated, through a comprehensive literature review, the state-of-art of machine learning applied to fans. The applications have been categorized based on their scopes in design and optimization, data-driven performance prediction, AI-driven control and fault analysis/prognosis. Among them, the biggest technological development can be found in the ML-assisted optimization, frequently coupled with evolutionary algorithm. Although NN-based method remain the most popular, other advanced ML algorithms are gradually being adopted and developed.

Despite this review of successful works on the topic, there is still a large opportunity to develop and improve existing methods, in particular for applications where the status of maturity is still low—as in the case of prognosis and control. To this scopes, the expertise derived in other turbomachinery relevant applications, e.g., pumps, compressors, may come in aid and easily transferred.

A survey carried out during Fan 2025 conference to an audience of experts has highlighted some of the limiting factors, challenges and opportunities that ML-TM is facing. In particular, although the general view is the future of a AI-fueled turbomachinery, several concerns on the reliability of this approach still exist. Nevertheless as shown in Figure 7, more proven cases studies, standardization procedures and education of novel hybrid ML/TM experts may greatly speed up the process leading to the full exploitation of data-driven methods.

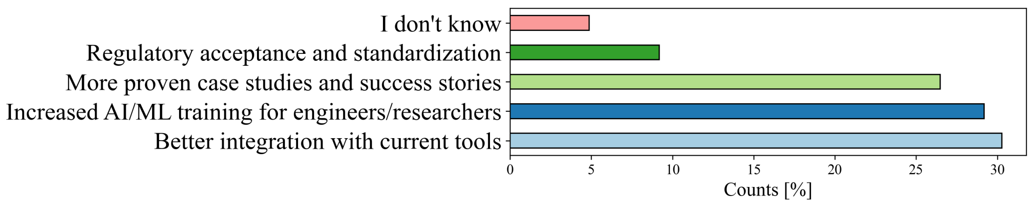

Figure 7. Survey answer to *"What would encourage you or your institution/company to adopt AI/ML in current workflow?"*—multiple answers were allowed.

Author Contributions: Conceptualization, L.T., G.D. and L.B.; methodology, L.T. and G.D.; software, L.T. and G.D.; formal analysis, L.T. and G.D.; investigation, L.T.; data curation, L.T.; writing, L.T., G.D. and L.B.; visualization, G.D. and L.T.; supervision, L.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Data Availability Statement: The data used in this research are available from the authors upon reasonable request.

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

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
ADS	Active Design Subspaces
ANOVA	Analysis of Variance
CFD	Computational Fluid Dynamics
CNN	Convolutional Neural Network
DNN	Dendrite Neural Network
DOE	Design of Experiments
ELM	Extreme Learning Machine
FFT	Fast Fourier Transformation
GA	Genetic Algorithm
GMDH	Group Method of Data Handling
GPR	Gaussian Process Regressor
GPU	Graphics Processing Unit
HPC	High-Performance Computing
IoT	Internet of Things
LSM	Least Squares Method
LS-SVM	Least Squares Support Vector Machine
ML	Machine Learning
NN	Neural Network
NSGA-II	Non-dominated Sorting Genetic Algorithm II
O&M	Operation and Maintenance
PCA	Principal Component Analysis
PLS	Partial Least Squares
POD	Proper Orthogonal Decomposition
PSO	Particle Swarm Optimization
RANS	Reynolds-Averaged Navier–Stokes
RBF	Radial Basis Function
RF	Random Forest
SCADA	Supervisory Control and Data Acquisition
TM	Turbomachinery
TKE	Turbulent Kinetic Energy
UKF	Unscented Kalman Filter
XGBoost	Extreme Gradient Boosting

References

1. Hammond, J.; Pepper, N.; Montomoli, F.; Michelassi, V. Machine learning methods in CFD for turbomachinery: A review. *Int. J. Turbomach. Propuls. Power* **2022**, *7*, 16.
2. Zou, Z.; Xu, P.; Chen, Y.; Yao, L.; Fu, C. Application of artificial intelligence in turbomachinery aerodynamics: Progresses and challenges. *Artif. Intell. Rev.* **2024**, *57*, 222.
3. de Castro-Cros, M.; Velasco, M.; Angulo, C. Machine-learning-based condition assessment of gas turbines—A review. *Energies* **2021**, *14*, 8468.
4. Rai, R.; Tiwari, M.K.; Ivanov, D.; Dolgui, A. Machine learning in manufacturing and industry 4.0 applications. *Int. J. Prod. Res.* **2021**, *59*, 4773–4778.
5. Gabsi, A.E.H. Integrating artificial intelligence in industry 4.0: Insights, challenges, and future prospects—A literature review. *Ann. Oper. Res.* **2024**. https://doi.org/10.1007/s10479-024-06012-6.
6. Leng, J.; Zhu, X.; Huang, Z.; Li, X.; Zheng, P.; Zhou, X.; Mourtzis, D.; Wang, B.; Qi, Q.; Shao, H.; et al. Unlocking the power of industrial artificial intelligence towards Industry 5.0: Insights, pathways, and challenges. *J. Manuf. Syst.* **2024**, *73*, 349–363.
7. Mehdi, G.; Roshchin, M.; Runkler, T. Internet of Turbines: An outlook on smart diagnostics. In Proceedings of the Annual Conference of Prognostics and Health Management Society, Dallas, TX, USA, 19–21 June 2017; pp. 1–7.
8. *FAN 2025, 8th International Conference on Fan Noise, Aerodynamics, Applications, and Systems*; Sigel Press: Antibes/Juan-les-Pins, France, 2024; ISBN 978-1-905941-32-2.

9. Pakatchian, M.R.; Ziamolki, A.; Alhuyi Nazari, M. Applications of machine learning approaches in aerodynamic aspects of axial flow compressors: A review. *Front. Energy Res.* **2023**, *11*, 1135055.
10. Lv, Q.; Yu, X.; Ma, H.; Ye, J.; Wu, W.; Wang, X. Applications of machine learning to reciprocating compressor fault diagnosis: A review. *Processes* **2021**, *9*, 909.
11. Dutta, N.; Kaliannan, P.; Subramaniam, U. Application of machine learning algorithm for anomaly detection for industrial pumps. In *Machine Learning Algorithms for Industrial Applications*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 237–263.
12. Xu, Y.; Gan, X.; Pei, J.; Wang, W.; Chen, J.; Yuan, S. Applications of artificial intelligence and computational intelligence in hydraulic optimization of centrifugal pumps: A comprehensive review. *Eng. Appl. Comput. Fluid Mech.* **2025**, *19*, 2474675.
13. Fausing Olesen, J.; Shaker, H.R. Predictive maintenance for pump systems and thermal power plants: State-of-the-art review, trends and challenges. *Sensors* **2020**, *20*, 2425.
14. Bamberger, K.; Carolus, T. Development, application, and validation of a quick optimization method for the class of axial fans. *J. Turbomach.* **2017**, *139*, 111001.
15. Liu, Y.L.; Nisa, E.C.; Kuan, Y.D.; Luo, W.J.; Feng, C.C. Combining deep neural network with genetic algorithm for axial flow fan design and development. *Processes* **2023**, *11*, 122.
16. Zhou, S.; Luo, Y.; Mao, Z.; Lu, L.; Feng, W. Machine-learning and CFD based optimization and comprehensive experimental study on diagonal flow fan for energy conservation and efficiency enhancement. *Eng. Appl. Comput. Fluid Mech.* **2024**, *18*, 2310608.
17. Khalkhali, A.; Farajpoor, M.; Safikhani, H. Modeling and multi-objective optimization of forward-curved blade centrifugal fans using CFD and neural networks. *Trans. Can. Soc. Mech. Eng.* **2011**, *35*, 63–79.
18. Angelini, G.; Bonanni, T.; Corsini, A.; Delibra, G.; Tieghi, L.; Volponi, D. On surrogate-based optimization of truly reversible blade profiles for axial fans. *Designs* **2018**, *2*, 19.
19. Ding, Y.; Wang, J.; Jiang, B.; Li, Z.; Xiao, Q.; Wu, L.; Xie, B. Multi-objective optimization for the radial bending and twisting law of axial fan blades. *Processes* **2022**, *10*, 753.
20. Biedermann, T.M.; Reich, M.; Paschereit, C.O. Multi-Objective Modeling of Leading-Edge Serrations Applied to Low-Pressure Axial Fans. *J. Eng. Gas Turbines Power* **2020**, *142*, 111009.
21. Cheng, H.; Zhang, Z.; Xingen, L.; Penghao, D.; Junqiang, Z. Aerodynamic robustness optimization of aeroengine fan performance based on an interpretable dynamic machine learning method. *Reliab. Eng. Syst. Saf.* **2025**, *254*, 110654.
22. Masi, M.; Danieli, P.; Lazzaretto, A. Overview of the best 2020 axial-flow fan data and inclusion in similarity charts for the search of the best design. In *Proceedings of the Turbo Expo: Power for Land, Sea, and Air*; American Society of Mechanical Engineers: Washington, DC, USA, 2021; Volume 84898, p. V001T10A008.
23. Moradihaji, K.; Ghassemi, M.; Pourbagian, M. Off-design performance analysis of a radial fan using experimental, computational, and artificial intelligence approaches. *Eur. J. Mech. B/Fluids* **2024**, *104*, 150–172.
24. Angelini, G.; Corsini, A.; Delibra, G.; Tieghi, L. A multidimensional extension of balje chart for axial flow turbomachinery using artificial intelligence-based meta-models. *J. Eng. Gas Turbines Power* **2019**, *141*, 111012.
25. Huang, Z.; Shen, H.; Kung, K.; Carvalho, L.; Thai, A.; Watchmann, B.; Ramsarran, T.; Winkler, J.; Reimann, A.; Joly, M.; et al. Fan Wake Prediction Via Machine Learning. In Proceedings of the 28th AIAA/CEAS Aeroacoustics 2022 Conference, Southampton, UK, 14–17 June 2022; p. 2883.
26. Li, N.; Winkler, J.; Reimann, C.A.; Voytovych, D.; Joly, M.; Lore, K.G.; Mendoza, J.; Grace, S.M. Machine learning aided fan broadband interaction noise prediction for leaned and swept fans. In Proceedings of the AIAA AVIATION 2023 Forum, San Diego, CA, USA, 12–16 June 2023; p. 4297.
27. Fesquet, J.; Bauerheim, M.; Rojda, L.; Bousquet, Y.; Binder, N. Application of deep learning for fan rotor blade performance prediction in turbomachinery. *J. Turbomach.* **2025**, *147*, 111002.
28. Duan Luong, C.; Quang, B.N.; Thanh, L.H.; Cao, D.T.; Trung, H.N.; Ngoc, M.N. Real-Time Fan Anomaly Detection Using Embedded Machine Learning. In *Proceedings of the 2024 9th International Conference on Applying New Technology in Green Buildings (ATiGB)*; IEEE: New York, NY, USA, 2024; pp. 362–367.
29. Kumar, A.; Manas, M.; Pradeep, A. Identification and classification of operating flow regimes and prediction of stall in a contra-rotating axial fan using machine learning. *Aeronaut. J.* **2022**, *126*, 1351–1369.
30. Peng, C.C.; Tsai, M.C.; Chen, T.Y. Nonlinearity modeling for online estimation of industrial cooling fan speed subject to model uncertainties and state-dependent measurement noise. *Nonlinear Eng.* **2024**, *13*, 20240049.
31. Tao, P.; Liu, J.; Liang, T. Research on fault diagnosis method of axial flow induced draft fan of power plant based on machine learning. In *Proceedings of the 2019 4th International Conference on System Reliability and Safety (ICSRS)*; IEEE: New York, NY, USA, 2019; pp. 325–330.
32. Corsini, A.; Delibra, G.; Giovannelli, M.; Lucherini, G.; Minotti, S.; Rossin, S.; Tieghi, L. Identification of poorly ventilated zones in gas-turbine enclosures with machine learning. In *Proceedings of the Turbo Expo: Power for Land, Sea, and Air*; American Society of Mechanical Engineers: Washington, DC, USA, 2019; Volume 58547, p. V001T09A003.

33. Iranfar, A.; Terraneo, F.; Csordas, G.; Zapater, M.; Fornaciari, W.; Atienza, D. Dynamic thermal management with proactive fan speed control through reinforcement learning. In *Proceedings of the 2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*; IEEE: New York, NY, USA, 2020; pp. 418–423.
34. Zhang, T.; Xu, F.; Jia, M. A centrifugal fan blade damage identification method based on the multi-level fusion of vibro-acoustic signals and CNN. *Measurement* **2022**, *199*, 111475.
35. Ciaburro, G.; Padmanabhan, S.; Maleh, Y.; Puyana-Romero, V. Fan fault diagnosis using acoustic emission and deep learning methods. *Informatics* **2023**, *10*, 24.
36. Angelini, G.; Bonanni, T.; Corsini, A.; Delibra, G.; Tieghi, L.; Volponi, D. A meta-model for aerodynamic properties of a reversible profile in cascade with variable stagger and solidity. In *Proceedings of the Turbo Expo: Power for Land, Sea, and Air*; American Society of Mechanical Engineers: Washington, DC, USA, 2018; Volume 50985, p. V001T09A009.
37. Corsini, A.; Delibra, G.; Tieghi, L.; Tucci, F.A. Cascade with Sinusoidal Leading Edges: Identification and Quantification of Deflection with Unsupervised Machine Learning. In *Proceedings of the ASME Turbo Expo 2021: Turbomachinery Technical Conference and Exposition. Volume 1*; ASME: New York, NY, USA, 2021; Volume 1, pp. 1–10.
38. Pan, T.; Shi, K.; Lu, H.; Yang, Z.; Zhang, J.; Li, Q. A region-segmentation combinational loss model based on data-driven machine learning for a boundary layer ingestion fan. *Aerosp. Sci. Technol.* **2024**, *144*, 108780.
39. Zhang, B.; Wang, T.; Gu, C.; Shu, X. Blade optimization design and performance investigations of an ultra-low specific speed centrifugal blower. *Sci. China Technol. Sci.* **2011**, *54*, 203–210.
40. Zhang, L.; Wang, S.; Hu, C.; Zhang, Q. Multi-objective optimization design and experimental investigation of centrifugal fan performance. *Chin. J. Mech. Eng.* **2013**, *26*, 1267–1276.
41. Bamberger, K.; Carolus, T.; Belz, J.; Nelles, O. Development, validation, and application of an optimization scheme for impellers of centrifugal fans using computational fluid dynamics-trained metamodels. *J. Turbomach.* **2020**, *142*, 111005.
42. Meng, F.; Wang, L.; Ming, W.; Zhang, H. Aerodynamics optimization of multi-blade centrifugal fan based on extreme learning machine surrogate model and particle swarm optimization algorithm. *Metals* **2023**, *13*, 1222.
43. Lopez, D.I.; Ghisu, T.; Shahpar, S. Global optimization of a transonic fan blade through AI-enabled active subspaces. *J. Turbomach.* **2022**, *144*, 011013.
44. Angelini, G.; Corsini, A.; Delibra, G.; Tieghi, L. Exploration of Axial Fan Design Space with Data-Driven Approach. *Int. J. Turbomach. Propuls. Power* **2019**, *4*, 11.
45. Moradihaji, K.; Ghassemi, M.; Pourbagian, M. Ensemble Transfer Learning for Rapid Fan Performance Prediction Across Configurations. *Int. J. Comput. Fluid Dyn.* **2024**, *38*, 557–578.
46. Ahmed, Z.; Sohail, M.U.; Javed, A.; Swati, R.F. Development of a predictive tool for the parametric analysis of a turbofan engine. *Appl. Sci.* **2023**, *13*, 10761.
47. Hati, A.S. A comprehensive review of energy-efficiency of ventilation system using Artificial Intelligence. *Renew. Sustain. Energy Rev.* **2021**, *146*, 111153.
48. Li, T.; Wu, X.; He, Y.; Peng, X.; Yang, J.; Ding, R.; He, C. Small samples noise prediction of train electric traction system fan based on a multiple regression-fuzzy neural network. *Eng. Appl. Artif. Intell.* **2023**, *126*, 106781.
49. Tieghi, L.; Becker, S.; Corsini, A.; Delibra, G.; Schoder, S.; Czwielong, F. Machine-learning clustering methods applied to detection of noise sources in low-speed axial fan. *J. Eng. Gas Turbines Power* **2023**, *145*, 031020.
50. Corsini, A.; Delibra, G.; Giovannelli, M.; Lucherini, G.; Minotti, S.; Rossin, S.; Tieghi, L. Prediction of Ventilation Effectiveness for LM9000 Package with Machine Learning. In *Proceedings of the Turbo Expo: Power for Land, Sea, and Air*; American Society of Mechanical Engineers: Washington, DC, USA, 2020; Volume 84201, p. V009T21A006.
51. Maldonado-Correa, J.; Martín-Martínez, S.; Artigao, E.; Gómez-Lázaro, E. Using SCADA data for wind turbine condition monitoring: A systematic literature review. *Energies* **2020**, *13*, 3132.
52. Barnabei, V.F.; Ancora, T.C.; Delibra, G.; Corsini, A.; Rispoli, F. Semi-Supervised Deep Learning Framework for Predictive Maintenance in Offshore Wind Turbines. *Int. J. Turbomach. Propuls. Power* **2025**, *10*, 14.
53. Anthropic. Claude Sonnet 4.5, 2025. Available online: https://www.anthropic.com/news/claude-sonnet-4-5 (accessed on 3 June 2026).