

Frequency Domain Prediction of NACA 0012 Self-noise Using Data-Driven Machine Learning

Rishav Mishra[1], Alvis Fong[2], and Bade Shrestha[1(✉)]

[1] Department of Mechanical and Aerospace Engineering, Western Michigan University, Kalamazoo, USA
bade.shrestha@wmich.edu

[2] Department of Computer Science, Western Michigan University, Kalamazoo, USA

Abstract. The paper presents the results of investigation on the feasibility of predicting the frequency domain of self-noise spectrum of a NACA 0012 airfoil employing a hybrid machine learning algorithm framework. The framework consisted of three learning architectures: CatBoost, Deep Neural Networks (DNN) and Support Vector Regression (SVR). These learning architectures were further optimized by Arithmetic Optimization Algorithm (AOA) using hyperparameter tuning. Instead of using widely available NASA Airfoil Self-Noise dataset, the models were trained and assessed employing the experimental data acquired from a low-speed anechoic wind tunnel at Western Michigan University. The experimental noise data sets of airfoil trailing edge were generated for various angle of attacks at a freestream velocity of 15 m/s. Each model was examined using baseline and a Fourier-enhanced formulation incorporating logarithmic and harmonic frequency features. The performance of each model was evaluated using standard regression metrics; root mean square error (RMSE), relative absolute error (RAE) and coefficient of determination (R^2). The results demonstrated that the Fourier-enhanced feature representations regularly predicated better results for all models tested for both broadband and tonal noise characteristics. Among the three architectures examined, CatBoost tuned with AOA achieved the best performance overall. These results support the effectiveness of using a hybrid machine learning framework with AOA based hyperparameter optimization to predict the frequency spectrums of trailing edge self-noise.

Keywords: Applied Computing · Physical Sciences and Engineering · Aerospace · Aeroacoustics · Machine learning

1 Introduction

Wind energy has become one of the widely adopted renewable and sustainable energy sources in the past few decades due to its competitiveness and reliability. As the global warming problem is exaggerating, expansions of the renewable and sustainable energy installations are critical in reducing and resolving these environmental challenges. As wind energy installations are rapidly increasing, growing numbers of communities find themselves in their proximity causing serious concern of noise pollution generated by

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026, LNCS 16767, pp. 110–129, 2027.
https://doi.org/10.1007/978-3-032-30491-9_8

these wind farms. These are generally low frequency noises that cause significant annoyance and long-term health problems. While there have been a large number of research and investigations conducted, still there is a challenge in developing viable new technology to overcome this issue. The research particularly focused on mitigating the adverse effects of turbine-induced noise pollution, or airfoil self-noise pollutions emphasizing the development of advanced technologies to minimize its environmental and societal impact. Airfoil self-noise is caused by the interaction between an airfoil blade and the turbulence produced in its own boundary layer and nearby wakes [1]. The basic mechanisms of sound generation by wind turbine blades can be categorized into many types such as vortex shedding noise, turbulence-structure interaction noise, and trailing edge noise, etc. as illustrated in Fig. 1. The dominant source of the trailing edge noise is generated by turbulent boundary layer vortices scatter at the sharp edge of the airfoil [1, 2].

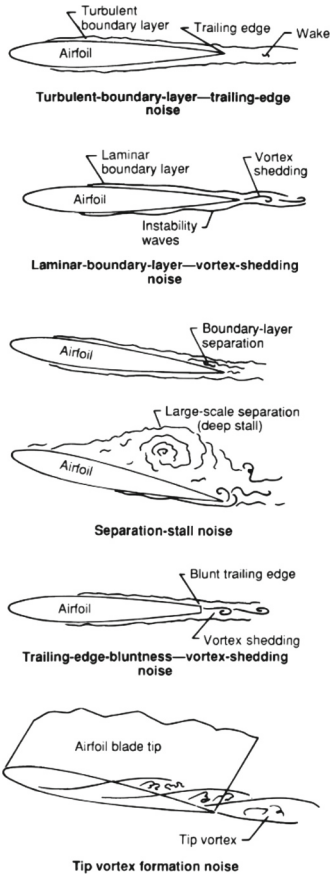

Fig. 1. Aerodynamic Noise Generation Mechanisms

The major challenge in noise reduction is a lack of models that can accurately determine the airfoil self-noise and its source terms. Since this noise is inherently connected with complex fluid-structure interactions with the solid surfaces, it depends on many factors like airfoil geometry, flow conditions, and boundary layer dynamics. The major challenge lies in modeling aerodynamic noise that accurately captures the complex, and stochastic dynamics of turbulent flows. However, simplified turbulence models offer easier implementation but limited accuracy, hindering the development of quieter airfoils.

Peterson et al. [3] conducted earlier experimental studies focusing on low Reynolds numbers and moderate angles of attack to reveal the role of boundary layer instabilities in tonal noise generation. Their experiments were conducted in an anechoic open jet wind tunnel at NASA Langley Research Center using NACA 0012 airfoils with chord lengths (C) up to 0.3 m. The airfoils were subjected to a 2D laminar inflow with varying freestream velocities (U) up to 71 m/s and angles of attack (α) up to 25°, corresponding to chord-based Reynolds numbers (Re_c) reaching 1.5×10^6. Their experiments uncovered a characteristic "*ladder-type*" acoustic signature, that corelated the dominant tonal frequency component with freestream velocity. This scaling behavior was quantified through an empirical relation for the characteristic shedding frequency f_s with kinematic viscosity of air ν, and it expressed as:

$$f_s = \frac{0.02 \times U_\infty^{\frac{3}{2}}}{\sqrt{c \times \nu}} \tag{1}$$

Along with the primary tone, several discrete and secondary tonal components were observed, each following a consistent frequency–velocity scaling proportional to $U_\infty^{0.8}$, creating rise to the ladder pattern. Subsequent experiments by Arbey and Bataille [4] and Arcondoulis et al. [5] reconfirmed the presence of this ladder pattern and provided further insight into its spectral composition. They demonstrated that the acoustic radiation consists of two principal components:(i) a broadband spectral hump associated with unstable boundary layer frequencies, and (ii) a superimposed series of discrete tonal peaks, as illustrated in Fig. 2. These observations suggested that trailing edge noise arises from a complex interaction between broadband aerodynamic processes and discrete instability driven by tonal mechanisms.

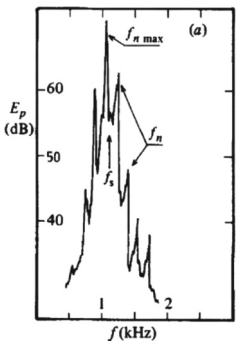

Fig. 2. Typical Noise Radiated Spectra NACA0012.8 @ U = 20.2m/s [4]

Though these experimental studies provided critical physical insight into the mechanisms underlying trailing edge noise, translating this understanding into predictive numerical models remains challenging. Advancements in computing now enable high fidelity turbulence simulations capable of resolving many of the flow features responsible for aerodynamic noise generation. However, numerical approaches still rely on experimental data for validation to ensure their accuracy. In prior studies, promising results were obtained using a similar data-driven methodology applied to NASA's NACA 0012 airfoil dataset (n0012-il) [2, 6]. The dataset comprises 1,503 samples collected under varying wind tunnel velocities and angles of attack while maintaining a fixed span and observer position. Although the dataset spans a broad range of operating conditions, the overall availability of labeled experimental data remains limited from the perspective of data-driven model development. In [2], CatBoost was combined with various metaheuristic optimization algorithms, while [6] employed a deep neural network framework to model airfoil self-noise. In the latter study, the relatively limited availability of experimental data motivated the exploration of multiple training and testing split ratios together with fivefold cross validation, in order to evaluate model robustness and reduce bias. Although better predictive performance was achieved, both studies highlighted the importance of expanding the availability of labeled experimental data to further improve generalization.

To tackle data availability constraints, alternative approaches have been explored using synthetically generated datasets. For example, in [7], a generalized airfoil database was constructed using a reference set of 31 commonly used wind turbine airfoils. A wide range of airfoil geometries were generated using airfoil functional integration theory, B spline curves, and the Class Shape Transformation method. Noise labels for this generalized airfoil set were then computed using the Brooks-Pope-Marcolini (BPM) semi empirical model, and a residual neural network (ResNet-18) was trained for noise prediction. While this approach enabled large scale data generation, the resulting labels remained dependent on semi-empirical modeling assumptions rather than direct experimental measurements.

All these studies accentuated the significance of labeled aeroacoustics data obtained under varying flow regimes for the development of reliable data-driven prediction models, especially when the objective is to capture the full noise spectra rather than isolated acoustic metrics. The current study used experimental data obtained from the low-speed wind tunnel, where the time-based raw pressure fluctuations using microphone was transformed into the frequency domain using Fast Fourier Transform (FFT). The FFT enabled decomposition of a time domain signal into its constituent frequency components, thereby revealing the spectral content of the input acoustic signal and allowing study of tonal and broadband features [8, 9]. To examine the robustness of the learning frameworks, spectral estimates were computed using varying FFT window lengths, enabling assessment of model sensitivity to spectral resolution.

The resulting full band noise spectra were used as learning targets, aiding a spectrum resolved assessment of each model's ability to reproduce the frequency domain characteristics of airfoil self-noise within the frequency range of interest. The proposed framework predicts the full distribution SPL(f), addressing a limitation of many data-driven aeroacoustics approaches that primarily report a single scalar noise metric. Since

the model learns the complete SPL(f) curve under controlled flow conditions, it captures frequency dependent behavior (including broadband trends and tonal components) that are essential for noise diagnosis and mitigation. The framework applies CatBoost tuned using AOA, while DNN and SVR tuned using the same strategy serve as benchmark models within a unified pipeline. Furthermore, rather than treating noise prediction as tabular regression, this work uses experimentally measured spectra and introduces a Fourier-enhanced frequency representation that encodes harmonic structure and multi scale spectral trends, together with angle of attack terms. Using a unified AOA based hyperparameter optimization across all models ensured consistent training and fair comparison, enabling practical frequency aware noise modeling for airfoil designs with lower noise and sustainable wind energy design.

2 Methodology

2.1 Experimental Setup

All experiments were performed in a controlled laboratory facility using a low-speed wind tunnel housed within a large (3 × 3 × 3 m) anechoic chamber. The internal walls of the chamber and all the reflective surfaces within the chamber are lined with sound absorbing foam wedges of varying shapes and sizes designed to reduce background noise contamination and limit the acoustic reflections within the chamber. The test section of the wind tunnel has a circular outlet of radius 0.115 m. Airflow enters the anechoic chamber through a circular inlet and exits via a 0.35 m diameter window on the opposite side. The airfoil model was installed at the center of the jet and mounted between two side plates that were flush with the wind tunnel outlet in order to suppress tip vortex formation and promote quasi two-dimensional flow, as shown in Fig. 3. Acoustic measurements were obtained using a $^1/_2$ " free field polarized microphone system consisting of a PCB Piezotronics Model 377A07 microphone coupled with a Model 426E01 preamplifier. The microphone was calibrated prior to testing with a Larson Davis CAL200 precision acoustic calibrator. During measurements, the microphone was located 2 chord lengths from the trailing edge of the airfoil, depending on the configuration at mid span and oriented perpendicular to the freestream direction. To reduce spurious acoustic contamination, the microphone mounting structure was treated with sound absorbing material to limit reflections from the support hardware.

The airfoil investigated in this study is a symmetric two-dimensional NACA 0012 profile. Experiments were conducted using nominal chord length, C = 0.1143 m, with a spanwise length of approximately 0.12 m. Measurements were carried out at five geometric angles of attack, $\alpha_g = 0°, 3°, 5°, 7°$ and 10°, under a fixed freestream velocity of $U_\infty = 15$ m/s. These conditions correspond to chord-based Reynolds numbers of $Re_c \approx 96{,}000$, assuming the air density $\rho = 1.169$ kg/m^3 at 28 °C within the testing facility. The finite width of the open jet test section introduces flow confinement effects that alter the effective aerodynamic angle of attack relative to the prescribed geometric value. As a result, two-dimensional correction factors were applied to account for jet contraction and blockage effects. In such configurations with limited jet height, the true angle of attack α differs from the prescribed geometric angle α_g, necessitating the

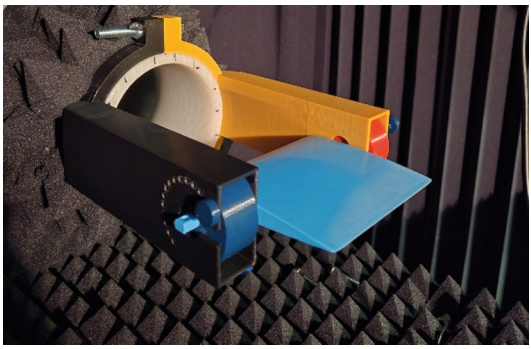

Fig. 3. Experimental Setup with NACA0012 Baseline Model

use of empirical corrections. Following the established approach outlined in [10], the relationship between α and α_g is given by:

$$\alpha = \frac{\alpha_g}{\zeta} \tag{2}$$

where:

$$\zeta = (1 + 2\sigma)^2 + \sqrt{12\sigma} \tag{3}$$

and

$$\sigma = \frac{\pi^2}{48}\left(\frac{C}{W}\right)^2 \tag{4}$$

Here C is the airfoil chord length and W = 0.115 m is the width of the circular jet within which the airfoil is placed. Acoustic data acquisition was performed using a National Instruments NI 6212 data acquisition system interfaced through the Python NI DAQ module. A custom graphical user interface developed using Node.js was employed to control the acquisition process and interface directly with the microphone system. The frontend allowed the user to specify key parameters such as sampling rate and acquisition duration. The measured voltage fluctuations over the acquisition period were converted into pressure fluctuations using the calibrated microphone sensitivity, and the resulting pressure signals were then transformed into the frequency domain using the fast Fourier transform.

2.2 CatBoost Classification Algorithm

CatBoost is a gradient boosted decision tree (GBDT) algorithm developed to improve learning performance on structured datasets that contain heterogeneous and categorical features. The algorithm was introduced by Prokhorenkova et al. with the primary objective of reducing bias caused by conventional categorical encoding strategies and improving generalization performance in supervised learning tasks [11]. A broader conceptual and empirical discussion of CatBoost within the family of gradient boosted decision trees is provided in the interdisciplinary survey by Hancock and Khoshgoftaar

[12]. A prominent distinguishing component of CatBoost is its native handling of categorical variables through ordered target statistics. Unlike traditional approaches that compute every category statistic using the full training dataset, which is prone to introducing target leakage, CatBoost computes target-based encodings using only samples that precede the current instance under a random permutation of the training data. This mechanism ensures that the encoded feature values do not contain information from the target of the current sample, thereby mitigating conditional shift and reducing overfitting [11, 12]. For a categorical feature x_i, the encoded numerical value for the k-th training instance is given by:

$$\hat{x}_i^{(k)} = \frac{\sum_{j<k} 1(x_i^{(j)}=x_i^{(k)})y_j + ap}{\sum_{j<k} 1(x_i^{(j)}=x_i^{(k)}) + a} \tag{5}$$

where $1(\cdot)$ denotes the indicator function, y_j is the target value associated with the j-th sample, p represents a prior estimate commonly chosen as the global mean of the target variable, and a is a regularization parameter introduced to stabilize the encoding for infrequently occurring categories [11, 12]. To complement ordered target statistics, CatBoost employs ordered boosting, where gradient estimates for each decision tree are computed using the same data ordering applied during categorical encoding. This strategy prevents biased gradient estimation caused by incorporating information from the current sample and improves generalization performance [11]. CatBoost further adopts symmetric decision trees, in which the same splitting rule is applied at each tree level. This structure produces balanced trees with predictable depth and inference cost, reducing overfitting and enabling efficient training and evaluation [11, 12].

CatBoost has been widely adopted in established application domains such as automation, healthcare, transportation, and education, where structured data with mixed feature types are common. Beyond these conventional use cases, its reliable performance has enabled effective deployment even in niche domains. For example, in geotechnical engineering, CatBoost has been successfully applied to slope stability prediction, where it effectively captures nonlinear relationships among geological and environmental parameters while maintaining stable predictive performance [13]. Similarly, CatBoost based hybrid models have been employed in educational data mining to assess student academic performance using datasets composed of mixed numerical and categorical attributes, achieving superior classification accuracy compared to conventional machine learning approaches [14]. Together, this evidence highlights CatBoost's suitability for structured data scenarios characterized by heterogeneity and complex feature interactions.

2.3 Deep Neural Network (DNN)

Deep Neural Networks (DNN) constitute a powerful class of data-driven models capable of learning complex nonlinear relationships present in the raw data through hierarchical feature extraction. DNNs automatically learn hierarchical representations directly from data by stacking multiple layers of nonlinear transformations, enabling effective modeling of structured data as shown in the example in Fig. 4 [15, 16].

In contrast to the other models that specifically rely on manual features, DNNs is a leading in a wide range of applications and disciplines regardless of the input data.

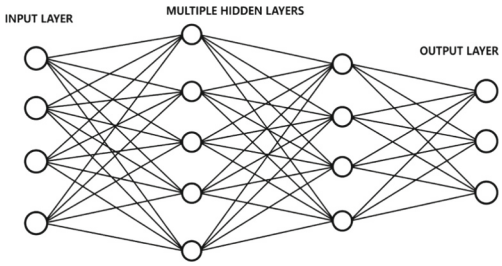

Fig. 4. Illustrative example of a DNN with input, hidden and output layers

Structurally, a standard feedforward DNNs architecture consists of an input layer, a sequence of hidden layers, and a final output layer. Within each hidden layer, the network performs an affine transformation followed by a nonlinear activation which helps the model with predictions. For a network with a depth of L layers, the forward propagation through a specific layer l is defined as:

$$a^{(l)} = g^{(l)}(W^{(l)}a^{(l-1)} + b^{(l)}) \tag{6}$$

where $\mathrm{W}^{(l)}$ and $\mathrm{b}^{(l)}$ denote the weights and biases of layer l, $g^{(l)}(\cdot)$ is a nonlinear activation function, and $\mathrm{a}^{(0)}$ is the initial input feature vector. This formulation follows the standard feedforward architecture used in DNNs, as described in [15, 16]. The training of DNNs involves optimizing the parameters involved in the network by minimizing a specific loss function through gradient-based methods. In the context of classification, the final layer typically utilizes the softmax activation function to generate class posterior probabilities. This relationship is defined as:

$$p_k = \frac{exp(z_k)}{\sum_j exp(z_j)} \tag{7}$$

where z_k represents the pre-activation logic for the k-th class. This formulation was popularized in early deep learning systems, where DNNs replaced traditional generative models for estimating posterior probabilities of latent states [17]. For regression problems, which are central to data-driven physical modeling, the learning objective is typically defined using the mean squared error (MSE):

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N} \|y^{(i)} - \hat{y}^{(i)}\|_2^2 \tag{8}$$

where $y^{(i)}$ and $\hat{y}^{(i)}$ denote the true and predicted outputs, respectively. The application of MSE as the standard loss metric for deep neural network regression is widely accepted in the field [15, 18]. The model is optimized by updating the parameters which are computed incorporating the gradients of the loss function via the backpropagation algorithm. This approach allows for the efficient training of deep architectures utilizing stochastic gradient descent (SGD) or its adaptive variants [15, 16] by systematically applying the chain rule across the network layers. A important aspect of DNNs is their capacity to develop layered internal representations. Early layers generally identify simple, localized patterns, and these features become progressively more abstract and specialized for

the task as they move through higher layers. This progressive feature refinement enables deep networks to model complex nonlinear patterns, forming the basis for their effectiveness in a wide range of perception tasks [16], including speech processing applications [17]. Beyond these traditional perception tasks, recent studies have demonstrated the effectiveness of deep hierarchical representations in diverse application areas such as lane detection in autonomous systems [19], flow prediction in turbomachinery and fluid dynamics [20], and complex medical image interpretation for disease diagnosis [21].

2.4 Support Vector Regression (SVR)

Support Vector Regression is a kernel based supervised learning algorithm that is rooted in the principles of statistical learning theory and derived from the Support Vector Machine Framework (SVM) methodology. In contrast to the traditional regression methods that aim to minimize the empirical loss on each data point directly, SVR formulates regression as a margin-based optimization problem. The objective is to discover a function that accurately reconstructs target values while simultaneously enforcing model complexity through regularization. The structural risk minimization principle underlying SVR confers robustness and excellent generalization properties, especially when applied to high-dimensional data or datasets with significant noise [22, 23]. Given a training set of samples $\{(x_i, y_i)\}_{i=1}^{N}$, where $x_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$, SVR constructs a regression function by minimizing a regularized objective subject to an ε-sensitive error tolerance. Following the formulation presented by [24], the SVR optimization problem is expressed as:

$$\underset{w,b,\xi_i,\xi_i^*}{min} \frac{1}{2}\|w\|^2 + C\sum_{i=1}^{N}(\xi_i + \xi_i^*) \tag{9}$$

subject to

$$\begin{aligned} y_i - (w^T x_i + b) &\leq \varepsilon + \xi_i, \\ (w^T x_i + b) - y_i &\leq \varepsilon + \xi_i^*, \\ \xi_i, \xi_i^* &\geq 0, \end{aligned} \tag{10}$$

In the above formulation, the regularization parameter $C > 0$ governs the critical equilibrium between the complexity of the model flatness and the penalty imposed on deviations from the training data. This optimization algorithm is the foundational concept of ε-sensitive error and functions as the baseline SVR model. These fundamental principles serve as the basis for numerous advanced regression variants that offer greater sophistication and robustness [24]. To expand SVR capabilities beyond linear regression, input data is often projected into a higher dimensional feature space using a nonlinear transformation. Within this generalized kernel-based framework, the regression function is expressed as:

$$f(x) = w^T \phi(x) + b \tag{11}$$

where $\phi(\cdot)$ represents the feature mapping generated by a kernel function, and b denotes the bias term. This explicit representation in kernel space is integral to the original

formulation of Support Vector Machines and serves as the structural foundation for kernelized SVR models [25]. By introducing Lagrange multipliers and applying the kernel trick, the SVR solution can be expressed in dual form as:

$$f(x) = \sum_{i=1}^{N} (\alpha_i - \alpha_i^*) K(x_i, x) + b \tag{12}$$

where $K(x_i, x) = \phi(x_i)^T \phi(x)$ is a kernel function. In this form, only a subset of training samples contributes to the regression function, yielding sparse solutions defined by support vectors and enabling computationally efficient modeling, particularly in large scale settings. In real world implementations, improvements in model adaptability are crucial. Therefore, various SVR modifications have been proposed in published literature. Refined tolerance mechanisms built upon the classical formulation is seen in [24] to better manage noise and heterogeneous data, while retaining the fundamental SVR objective. Elsewhere, linear SVR formulations have deemed to be cost-effective with competitive predictive performance alternative in high-dimensional feature spaces [23]. Applications of SVR are diverse, and it has been utilized in domains characterized by nonlinear relationships and structured input spaces. In engineering, it functions as a surrogate model to accelerate CFD workflows, including aerodynamic design problems such as wind turbine blade twist optimization [26]. In manufacturing and process modeling, it has demonstrated strong predictive accuracy for deformation estimation and bending behavior under various operating conditions [27]. Beyond engineering applications, SVR has also been widely adopted in data-driven decision and forecasting problems. In the medical field, it is utilized to model disease progression and continuous clinical outcomes, capitalizing on its resilience to noise and efficacy with limited sample sizes [28, 29].

2.5 Arithmetic Optimization Algorithm

The Arithmetic Optimization Algorithm (AOA) is a population-based metaheuristic optimization technique inspired by the fundamental arithmetic operators: addition, subtraction, multiplication, and division [30]. In contrast to the numerous prevalent optimization methodologies that rely on metaphors from biology or physics, AOA utilizes these mathematical operations directly as search strategies. This allows the algorithm to proficiently guide candidate solutions through the search space, culminating in a structure that is conceptually straightforward yet possesses the capacity for robust and flexible performance. AOA begins by initializing a population of candidate solutions uniformly within search bounds. The j-th dimension of the i-th candidate is initialized as:

$$X_{i,j} = LB_j + r_{i,j}(UB_j - LB_j) \tag{13}$$

where LB_j and UB_j denote the lower and upper bounds, and $r_{i,j} \in (0,1)$ is a random number [30]. Mathematical Optimization Accelerator (MOA) controls the balance between exploration and exploitation defined as:

$$MOA(t) = Min + t \times \frac{Max - Min}{T} \tag{14}$$

where t is the current iteration, T is the maximum number of iterations, *Min* and *Max* are control parameters. While, Mathematical Optimization Probability (MOP) governs the step size of population updates:

$$MOP(t) = 1 - (\frac{t}{T})^{\frac{1}{\alpha}} \tag{15}$$

with α controlling the rate of decay [30]. During the exploration phase, multiplication and division operators are used to generate large search steps:

$$X_{i,j}^{t+1} = Best_j \times \frac{r}{MOP} \text{ or } X_{i,j}^{t+1} = \frac{Best_j}{r \times MOP} \tag{16}$$

where $Best_j$ is the best solution so far. In the exploitation phase, addition and subtraction operators refine solutions locally:

$$X_{i,j}^{t+1} = Best_j \pm r \times MOP \tag{17}$$

The transition between these phases and the selection of arithmetic operators are determined probabilistically using the MOA value, allowing AOA to adapt its search behavior over iterations [30]. While the standard AOA delivers good results, it may suffer from slow convergence and a tendency to become trapped in local optima in complex search spaces [31]. To mitigate such issues, Fang et al. reported that utilizing AOA to optimize SVM parameters leads to superior convergence characteristics [32], while incorporation of multi strategy fusion techniques such as opposition based learning and chaotic initialization strengthened robustness and convergence stability [33]. Recently, hybrid learning frameworks that combine AOA with advanced models like CatBoost have shown higher predictive accuracy for aeroacoustic applications validating the effectiveness of AOA as a potent optimization strategy for data-driven noise prediction [2]. This specific paradigm directly motivates the approach adopted in the present work.

2.6 Proposed Hybrid AOA Optimization for CatBoost, DNN and SVR

While Arithmetic Optimization Algorithm (AOA) has previously been combined with CatBoost for airfoil self-noise prediction, most notably in [2], the present work extends this paradigm in both methodological scope and data representation. Prior implementations utilized AOA within a singular learning architecture, often constrained by limited feature sets. In contrast, the framework proposed here introduces a generalized hybrid AOA strategy capable of simultaneously optimizing CatBoost, DNN and SVR. This mutli model approach enables a systematic and rigorous comparison across fundamentally distinct algorithmic classes. Furthermore, accurately predicting airfoil self-noise requires a deep understanding of how different flow regimes affect the acoustic signature of the airfoil, specifically distinguishing between broadband turbulence and narrowband tonal instabilities [5, 34]. Consequently, the proposed approach is specifically tailored for frequency domain datasets, where the objective is not merely a scalar SPL prediction but rather the reconstruction of the entire Sound Pressure Level (SPL) spectrum. This requirement necessitates specialized feature engineering and evaluation metrics designed to preserve both broadband characteristics and critical narrowband tonal peaks.

Therefore, a feature has been embedded into the proposed model, which is critical for meaningful aeroacoustics modeling, as given below.

Proposed Algorithm: Hybrid AOA Optimization for CatBoost, DNN and SVR

1 Input:

Train file $\mathbf{D_{tr}} = \{(f_i, \alpha_i, y_i)\}_{i=1}^{N_{tr}}$, test file $\mathbf{D_{te}} = \{(f_i, \alpha_i, y_i)\}_{i=1}^{N_{te}}$ max AOA iterations T, population size N_p.

2 Preprocess:

Convert (f, α, y) to numeric, remove non-finite rows, keep $y \geq 0$, clip $f \geq 10^{-9}$.

3 Feature Engineering (two variants):

3A Baseline:

$$X^{(B)} = [log\,(1+f), \alpha, \alpha f, (\alpha f)^2]$$

3B Fourier-enhanced: Let $x = \log_{10}(f)$. Construct

$X^{(F)} = [log\,(1+f), x, x^2, \ldots\ldots, x^P, sin(2\pi k x), cos(2\pi k x), \alpha, \alpha^2, \alpha x, \alpha x^2, (f)]$

For $k = 1, \ldots., K$, with optional raw f.

4 Standardize:

Fit StandardScaler on X_{tr}, transform X_{tr}, X_{te}.

5 AOA Setup (model specific search vectors):

Initialize a population of candidates x_i within bounds:

CatBoost: $\{\text{depth}, \eta, \lambda, T_{bag}\}$, **DNN:** $\{\text{depth}, \log_{10}(\eta), \log_{10}\lambda\}$,

SVR: $\{\log_{10}(C), \log_{10}(\gamma), \epsilon\}$

6 While $t < T$ **do:**

a. Compute MOA and MOP to balance exploration and exploitation.

b. Update candidate solutions using arithmetic operators

c. Evaluate fitness using validation RMSE with *fidelity-controlled subsampling*.

d. Update the global best solution.

e. Set $t = t + 1$

7 End While

8 Final Training (using best AOA hyperparameters):

Retrain each model on the full standardized training set using the optimal AOA-selected hyperparameters (with early stopping behavior), then predict on X_{te}.

9 Output:

Optimized models for both $X^{(B)}$ **and** $X^{(F)}$, frequency domain SPL predictions and RMSE based evaluation on test set.

The developed framework effectively bridges the gap between physics informed frequency domain feature with model specific learning architectures, while employing AOA as a unified, model agnostic optimizer for hyperparameter selection. To ensure proper assessment of this hybrid methodology, a comprehensive strategy is devised:

i. controlled dataset partitioning based on repeated wind tunnel experiments,
ii. systematic comparison of model performance before and after the introduction of Fourier-enhanced frequency features, and
iii. quantitative evaluation using established statistical error indices computed over the full frequency spectrum.

At each AOA iteration, candidate hyperparameter sets are refined using the arithmetic operators governed by exploration-exploitation control parameters. Model performance is assessed based on validation error metrics computed on predicted frequency domain SPL values. To reduce computational load, a fidelity controlled subsampling approach is implemented during optimization. Initial iterations evaluate candidates on reduced data subsets, while subsequent iterations gradually enhance evaluation fidelity. This

ensures efficient exploration of hyperparameter search space while preserving stability and accuracy near convergence (also refer to Appendix Fig. 9).

2.7 Dataset Description and Preprocessing

The datasets used in this study were obtained from the experiments conducted in a low-speed anechoic wind tunnel facility, where acoustic pressure signals were measured for symmetric NACA 0012 airfoils under controlled operating conditions. To evaluate the model's generalization capability across geometric configurations, airfoil model with chord length of C = 0.1143 m was subjected to a constant free-stream velocity of U_∞ = 15 m/s, sweeping through a range of geometric angles of attack $\alpha_g = 0°, 3°, 5°, 7°$ and 10°. For each test condition, a free-field pre-polarized microphone system was positioned at the midspan location of the airfoil and at a distance of 2C above the trailing edge to isolate the radiated noise while minimizing near-field hydrodynamic interference [35]. The microphone was shielded using a standard aerodynamic windscreen to minimize the influence of flow-induced pseudo sound, which contains of hydrodynamic pressure fluctuations caused by local turbulence [36]. To ensure sufficient temporal and statistical convergence, data acquisition was performed at a sampling rate of 20kHz with a recording duration of 30 s per run. The raw microphone pressure time histories were transformed into the frequency domain using FFT to obtain the power spectral density (PSD) of the acoustic pressure fluctuations. Following the standard aeroacoustics signal processing practice, the Sound pressure level (SPL) was computed from PSD using:

$$\text{SPL} = 10log_{10}(\frac{P_{xx}}{p_{\text{ref}}^2}) \tag{18}$$

where P_{xx} denotes the mean squared pressure density and $p_{\text{ref}} = 20\mu Pa$ is the reference acoustic pressure in air [37]. The resulting SPL spectra constitute the primary modeling target. The input vector utilized within the learning frameworks comprises the frequency domain SPL spectra, geometric angles of attack, free-stream velocity and airfoil chord length. To mitigate data leakage and ensure an unbiased evaluation, the dataset was partitioned into equal training and testing subsets based on two distinct, sequential experimental runs performed under identical nominal flow conditions as shown in Fig. 5. This partitioning strategy maintains closely matched experimental conditions while allowing model evaluation under the quasi-transient nature of the flow, thereby assessing predictive performance on unseen realizations that reflect inherent flow variability.

2.8 Effect of Frequency Representation on Model Performance

To study the impact of the proposed frequency domain feature engineering across all the models, each model is assessed using two distinct configurations that vary only in their representation of the frequency input. In the baseline configuration, frequency is provided to the models as a compact, low order representation based on $log(1+f)$ and simple interaction terms with angle of attack, i.e., $[log(1+f), \alpha, \alpha f, (\alpha f)^2]$. In doing so, frequency is a simple monotonic scaling and low order couplings, which forces the learning algorithms to derive the spectral structure, periodicity, and harmonic content implicitly from

the limited frequency parametrization. Conversely, the Fourier-enhanced configuration substitutes the low order representation with a structured basis built on $x = log_{10}(f)$, polynomial expansions x, harmonic sine/cosine components $sin(2\pi kx) and cos(2\pi kx)$, and α-interaction terms (eg., $\alpha, \alpha^2, \alpha x, \alpha x^2$), with the optional inclusion of raw f. By explicitly encoding spectral periodicity and multi-scale trends in the feature space, the burden on the models to discover oscillatory and broadband characteristics purely from data is reduced, yielding an efficient and physically meaningful representation of airfoil self-noise behavior. Similar hybrid strategies utilizing Fourier based spectral representations have been shown to significantly enhance feature extraction and classification accuracy in other acoustic monitoring applications, such as industrial pipeline leak detection [38].

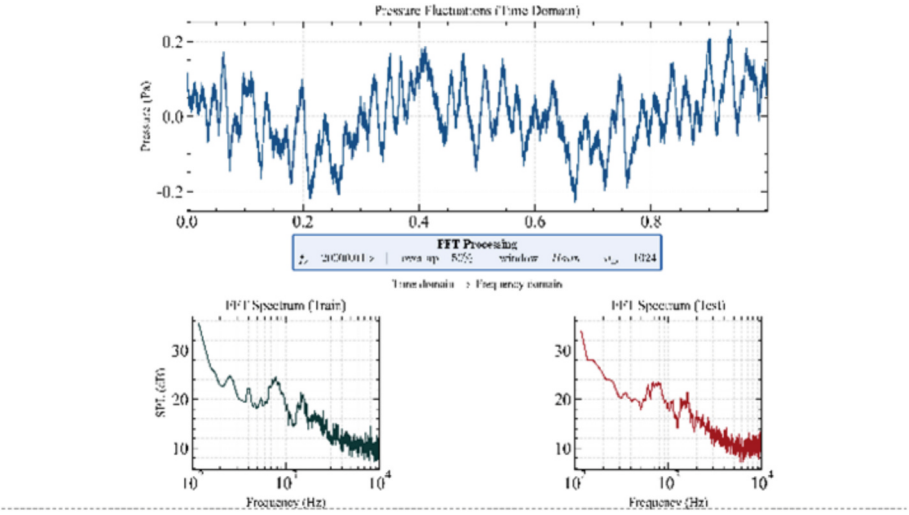

Fig. 5. Example Frequency domain transformation and resulting SPL spectra split

3 Results

The predictive performance of the hybrid CatBoost–AOA, DNN–AOA, and SVR–AOA frameworks was evaluated using both baseline and Fourier-enhanced hybrid models. The quantitative metrics were summarized in Tables 1 and 2, together with the parity plots and the SPL spectra comparisons shown in Figs. 6 and 7. It provided insight into how chosen frequency representation influences each model's ability to reproduce the broadband and tonal noise characteristics. In the baseline configuration, frequency was represented using low order transformations and interaction terms, which emphasized smooth spectral variation rather than explicit harmonic structure. CatBoost-AOA captured the broadband and dominant tonal peaks with reasonable accuracy; although the tonal peaks were slightly smoothed, their locations and relative magnitudes were largely preserved, showing that the model can exploit structured feature interactions even without harmonic encoding. DNN-AOA reproduced the broadband structure and follows general trend of the tonal peaks, but failed to resolve them as sharp narrowband peaks,

particularly at higher frequencies. In contrast, SVR-AOA only captured the broadband behavior, while the tonal components were strongly attenuated or absent in the predicted spectra. Overall, these results suggest that low order frequency features support learning of broadband trends, whereas accurate reconstruction of narrowband tonal peaks depends strongly on the model architecture, with CatBoost showing the highest inherent capability and DNN and SVR exhibiting progressively reduced narrowband spectral fidelity.

Table 1. Quantitative Performance of baseline hybrid models without explicit Fourier-enhanced features.

	Train				Test			
Model	**RMSE**	**MAE**	**RAE**	R^2	**RMSE**	**MAE**	**RAE**	R^2
CatBoost-AOA	0.5350	0.3572	0.0662	0.99519	0.5523	0.3744	0.0696	0.99486
DNN-AOA	0.7571	0.5044	0.0935	0.99037	0.7572	0.5030	0.0935	0.99034
SVR-AOA	1.1575	0.7028	0.1303	0.97749	1.1544	0.6982	0.1298	0.97755

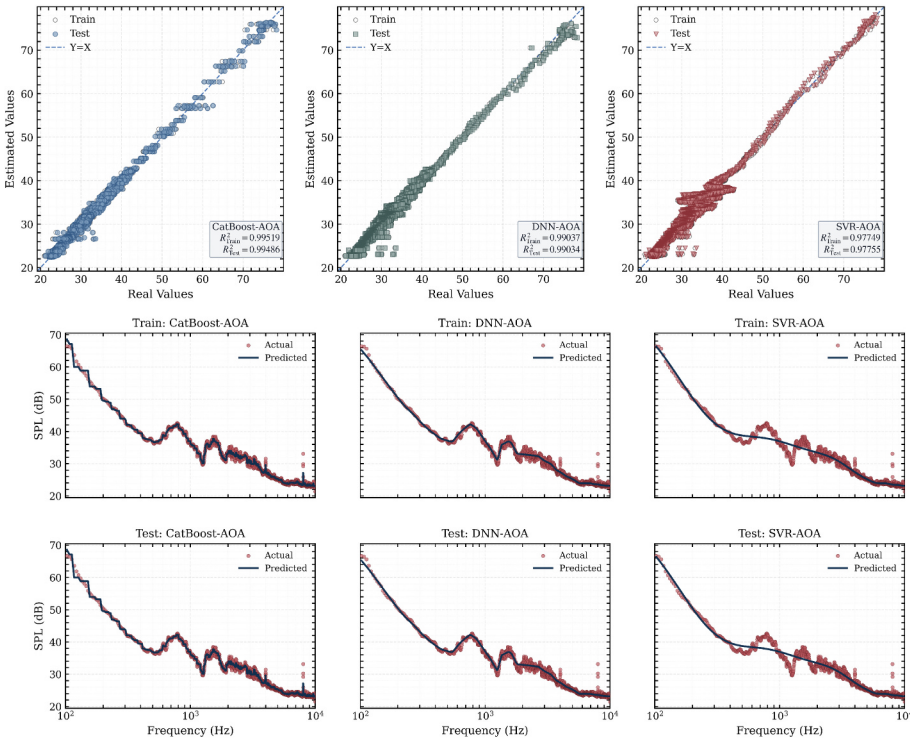

Fig. 6. Parity plots of baseline hybrid models without explicit Fourier-enhanced features (Top) and comparison of measured and predicted SPL spectra obtained from the same models (Bottom)

Table 2. Quantitative Performance of baseline hybrid models with Fourier-enhanced features.

	Train				Test			
Model	**RMSE**	**MAE**	**RAE**	R^2	**RMSE**	**MAE**	**RAE**	R^2
CatBoost-AOA (Fourier)	0.3079	0.2399	0.0445	0.99841	0.3453	0.2688	0.0500	0.99799
DNN-AOA (Fourier)	0.9158	0.4684	0.0869	0.98591	0.9189	0.4787	0.0890	0.98577
SVR-AOA (Fourier)	0.5320	0.3252	0.0603	0.99525	0.5514	0.3456	0.0642	0.99488

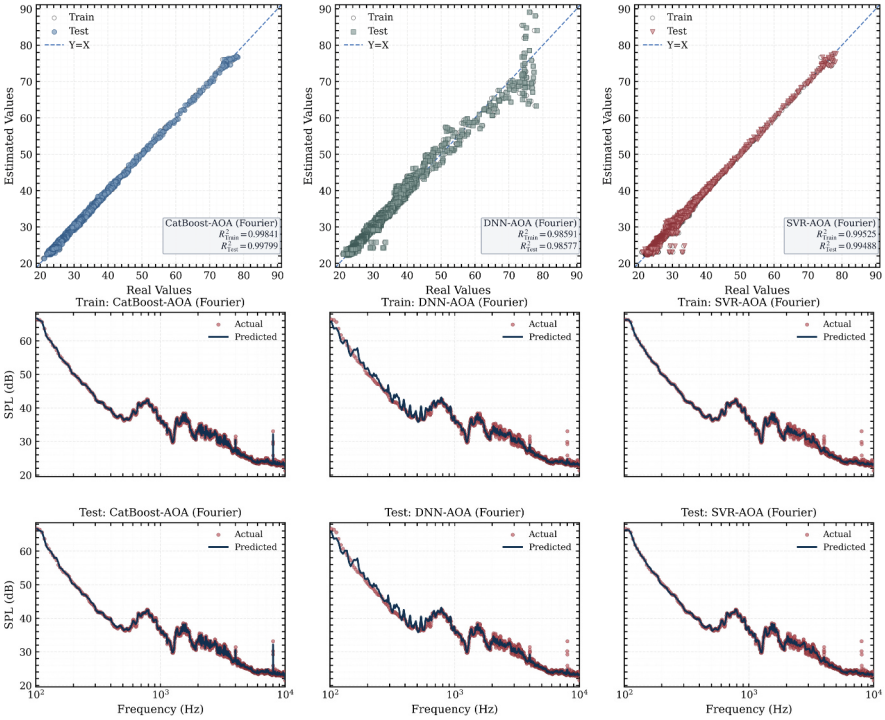

Fig. 7. Parity plots of baseline hybrid models with Fourier-enhanced features (Top) and comparison of measured and predicted SPL spectra obtained from the same models (Bottom)

With the introduction of Fourier-enhanced frequency features, the models exhibited different responses depending on the architecture. Logarithmic frequency scaling combined with harmonic sine and cosine terms provided a structured representation of frequency variation that improved the representation of tonal content relative to the baseline formulation. In the cases of the CatBoost-AOA and SVR-AOA, this yielded better agreement with measured tonal peaks, while preserving broadband trends, reflected by lower RMSE, MAE, and RAE and higher R^2 values. Although the DNN-AOA model with Fourier features better encoded the locations of tonal peaks, it introduced localized low frequency oscillations that manifest as nonphysical spectral artifacts. These

oscillations increased overall prediction error, leading to higher RMSE and lower R^2 compared to the baseline DNN formulation despite improved tonal peak visibility. This behavior suggested increased sensitivity of the DNNs architecture to harmonic feature interactions, which degraded global error metrics even when it was able to capture the local spectral features.

Despite the SVR-AOA model produced smoother predictions without such oscillations, its ability to resolve sharp tonal peaks diminished at higher frequencies, indicating limited capacity to model abrupt spectral transitions within the chosen kernel space. Overall, the Fourier-enhanced CatBoost-AOA model provided the most accurate and stable performance, consistently capturing both broadband trends and discrete tonal peaks across the full frequency range, which were supported by strong parity agreement and the lowest error metrics. These results indicated that Fourier based frequency encoding improved tonal representation, and its effectiveness depended strongly on model architecture.

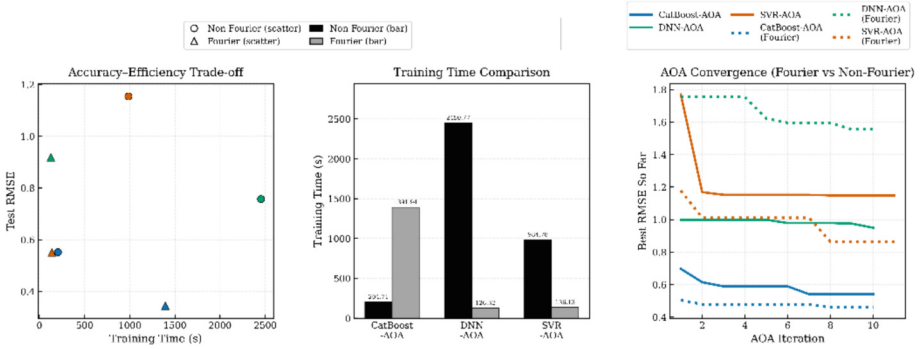

Fig. 8. Accuracy-Efficiency trade off (Left), Training Time Comparison (Middle) and Convergence behavior of the AOA for different learning models (Right)

Furthermore, Fig. 8 collectively evaluated the efficiency, accuracy-efficiency trade off, and convergence behavior of the proposed models. Fourier-enhanced configuration substantially reduced training time for DNN-AOA and SVR-AOA, while increasing computational cost for CatBoost-AOA due to higher complexity of tree-based optimization. In terms of accuracy, CatBoost-AOA (Fourier) achieved the lowest average RMSE, whereas SVR-AOA (Fourier) provided balanced trade-off between accuracy and efficiency. The convergence analysis showed that Fourier features improved the AOA search for CatBoost-AOA and SVR-AOA, while the DNN-AOA exhibited less consistent convergence, indicating higher sensitivity to feature encoding and optimization settings.

4 Conclusion and Future Work

With the integration of Fourier based input, the CatBoost-AOA model proved to be the most robust hybrid model delivering precise and physically consistent predictions, successfully capturing both broadband noise and discrete tonal peaks associated with trailing edge instabilities. Although SVR-AOA offered a competitive compromise between

accuracy and computational efficiency, the CatBoost-AOA demonstrated superior ability to resolve the narrowband spectral details vital to aeroacoustics research. In contrast DNN-AOA exhibited less consistent convergence across runs and produced localized low frequency oscillations that appeared as nonphysical tonal peaks, suggesting that the network was more sensitive to harmonic feature interactions and might require additional regularization or architecture tuning to ensure stable spectral behavior.

To improve resolution and accuracy of the models, more experimental data will be incorporated with wide ranges of stream velocities, chord lengths and clusters of microphones in future research. It will provide higher FFT resolution, yielding finer frequency spacing, that will improve repeatability of weak tonal peak identification and background noise separation. Expanding the Fourier feature with higher order harmonic terms can provide a better frequency representation and capture the multiscale oscillation behaviors. All these expansions are expected to increase the robustness and the generalization of the models to reliably and meaningfully predict the spectrum of airfoil trailing edge noise.

Appendix

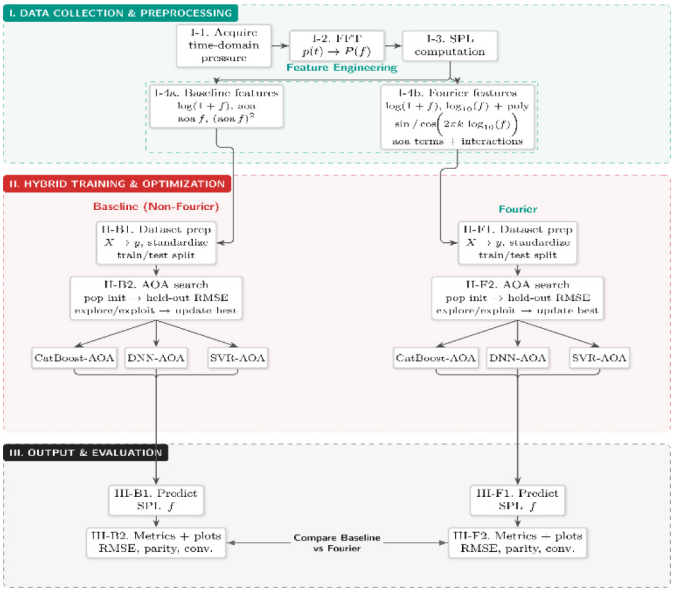

Fig. 9. Overview of the Hybrid AOA-Optimized Machine Learning Framework

References

1. Brooks, T.F., Stuart Pope, D., Marcolini, M.A.: Airfoil self-noise and prediction. Technical Report No. L-16528. NASA, Washington, DC, USA (1989)

2. Rastgoo, A., Khajavi, H.: A novel study on forecasting the airfoil self-noise, using a hybrid model based on the combination of CatBoost and arithmetic optimization algorithm. Expert Syst. Appl. **229**(2023), 120576 (2023)
3. Paterson, R.W., Vogt, P.G., Fink, M.R., Lee Munch, C.: Vortex noise of isolated airfoils. J. Aircraft **10**(5), 296–302 (1973)
4. Arbey, H., Bataille, J.: Noise generated by airfoil profiles placed in a uniform laminar flow. J. Fluid Mech. **134**(1983), 33–47 (1983)
5. Arcondoulis, E., Doolan, C., Zander, A.C.: Airfoil noise measurements at various angles of attack and low Reynolds number. In: Proceedings of ACOUSTICS. Australian Acoustical Society, pp. 23–25 (2009)
6. Redonnet, S., Bose, T., Seth, A., Li, L.K.B.: Airfoil self-noise prediction using deep neural networks. Eng. Anal. Boundary Elem. **159**(2024), 180–191 (2024)
7. Wang, Q., Zhang, H., Wang, X., Ni, Y.: Wind turbine airfoil noise prediction method based on generalized airfoil database and residual neural network. Appl. Sci. **15**(9), 5123 (2025)
8. Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex fourier series. Math. Comput. **19**(90), 297–301 (1965)
9. Oppenheim, A.V.: Discrete-Time Signal Processing. Pearson Education India (1999)
10. Arce León, C., Ragni, D., Pröbsting, S., Scarano, F., Madsen, J.: Flow topology and acoustic emissions of trailing edge serrations at incidence. Exp. Fluids **57**(5), 91 (2016)
11. Prokhorenkova, L., Gusev, G., Vorobev, A., Veronika Dorogush, A., Gulin, A.: CatBoost: unbiased boosting with categorical features. In: Advances in Neural Information Processing Systems (NeurIPS 2018), vol. 31 (2018)
12. Hancock, J.T., Khoshgoftaar, T.M.: CatBoost for big data: an interdisciplinary review. J. Big Data **7**(1), 94 (2020)
13. Cai, Y., Yuan, Y., Zhou, A.: Predictive slope stability early warning model based on CatBoost. Sci. Rep. **14**(1), 25727 (2024)
14. Hao, D., Yang, X., Qi, T.: Hybrid machine learning models based on CatBoost classifier for assessing students' academic performance. Int. J. Adv. Comput. Sci. Appl. **15**, 7 (2024)
15. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge, MA, USA (2016)
16. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature **521**(2015), 436–444 (2015)
17. Hinton, G., et al.: Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. IEEE Signal Process. Mag. **29**(6), 82–97 (2012)
18. Das, L., Sivaram, A., Venkatasubramanian, V.: Hidden representations in deep neural networks: part 2 regression problems. Comput. Chem. Eng. **139**(2020), 106895 (2020)
19. Kim, J., Kim, J., Jang, G.-J., Lee, M.: Fast learning method for convolutional neural networks using extreme learning machine and its application to lane detection. Neural Netw. **87**(2017), 109–121 (2017)
20. Gao, G., Wang, G.: A deep neural network with attention mechanism for flow prediction of compressor blade. Sci. Rep. **15**(1), 16336 (2025)
21. Anderson, P.G., et al.: Deep learning improves physician accuracy in the comprehensive detection of abnormalities on chest X-rays. Sci. Rep. **14**(1), 25151 (2024)
22. Sifaou, H., Kammoun, A., Alouini, M.-S.: A precise performance analysis of support vector regression. In: Proceedings of the International Conference on Machine Learning (ICML 2021), pp. 9671–9680. PMLR (2021)
23. Ho, C.-H., Lin, C.-J.: Large-scale linear support vector regression. J. Mach. Learn. Res. **13**(1), 3323–3348 (2012)
24. Wei, J., He, X.: Support vector regression model with variant tolerance. Measur. Control **56**(9–10), 1705–1719 (2023)
25. Hsu, C.-W., Lin, C.-J.: A comparison of methods for multiclass support vector machines. IEEE Trans. Neural Netw. **13**(2), 415–425 (2002)

26. Kaya, M.: A CFD based application of support vector regression to determine the optimum smooth twist for wind turbine blades. Sustainability **11**(16), 4502 (2019)
27. Safari, M., Hossein Rabiee, A., Joudaki, J.: Developing a support vector regression (SVR) model for prediction of main and lateral bending angles in laser tube bending process. Materials **16**(8), 3251 (2023)
28. Du, W., Cheung, H., Johnson, C.A., Goldberg, I., Thambisetty, M., Becker, K.: A longitudinal support vector regression for prediction of ALS score. In: Proceedings of the 2015 IEEE International Conference on Bioinformatics and Biomedicine (BIBM 2015), pp. 1586–1590. IEEE (2015)
29. Akay, Ö., Tunçeli, M.: Use of the support vector regression in medical data analysis. Exp. Appl. Med. Sci. **2**(4), 242–256 (2021)
30. Abualigah, L., Diabat, A., Mirjalili, S., Elaziz, M.A., Gandomi, A.H.: The arithmetic optimization algorithm. Comput. Methods Appl. Mech. Eng. **376**(2021), 113609 (2021)
31. Dhal, K.G., Sasmal, B., Das, A., Ray, S., Rai, R.: A comprehensive survey on arithmetic optimization algorithm. Arch. Comput. Methods Eng. **30**(5), 3379–3404 (2023)
32. Fang, H., Fu, X., Zeng, Z., Zhong, K., Liu, S.: An improved arithmetic optimization algorithm and its application to determine the parameters of support vector machine. Mathematics **10**(16), 2875 (2022)
33. Feng, C., Cai, M., Luo, J., Li, B.: Arithmetic optimization algorithm improved by multi-strategy fusion. In: Proceedings of the 4th International Conference on Artificial Intelligence, Big Data and Algorithms (AIBDA 2024), pp. 896–901 (2024)
34. Nash, E.C., Lowson, M.V., McAlpine, A.: Boundary-layer instability noise on aerofoils. J. Fluid Mech. **382**(1999), 27–61 (1999)
35. Zajamsek, B., Doolan, C.J., Moreau, D.J., Fischer, J., Prime, Z.: Experimental investigation of trailing edge noise from stationary and rotating airfoils. J. Acoust. Soc. Am. **141**(5), 3291–3301 (2017)
36. Barnard, A.R.: Flow induced noise reduction techniques for microphones. Sound Vibr. **48**, 10 (2014)
37. Al Tlua, B., Rocha, J.: Experimental investigation of NACA-0012 airfoil instability noise with sawtooth trailing edges. Wind Eng. **46**(3), 972–982 (2022)
38. Siddique, M.F., Ahmad, Z., Ullah, N., Kim, J.: A hybrid deep learning approach: integrating short-time Fourier transform and continuous wavelet transform for improved pipeline leak detection. Sensors **23**(19), 8079 (2023)