electricity

MDPI

Article

Fourier Feature-Enhanced Neural Networks for Wind Turbine Power Modeling

Theofanis Aravanis [1,*], Polydoros Papadopoulos [2] and Dimitrios Georgikos [1]

[1] Department of Digital Systems, University of the Peloponnese, 231 00 Sparta, Greece; d.georgikos@go.uop.gr
[2] Department of Mechanical Engineering, University of the Peloponnese, 263 34 Patras, Greece; ppapadop@iceht.forth.gr
* Correspondence: taravanis@uop.gr

Abstract

Accurate prediction of wind turbine power output is essential for optimizing renewable energy generation, enhancing grid integration, and improving the efficiency of wind farms. However, the inherent non-linearities of wind speed–power relationships, combined with abrupt cut-in, rated, and cut-out effects, pose a significant modeling challenge. In this study, we investigate the use of artificial neural networks (ANNs) to model the power curve of a 1 kW wind turbine, using an open-access dataset of real operational measurements recorded at 10 min intervals over the course of 2011. In particular, we compare a conventional multilayer perceptron (MLP) trained on raw wind speed inputs with a Fourier-feature-encoded MLP designed to mitigate spectral bias—the tendency of neural networks to favor smooth, low-frequency patterns over sharp, high-frequency variations. Experimental results show that the Fourier-enhanced MLP substantially improves predictive performance, reducing the mean absolute error (MAE) by more than 65% and achieving an R^2 score of 0.999. The proposed approach demonstrates that Fourier feature encoding enables neural networks to capture sharp non-linearities in wind-turbine power curves, representing one of the first applications of this technique to wind-turbine power-curve modeling.

Keywords: wind turbine power curve; renewable energy sources; data-driven modeling; machine learning; artificial neural networks; Fourier feature encoding

MSC: 68T01; 68T07; 68T09

check for updates

Academic Editor: Andreas Sumper

Received: 9 October 2025
Revised: 12 November 2025
Accepted: 28 November 2025
Published: 1 December 2025

Citation: Aravanis, T.; Papadopoulos, P.; Georgikos, D. Fourier Feature-Enhanced Neural Networks for Wind Turbine Power Modeling. *Electricity* **2025**, *6*, 70. https://doi.org/10.3390/electricity6040070

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The transition toward sustainable energy systems has placed renewable resources at the forefront of global strategies to mitigate climate change and reduce dependence on fossil fuels. Among them, wind energy has emerged as one of the most mature and scalable technologies, with installed capacity expanding rapidly worldwide [1]. Wind turbines, which convert the kinetic energy of moving air into electrical power, represent the cornerstone of this transition. Accurately modeling their performance is crucial for optimizing energy production, facilitating grid integration, and ensuring the reliable operation of wind farms.

Nevertheless, predicting wind turbine power output remains a challenging task. The power curve is inherently non-linear and influenced by multiple meteorological, geographical, and mechanical factors. Classical physics-based formulations, such as Betz's law, provide a theoretical upper bound but fail to capture the variability observed in real-world conditions. This complexity has motivated the adoption of data-driven approaches.

In particular, machine learning (ML) methods have attracted increasing attention due to their ability to model non-linear relationships and learn directly from empirical data [2,3].

Among ML techniques, artificial neural networks (ANNs) [4–6] stand out as universal function approximators [7–9], capable of capturing intricate dependencies between input and output variables. However, standard multilayer perceptrons (MLPs)—feedforward fully-connected ANNs—often suffer from *spectral bias*, which limits their ability to approximate rapidly varying or oscillatory functions [10,11]. This limitation is particularly relevant in wind energy modeling, where abrupt transitions occur near the cut-in, rated, and cut-out wind speeds of a turbine.

To address this challenge, we propose using *Fourier feature encoding* (FFE) as a means to enrich the input representation of wind speed measurements. By projecting scalar inputs into a high-dimensional sinusoidal space, FFE alleviates spectral bias and enables MLPs to better approximate complex, non-linear power curves. Our approach builds upon the influential work of Tancik et al. [12], who demonstrated that Fourier feature mappings substantially improve the ability of MLPs to learn high-frequency functions in low-dimensional domains. While their study focused mainly on computer vision and graphics applications, we *extend* this concept to the *energy domain*, where, to the best of our knowledge, the use of input-level random Fourier feature encodings with simple MLPs has received little to no prior attention.

Specifically, we present a systematic comparison between a conventional plain-input MLP and an FFE-enhanced MLP, both trained on the real-world, open-source dataset of a 1 kW wind turbine, provided by the Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development (GECAD) [13]. Our experiments show that Fourier-based encoding yields a substantial improvement in predictive accuracy, reducing the mean absolute error (MAE) by more than 65%, and achieving an R^2 score of 0.999 on the *validation* dataset. The obtained results highlight the effectiveness of FFE in overcoming the spectral bias of standard MLPs, and establish it as a novel methodology for accurate wind turbine power prediction.

The remainder of this article is organized as follows. Section 2 reviews prior work on the application of ANNs to wind-turbine power-curve modeling, as well as recent advances in integrating FFE within neural network architectures. Section 3 provides an overview of wind-turbine operation. Section 4 details the research methodology, including the dataset characteristics and preprocessing steps. Section 5 presents the architecture of the predictive models developed in this study—namely, a baseline MLP with plain wind speed input and an FFE-enhanced MLP designed to mitigate spectral bias. Section 6 reports and interprets the experimental outcomes, comparing the predictive performance of the two approaches. Finally, Section 7 summarizes the main findings and outlines potential directions for future research and model generalization.

2. Related Work

This section situates our work within two interconnected strands of research, namely, the extensive use of ANNs for wind-turbine power-curve modeling, and the recent advances in FFE as a means to mitigate spectral bias in neural networks.

A substantial body of wind energy research has established ANNs as strong models for turbine power-curve estimation. Compared with manufacturer S-curves and other parametric fits, ANN-based power curves capture site-specific non-linearities and operational effects more accurately, as summarized by comprehensive reviews of power-curve methodologies and applications [14]. Pelletier, Masson, and Tahan, for example, trained ANN power-curve models with multiple meteorological/operational inputs (wind speed, air density, turbulence intensity, shear, etc.) and reported markedly lower absolute and ran-

dom errors than parametric baselines [15]. Manobel et al. combined Gaussian process (GP) prefiltering with ANN regression for power-curve modeling and achieved further accuracy gains over GP-only or ANN-only approaches [16]. Extending beyond one-dimensional curves, Yan et al. proposed two-dimensional (speed-direction) ANN power curves at the wind-farm level, reducing bias and improving accuracy relative to one-dimensional mappings [17]. More recently, explainable-AI frameworks tailored to data-driven power-curve models have been introduced to evaluate physical plausibility and robustness, addressing transparency concerns in operational settings [18].

Despite these advances, most ANN-based power-curve models rely on raw wind speed inputs, which, as will be discussed in detail, can limit their ability to represent rapid, high-frequency, or sharply varying patterns present in real measurement data. Such behaviors are precisely where standard MLPs tend to struggle, due to their inherent spectral bias—i.e., their tendency to learn smooth, low-frequency relationships more easily than high-frequency or rapidly changing behaviors [10,11]. This challenge motivates the incorporation of Fourier-based input encodings.

A first, influential use of fixed sinusoidal bases in deep learning appears in the Transformer models. In particular, Vaswani et al. introduced sinusoidal positional encodings to inject token order directly into the representation space [19]. Although not framed as "Fourier features", these deterministic sine–cosine embeddings effectively instantiate a Fourier basis and demonstrate how frequency-structured priors can improve learning and generalization. Motivated by such observations, Rahaman et al. provided empirical and theoretical evidence that standard MLPs exhibit spectral bias, fitting low-frequency components far more readily than high-frequency content [10]. This clarified why naïve-coordinate-based networks blur fine detail and why explicit frequency encodings might help.

Against this backdrop, Tancik et al. crystallized Fourier feature mappings for coordinate-based MLPs, showing that passing inputs through multifrequency sinusoidal embeddings expands a network's effective bandwidth and counters spectral bias [12]. With these encodings, plain MLPs reliably capture high-frequency structure in low-dimensional domains that they otherwise fail to learn.

A prominent application is neural radiance fields (NeRF), where an MLP represents a continuous radiance field for novel-view synthesis [20]. NeRF's positional (Fourier-style) encodings of 3D coordinates and viewing directions are crucial, as they enable the network to capture high-frequency variations in appearance and geometry. Without these encodings, the model produces blurred structures and textures, whereas their inclusion allows it to reconstruct sharp edges and fine scene details with photorealistic accuracy.

Beyond vision, Fourier features have proven valuable in modeling oscillatory physics. Jandrell and Cox proposed a data-efficient framework for perturbation-dependent optical transmission in multimode fibers, encoding physical perturbations with a Fourier feature basis and training a compact MLP to learn highly oscillatory mappings [21]. On experimental fiber data, their model attained a complex correlation of 0.995 to ground truth, while using 85% fewer parameters than conventional baselines, underscoring how explicit frequency bases stabilize optimization and improve fidelity in real systems.

In the domain of physics-informed neural networks (PINNs), designed for problems governed by partial differential equations [22], incorporating Fourier features has proven effective in capturing multiscale solutions [23]. For instance, Wang et al. showed that Fourier feature MLPs alleviate the spectral bias in PINNs, enabling them to solve high-frequency or multiscale problems much more accurately [24].

Taken together, while ANNs are well established for power-curve modeling, Fourier feature encoding has not yet been explored in this domain, leaving an open opportunity

for further improvement. This is precisely where our contribution lies, as outlined in the sections that follow.

3. Wind Turbines

Before presenting the core of the proposed methodology, we briefly outline how a modern horizontal-axis wind turbine converts wind energy into electricity and why the resulting power curve exhibits a highly non-linear, piecewise structure. A wind turbine transforms the kinetic energy of moving air into electrical power through a sequence of aerodynamic, mechanical, and electrical processes. The rotor blades extract energy from the wind and transfer it through the drivetrain to the generator, while control systems regulate this conversion to maintain efficient and safe operation [25].

Structurally, a wind turbine comprises several interdependent components, each affecting the overall power output. The tower elevates the rotor to higher altitudes where wind speeds are stronger and steadier, thereby increasing the kinetic energy available for conversion. Mounted atop the tower, the nacelle houses the gearbox, generator, and control systems that together transform mechanical rotation into electricity. The rotor blades, typically made of lightweight composite materials, are responsible for capturing the wind's kinetic energy; their aerodynamic design determines how effectively this energy is converted into torque. The gearbox adjusts the low-speed, high-torque rotation of the rotor into the high-speed rotation required by the generator, which converts mechanical energy into electrical energy through electromagnetic induction. Meanwhile, control systems—including pitch and yaw mechanisms—continuously optimize the orientation and rotational speed of the blades to maximize energy capture under varying wind conditions, while preventing mechanical overload. Finally, the anemometer and wind vane provide real-time measurements of wind speed and direction, enabling the control systems to adaptively adjust turbine operation for efficiency and safety [26].

The electrical power output P of a wind turbine depends on the air density, rotor swept area, and wind speed, as well as on the turbine's aerodynamic, mechanical, and electrical efficiencies. This relationship is expressed by

$$P = \eta_E \cdot \eta_M \cdot C_p \cdot P_{wind} = \eta_E \cdot \eta_M \cdot C_p \cdot \left(\frac{1}{2} \cdot \rho \cdot A \cdot v^3\right), \tag{1}$$

where ρ is the air density (kg/m^3), A is the swept area of the rotor blades (m^2), and v is the wind speed (m/s). The term C_p denotes the power coefficient, a dimensionless measure of how efficiently the rotor converts the wind's kinetic power ($P_{wind} = \frac{1}{2} \cdot \rho \cdot A \cdot v^3$) into mechanical power. According to Betz's law, the theoretical upper limit of C_p is 0.593, meaning that no turbine can capture more than 59.3% of the wind's kinetic energy. In practice, however, aerodynamic imperfections, frictional losses, and wake effects reduce this value to typically between 0.3 and 0.4, depending on the rotor geometry, airfoil design, and operating conditions. The parameters η_M and η_E denote the mechanical and electrical efficiencies, representing drivetrain losses (bearings, shafts, gearbox) and conversion losses in the generator and power electronics, respectively.

Equation (1) is only valid for a *certain range* of wind speeds, as the electrical power that can finally be generated is *at most equal* to the rated electrical power of the generator. Thus, when the wind speed is large enough that the electrical power calculated from Equation (1) is greater than the aforementioned rated electrical power, the electrical power actually produced is equal to that rated value. Some important values of wind speed, related to the power generation of the wind turbine, are also noted: Firstly, the *cut-in speed* is the minimum wind speed (typically 3–4 m/s) at which the wind turbine starts to produce mechanical power, as it is hindered by friction and other mechanical losses. Secondly,

the *rated speed* is the wind speed at which the generator starts to produce the rated (i.e., maximum) electrical power. Lastly, from the rated, up to a wind speed called *cut-out speed*, the generator limits its output to the rated electrical power. For speeds greater than the cut-out speed, power generation is interrupted to protect the installation. Against this background, and for constant values for η_E, η_M, C_p, ρ, and A, the *power curve* of a wind turbine, in view of Equation (1), is depicted in Figure 1.

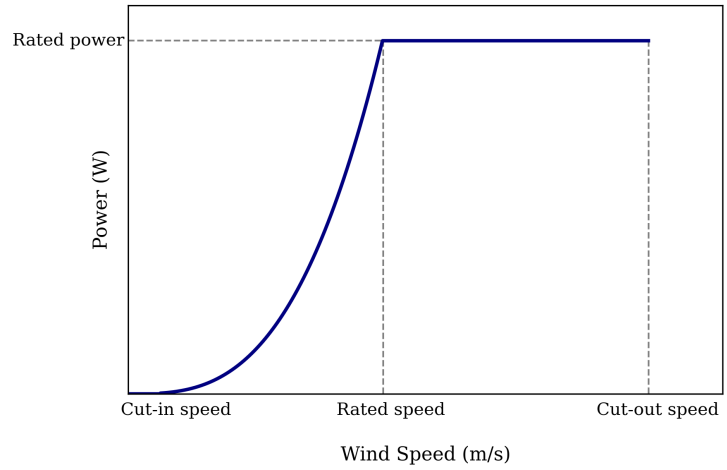

Figure 1. Typical power curve of a wind turbine derived from Equation (1).

As indicated by Equation (1) and illustrated in Figure 1, between the cut-in speed and the rated speed, the generated electrical power P grows as a cubic function of the wind speed v, before saturating at the rated power and dropping abruptly at the cut-out speed. This combination of smooth cubic growth and abrupt transitions in the power curve leads to a highly complex power production pattern that is difficult to approximate with conventional linear models or plain-input neural networks. Such non-linearities provide the main motivation for employing advanced input encodings, such as FFE, which, as we will demonstrate, enrich the input space with sinusoidal components and allow neural networks to better capture the sharp, piecewise behavior of real turbine power curves.

4. Methodology

The methodology adopted in this work is presented in this section, covering the research design, data profile, and preprocessing procedures.

4.1. Research Design

The research design of this study follows a structured, stepwise methodology aimed at developing, training, and evaluating neural network models for wind turbine power prediction. The workflow begins with the selection and preparation of a real-world dataset containing synchronized wind speed and power measurements from a 1000 W turbine. Following data acquisition, the dataset undergoes a standardized preprocessing pipeline that includes cleaning and chronological splitting. Particular attention is given to preventing temporal leakage by ensuring that all preprocessing operations are fitted exclusively on the training period.

The study then proceeds with the development of two predictive models, namely, a baseline MLP using raw wind speed inputs, and an enhanced MLP that incorporates FFE. Model training, hyperparameter selection, and early-stopping strategies are applied consistently across both architectures, enabling a controlled comparison. The evaluation phase uses a strictly held-out, chronologically later portion of the dataset to assess generalization performance under realistic operational conditions.

4.2. Data Profile

For the present study, we employ an open-source, real-world dataset provided by the Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development (GECAD) [13]. The dataset contains measurements of wind speed and corresponding electrical power output from a 1000 W rated wind turbine, recorded throughout the year 2011. Each entry consists of a pair (v, P), where v is the wind speed (in m/s) and P is the electrical power generated (in W), sampled at 10 min intervals. An excerpt of the dataset is shown in Table 1, while Figure 2 illustrates the time series of both the turbine's electrical power output and the corresponding wind speed over a representative portion of the dataset.

Table 1. Excerpt from the GECAD dataset. Each row contains a timestamped measurement of wind speed (v) and the corresponding electrical power output (P) of the 1 kW wind turbine.

Timestamp	Wind Speed (m/s)	Power Output (W)
01-01-2011 00:00:00	5.26	103.34
01-01-2011 00:10:00	5.5	114.83
01-01-2011 00:20:00	4.76	79.42
01-01-2011 00:30:00	4.94	88.03
⋮	⋮	⋮

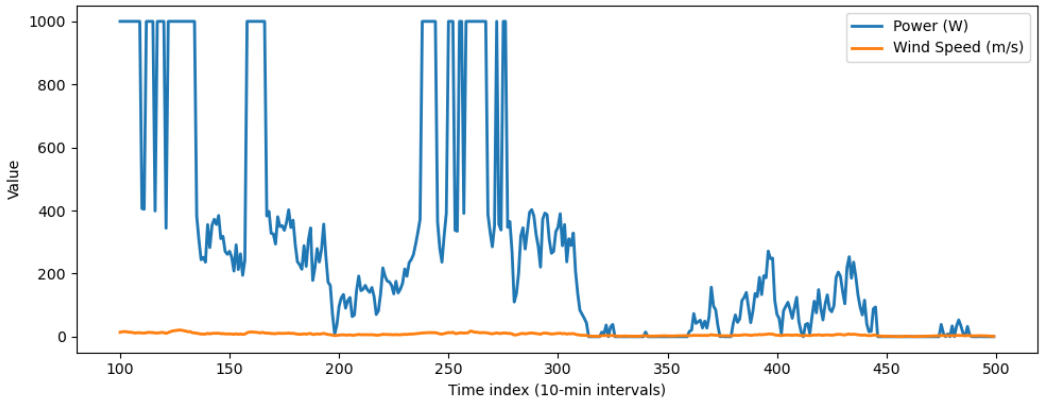

Figure 2. Time series of the wind turbine's electrical power output (P) and wind speed (v) over a representative segment of the GECAD dataset, sampled at 10 min intervals.

The dataset covers the entire year 2011 at 10 min intervals, yielding a total of 50,364 samples. Of these, the earliest 35,254 samples (approximately 70%) were used for training, while the most recent 15,110 samples (approximately 30%) formed the held-out validation set. This chronological separation ensures that no information from the future leaks into the training process and provides a realistic assessment of predictive performance.

5. Design and Configuration of the Models

In this section, we describe the modeling approach, outlining the structure, training procedure, and key design choices of both the plain-input and Fourier-enhanced ML models.

5.1. Multilayer Perceptron with Plain Input

The baseline ML model considered for predicting wind turbine power output is a standard multilayer perceptron (MLP) [6]. The network architecture comprises four layers,

namely, one input layer, two fully connected (dense) hidden layers, and one output layer (Figure 3). The input layer consists of a single linear unit that receives plain wind speed values, after min–max normalization into the range $[0, 1]$. Each hidden layer contains 128 neurons with rectified linear unit (ReLU) activations, providing the non-linear capacity required to approximate the power curve. The output layer consists of a single linear unit that produces an estimate of the normalized electrical power, which is subsequently inverse-transformed to Watts for evaluation.

The network was trained using the Adam optimization algorithm [27], with a mean squared error (MSE) loss function, and a maximum of 100 training epochs. To avoid overfitting and ensure robust convergence, an early stopping criterion was applied with a patience of 10 epochs, restoring the best-performing weights. All experiments were repeated with fixed random seeds to guarantee reproducibility.

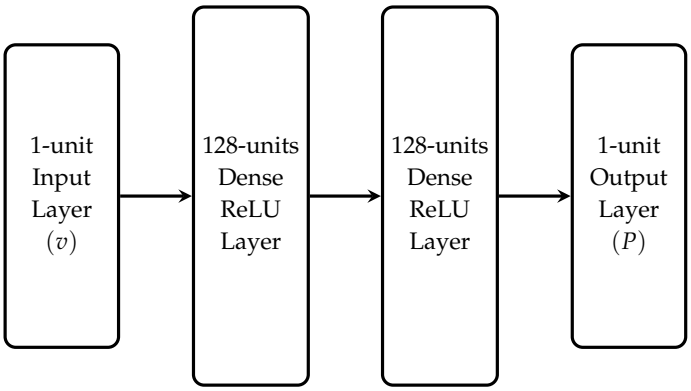

Figure 3. Architecture of the plain-input MLP.

5.2. Multilayer Perceptron with Fourier Feature Encoding

Standard MLPs exhibit what is known as *spectral bias*, namely, they tend to learn low-frequency functions more easily than high-frequency ones, which leads to poor representation of rapidly varying signals. This limitation has been extensively studied in the Neural Tangent Kernel (NTK) framework, which shows that the eigenvalue spectrum of a conventional MLP decays quickly with frequency, resulting in very slow convergence of high-frequency components [10,11]. As a consequence, MLPs trained directly on scalar inputs, such as wind speed, may fail to capture the sharp non-linearities that characterize the wind speed–power curve of a real wind turbine.

To overcome this limitation, a line of work beginning with Rahimi and Recht [28] proposed the use of *random Fourier features* to approximate stationary kernels (i.e., similarity functions that depend only on the difference between inputs) via sinusoidal mappings. Building on this idea, Tancik and colleagues [12] adapted Fourier features to coordinate-based MLPs, showing that this encoding alleviates spectral bias by making the effective NTK stationary and tunable, thereby enabling faster learning of high-frequency components and improving generalization across low-dimensional regression tasks. Specifically, each scalar input x is transformed as

$$\gamma(x) = \left[\sin(2\pi x B), \cos(2\pi x B)\right], \tag{2}$$

where B is a row vector of randomly sampled frequencies. This mapping enriches the input domain with harmonics at multiple scales, allowing the MLP to better approximate complex non-linear functions.

Originally developed for coordinate-based neural representations in computer vision and graphics, Fourier features have been shown to dramatically improve the accuracy of neural networks in representing images, 3D shapes, and volumetric data [12]. In the

present work, we transfer this methodology to the energy domain by encoding raw wind speed values using random Fourier features before passing them into the MLP. To the best of our knowledge, such an approach has not been previously investigated for turbine power-curve modeling. As our results demonstrate, the proposed method enables the network to capture the non-linear relationship between wind speed and power output more effectively, yielding improved predictive performance compared to models trained only on normalized raw inputs.

In our implementation, the Fourier feature mapping is applied directly to the raw wind speed samples. Let $X = (x_1, \ldots, x_N)^\top \in \mathbb{R}^{N\times 1}$ denote the vector of training inputs, where each $x_i \in \mathbb{R}$ is a single wind speed value. We first compute their empirical standard deviation $s_x = \mathrm{std}(X)$, which is then used to scale the distribution from which the random frequencies are drawn. A frequency row vector $B = (B_1, \ldots, B_D) \in \mathbb{R}^{1\times D}$ is sampled with independent entries

$$B_j \sim \mathcal{N}\left(0, \left(\frac{\sigma}{s_x}\right)^2\right), \qquad j = 1, \ldots, D, \tag{3}$$

that is, each B_j is drawn from a normal distribution with mean zero and variance $(\sigma/s_x)^2$. Here D denotes the number of Fourier features and σ is a tunable scale parameter. Each input x_i is then projected as

$$P_{ij} = 2\pi\, x_i\, B_j, \qquad i = 1, \ldots, N, \;\; j = 1, \ldots, D, \tag{4}$$

yielding a projection matrix $P \in \mathbb{R}^{N\times D}$. The Fourier feature embedding of x_i is defined as

$$\gamma(x_i) = \big(\sin(P_{i1}), \ldots, \sin(P_{iD}),\ \cos(P_{i1}), \ldots, \cos(P_{iD})\big) \in \mathbb{R}^{2D}. \tag{5}$$

Stacking these row vectors produces the encoded design matrix

$$\Gamma = \begin{bmatrix} \gamma(x_1)^\top \\ \vdots \\ \gamma(x_N)^\top \end{bmatrix} \in \mathbb{R}^{N\times 2D}. \tag{6}$$

In practice, based on preliminary experimentation across several candidate values ($D \in \{8, 16, 32, 64\}$), we found that $D = 32$ provides the best balance between predictive accuracy and model complexity, yielding stable performance improvements without unnecessarily increasing the input dimensionality. With $D = 32$, each scalar wind speed input is expanded into a 64-dimensional vector of sinusoidal features. The same frequency vector B sampled during training is reused when encoding the validation data, ensuring consistency across both phases.

To ensure a fair comparison, the MLP following the Fourier feature encoding employs the same architecture as the plain-input network described in Section 5.1. In particular, it consists of one input layer, two fully connected hidden layers with 128 neurons each and ReLU activations, and a final linear output neuron. Training was again performed with the Adam optimizer, using an MSE loss function, and a maximum of 100 epochs. Early stopping with a patience of 10 epochs was applied to prevent overfitting, and the same fixed random seeds were used to ensure reproducibility. The resulting architecture is illustrated in Figure 4.

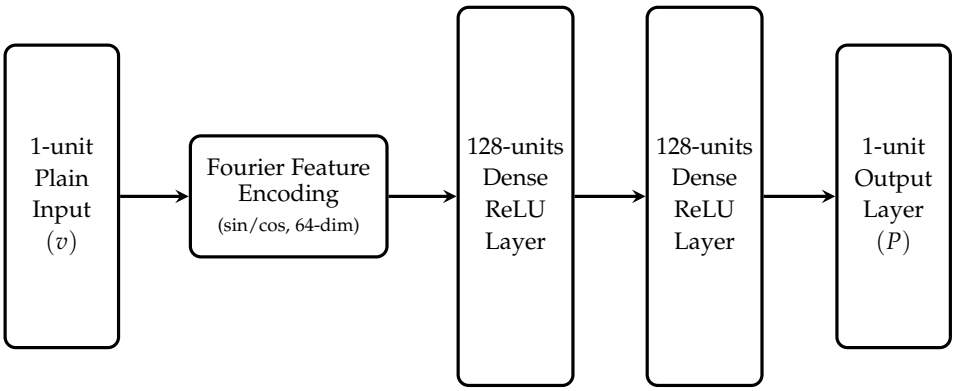

Figure 4. Architecture of the MLP with FFE.

6. Results and Discussion

Having introduced the architecture and parameters of the developed ANNs, this section presents and interprets the experimental results. Accordingly, Table 2 summarizes the error metrics obtained for both models on the *validation* dataset. Recall that the dataset was split chronologically, with the first 70% of samples used for training and the remaining 30% reserved as a held-out validation set, ensuring a leakage-free and operationally realistic evaluation. On that basis, the plain-input MLP achieves an R^2 of 0.991 and a mean absolute error (MAE) of 8.867 W, indicating that while it captures the overall trend of the power curve, residual deviations persist (as shall be illustrated subsequently). By contrast, the FFE-MLP delivers substantially better performance, reducing the MAE to only 2.917 W and lowering the mean squared error (MSE) by nearly an order of magnitude. Its R^2 score of 0.999 confirms an almost perfect fit to the true power values.

Table 2. Performance metrics of the two MLP models on the validation dataset.

Machine Learning Model	MSE (W^2)	MAE (W)	R^2
MLP with Plain Input	1543.431	8.867	0.991
MLP with FFE Input	183.026	2.917	0.999

The training dynamics of the developed models are illustrated in Figure 5, which shows the evolution of the validation MAE (computed on normalized outputs) across epochs for both models. It is evident that the plain-input MLP converges slowly and stabilizes at a higher validation error, while the Fourier feature MLP exhibits both faster convergence and consistently lower validation error. These observations are fully consistent with the error metrics of Table 2, where the Fourier-based model clearly outperforms the plain-input baseline.

Thereafter, Figure 6a illustrates the temporal behavior of the predicted power compared to the actual (ground truth) output in a representative segment of the validation dataset. The plain-input MLP (orange) succeeds in capturing the general trend of the power curve, but noticeable deviations arise, especially at the transition points between low- and high-power regions. In contrast, the MLP enriched with FFE (green) aligns almost perfectly with the true signal, exhibiting significantly reduced error and enhanced stability.

A complementary perspective is provided in Figure 6b, which presents the scatter plot of estimated power against wind speed. The plain MLP (orange) approximates the power curve reasonably well, but introduces discrepancies around the cut-in, rated, and cut-out regions. These regions are characterized by abrupt non-linearities that plain-input networks tend to smooth out. On the other hand, the FFE-MLP (green) almost exactly overlaps with the ground truth, showing only minor deviations at extreme wind speeds.

In particular, a few scattered predictions appear above the cut-out region, which are inconsistent with turbine physics, as power generation must be strictly zero beyond the safety threshold. This can be remedied by applying a simple corrective rule, according to which the predicted power values are clamped to zero slightly beyond the nominal cut-out speed—acknowledging that this threshold is not precisely defined and may vary across turbine models and operating conditions. With this adjustment, the Fourier-based model remains both highly accurate and physically consistent. This confirms that Fourier features effectively mitigate spectral bias, enabling the model to learn both the cubic relationship of Equation (1) and the sharp discontinuities that define the turbine's power curve.

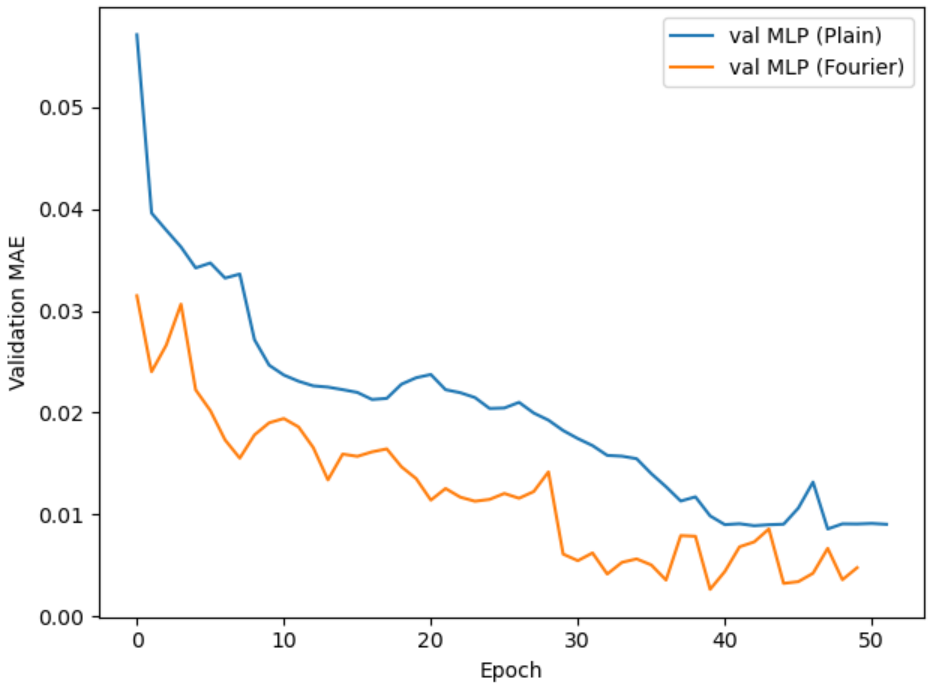

Figure 5. Validation MAE learning curves (normalized outputs) of the two MLP models.

The demonstrated improvement achieved by FFE is not only of methodological interest, but also of practical importance for renewable energy applications. Accurate prediction of wind turbine power output directly affects energy scheduling, reserve allocation, and grid stability. By substantially reducing prediction error, the FFE-MLP model enables operators to anticipate output fluctuations more reliably, particularly around the cut-in, rated, and cut-out speeds where abrupt changes occur. In addition, enhanced modeling accuracy facilitates more efficient turbine control strategies, which rely on precise characterization of turbine behavior. Equally importantly, the improved representation of sharp non-linear transitions enables a more reliable design of protection curves, alarm thresholds, and shut-down logic. As protection relays and power-electronic controllers depend on timely detection of overspeed and overproduction events, a model that captures these high-frequency variations reduces nuisance trips and improves the selectivity of alarms. This leads to more stable operation during gusts, turbulence bursts, or sudden grid disturbances, and supports better event logging in supervisory control systems. The approach also holds promise for predictive maintenance, as deviations between predicted and actual power could serve as early indicators of mechanical or aerodynamic faults.

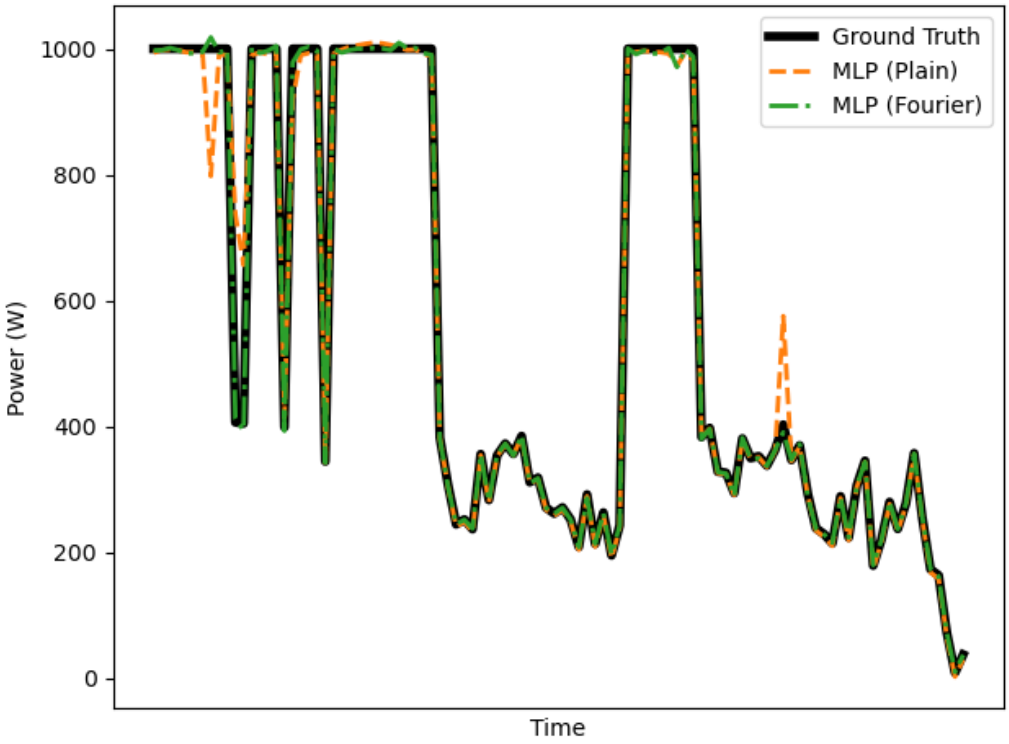

(**a**) Predicted versus actual power in a representative time series segment.

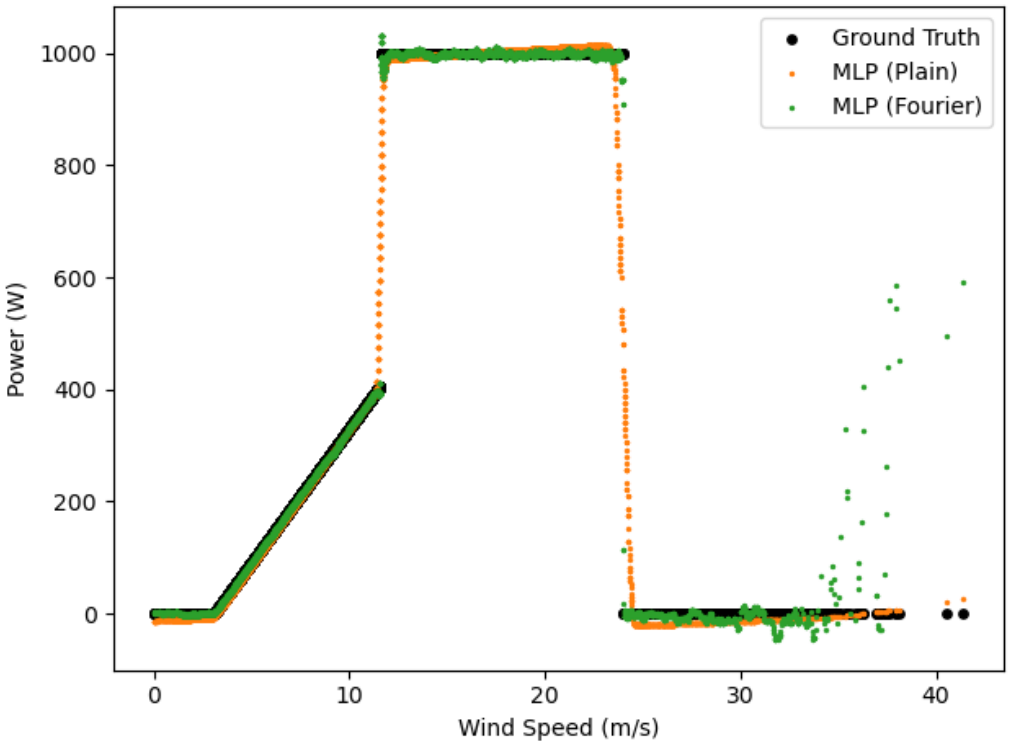

(**b**) Scatter plot of predicted versus actual power across wind speeds.

Figure 6. Comparison of the plain-input MLP and the FFE-MLP on the validation dataset.

Beyond these operational benefits, the proposed approach also highlights a broader innovation pathway for wind energy systems. By improving the fidelity with which high-frequency events and abrupt transitions are captured, Fourier-enhanced models support more resilient plant operation, enabling turbines to maintain stable production even under

rapidly changing atmospheric or grid conditions. This improved robustness reduces the likelihood of unnecessary curtailments, mitigates cascading effects during disturbances, and contributes to maintaining continuous power delivery from converter-dominated renewable fleets. As wind penetration increases and system operators place greater emphasis on reliability, inertia emulation, and fault-ride-through compliance, accurate modeling of fast dynamic behavior becomes a key enabler for secure integration. In this context, the FFE-MLP framework represents not only a modeling refinement, but also a step toward next-generation forecasting and control strategies capable of enhancing production efficiency, reducing operational risk, and strengthening overall grid stability.

7. Conclusions

This work presented a comparative study of artificial neural networks for modeling wind turbine power generation, with a particular focus on the effect of Fourier feature encoding (FFE) on predictive performance. Using a real-world dataset of a 1 kW wind turbine, we evaluated two multilayer perceptron (MLP) architectures, that is, one trained directly on raw wind speed inputs, and one trained on Fourier-encoded inputs.

The obtained results clearly indicate that Fourier features substantially enhance the representational power of MLPs. While the plain-input MLP was able to capture general trends, it struggled with the highly non-linear regions of the wind speed–power curve, especially near the cut-in, rated, and cut-out speeds. By contrast, the Fourier-enhanced MLP achieved nearly perfect alignment with the ground truth, as demonstrated by both visual comparisons and quantitative metrics (mean absolute error reduced by more than 65%, R^2 close to unity).

While our experiments focus specifically on wind-turbine power modeling, the underlying mechanism—Fourier feature inputs mitigating spectral bias—suggests broader applicability to regression tasks in the energy domain that involve oscillatory or sharply non-linear input–output relationships. This includes other renewable energy applications such as photovoltaic generation modeling, where non-linearities and threshold effects similarly challenge standard machine learning models. Future work will extend the present single-turbine framework to a multiturbine collaborative model, explicitly incorporating farm-level interactions such as turbine spacing, wake effects, relative positioning, and intra-farm turbulence patterns. We also aim to integrate additional meteorological variables relevant at the wind farm scale (such as turbulence intensity, air density variations, and directional shear), and to systematically investigate the impact of hyperparameter choices (e.g., number of Fourier features, frequency scaling, and network depth) on model performance. In addition, we plan to explore hybrid architectures that combine Fourier features with symbolic regression for enhanced interpretability [29].

Author Contributions: Conceptualization, T.A. and P.P.; methodology, T.A.; software, T.A.; validation, T.A.; formal analysis, T.A. and D.G.; investigation, T.A.; resources, T.A.; data curation, T.A.; writing—original draft, T.A.; writing—review and editing, P.P., D.G. and T.A.; visualization, T.A.; project administration, T.A. 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 presented in this study are openly available from the Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development (GECAD) [13]. No new data were generated in this work.

Acknowledgments: The authors extend their appreciation to the anonymous reviewers for their valuable and constructive feedback, which has significantly enhanced the quality of this article.

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

References

1. Qazi, A.; Hussain, F.; Rahim, N.A.; Hardaker, G.; Alghazzawi, D.; Shaban, K.; Haruna, K. Towards sustainable energy: A systematic review of renewable energy sources, technologies, and public opinions. *IEEE Access* **2019**, *7*, 63837–63851. [CrossRef]
2. Rashid, H.; Haider, W.; Batunlu, C. Forecasting of wind turbine output power using machine learning. In Proceedings of the 10th International Conference on Advanced Computer Information Technologies (ACIT 2020), Deggendorf, Germany, 16–18 September 2020; pp. 396–399.
3. Clifton, A.; Kilcher, L.; Lundquist, J.K.; Fleming, P. Using machine learning to predict wind turbine power output. *Environ. Res. Lett.* **2023**, *8*, 024009. [CrossRef]
4. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [CrossRef] [PubMed]
5. Schmidhuber, J. Deep learning in neural networks: An overview. *Neural Netw.* **2015**, *61*, 85–117. [CrossRef] [PubMed]
6. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; The MIT Press: Cambridge, MA, USA, 2016.
7. Hornik, K.; Stinchcombe, M.; White, H. Multilayer feedforward networks are universal approximators. *Neural Netw.* **1989**, *2*, 359–366. [CrossRef]
8. Sonoda, S.; Murata, N. Neural network with unbounded activation functions is universal approximator. *Appl. Comput. Harmon. Anal.* **2017**, *43*, 233–268. [CrossRef]
9. Shaham, U.; Cloninger, A.; Coifman, R.R. Provable approximation properties for deep neural networks. *Appl. Comput. Harmon. Anal.* **2018**, *44*, 537–557. [CrossRef]
10. Rahaman, N.; Baratin, A.; Arpit, D.; Dräxler, F.; Lin, M.; Hamprecht, F.A.; Bengio, Y.; Courville, A.C. On the spectral bias of neural networks. In Proceedings of the 36th International Conference on Machine Learning (ICML 2019), Long Beach, CA, USA, 9–15 June 2019.
11. Cao, Y.; Fang, Z.; Wu, Y.; Zhou, D.X.; Gu, Q. Towards understanding the spectral bias of deep learning. In Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI 2021), Virtual, 19–27 August 2021; pp. 2205–2211.
12. Tancik, M.; Srinivasan, P.P.; Mildenhall, B.; Fridovich-Keil, S.; Raghavan, N.; Singhal, U.; Ramamoorthi, R.; Barron, J.T.; Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. In Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Virtual, 6–12 December 2020.
13. Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development (GECAD). Open-Source Dataset for Wind-Based Power Generation. 2011. Available online: https://site.ieee.org/pes-iss/data-sets/#wind (accessed on 9 October 2025).
14. Lydia, M.; Kumar, S.S.; Selvakumar, A.I.; Kumar, G.E.P. A comprehensive review on wind turbine power curve modeling techniques. *Renew. Sustain. Energy Rev.* **2014**, *30*, 452–460. [CrossRef]
15. Pelletier, F.; Masson, C.; Tahan, A. Wind turbine power curve modelling using artificial neural network. *Renew. Energy* **2016**, *89*, 207–214. [CrossRef]
16. Manobel, B.; Sehnke, F.; Lazzús, J.A.; Salfate, I.; Felder, M.; Montecinos, S. Wind turbine power curve modeling based on Gaussian processes and artificial neural networks. *Renew. Energy* **2018**, *125*, 1015–1020. [CrossRef]
17. Yan, C.; Pan, Y.; Archer, C.L. A general method to estimate wind farm power using artificial neural networks. *Wind Energy* **2019**, *22*, 1421–1432. [CrossRef]
18. Letzgus, S.; Müller, K.R. An explainable AI framework for robust and transparent data-driven wind turbine power curve models. *Energy AI* **2024**, *15*, 100328. [CrossRef]
19. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS 2017), Long Beach, CA, USA, 4–9 December 2017; pp. 6000–6010.
20. Mildenhall, B.; Srinivasan, P.P.; Tancik, M.; Barron, J.T.; Ramamoorthi, R.; Ng, R. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Lecture Notes in Computer Science, Proceedings of the European Conference on Computer Vision (ECCV), Glasgow, UK, 23–28 August 2020*; Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M., Eds.; Springer: Cham, Switherland, 2020; Volume 12346, pp. 405–421.
21. Jandrell, J.R.; Cox, M.A. High-fidelity prediction of perturbed optical fields using Fourier feature networks. *arXiv* **2025**, arXiv:2508.19751v2.
22. Raissi, M.; Perdikaris, P.; Karniadakis, G. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **2019**, *378*, 686–707. [CrossRef]
23. Xu, Z.Q.J.; Zhang, L.; Cai, W. On understanding and overcoming spectral biases of deep neural network learning methods for solving PDEs. *J. Comput. Phys.* **2025**, *530*, 113905. [CrossRef]
24. Wang, S.; Wang, H.; Perdikaris, P. On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with Physics-Informed Neural Networks. *Comput. Methods Appl. Mech. Eng.* **2021**, *384*, 113938. [CrossRef]
25. Nwaigwe, K.N. Assessment of wind energy technology adoption, application and utilization: A critical review. *Environ. Res. Lett.* **2022**, *19*, 4525–4536. [CrossRef]

26. Anderson, C. *Wind Turbines: Theory and Practice*; Cambridge University Press: Cambridge, UK, 2020.
27. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015), San Diego, CA, USA, 7–9 May 2015.
28. Rahimi, A.; Recht, B. Random features for large-scale kernel machines. In *Proceedings of the Advances in Neural Information Processing Systems*; Platt, J., Koller, D., Singer, Y., Roweis, S., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2007; Volume 20.
29. Kronberger, G.; Burlacu, B.; Kommenda, M.; Winkler, S.M.; Affenzeller, M. *Symbolic Regression*; CRC Press: Boca Raton, FL, USA, 2024.

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