

Rüzgâr Türbini Güç Eğrisinin Modellenmesi: Yapay Sinir Ağları ve Klasik Regresyonun Karşılaştırılması

Mahmud ASİLSOY^{1*}

¹Şırnak Üniversitesi Mühendislik Fakültesi Bilgisayar Mühendisliği Bölümü, Şırnak

¹<https://orcid.org/0000-0003-1718-5075>

*Sorumlu yazar: mahmudasilsoy@gmail.com

Araştırma Makalesi

Makale Tarihi:

Geliş tarihi: 10.06.2025

Kabul tarihi: 26.10.2025

Online Yayınlanma: 16.03.2026

Anahtar Kelimeler:

Rüzgâr türbini güç eğrisi modellemesi
Veriye dayalı yaklaşımlar
Yapay sinir ağı (YSA)
Rüzgâr enerjisi tahmini
Makine öğrenimi

ÖZ

Bir rüzgâr türbininin güç eğrisinin—rüzgâr hızı ile elektriksel güç çıkışı arasındaki fonksiyonel ilişkinin—doğru biçimde modellenmesi, güvenilir rüzgâr enerjisi tahmini ve gelişmiş türbin kontrolü için büyük önem taşımaktadır. Bu çalışmada, iki tamamlayıcı veri kümesi kullanılarak kapsamlı bir karşılaştırmalı analiz gerçekleştirilmiştir: (i) 10 dakikalık ortalamalara dayalı 48.000 gözleme sahip gerçek bir SCADA veri kümesi, ve (ii) National Renewable Energy Laboratory (NREL) tarafından oluşturulmuş, Texas bölgesine ait simüle edilmiş bir rüzgâr türbini veri kümesi. İkinci veri kümesi, 3.6 MW kapasiteli kara tipi bir General Electric türbini için tam yıllık saatlik ölçümler içermektedir. Her iki veri kümesi üzerinde üç veri odaklı model—Doğrusal Regresyon, Dördüncü Dereceden Polinom Regresyon ve Yapay Sinir Ağı (YSA)—uygulanmış, modellerin doğruluk ve genellebilirlik performansları farklı işletim koşullarında karşılaştırılmıştır. YSA modeli, her iki veri kümesinde de en yüksek başarıyı göstermiştir (SCADA için $R^2 = 0.993$ ve $RMSE = 107$ kW; Texas için $R^2 = 0.999$ ve $RMSE = 27$ kW). Doğrusal model, türbinin karakteristik S-şekilli eğrisindeki doğrusal olmayan yapıyı yakalamakta yetersiz kalırken, polinom model doğruluğu artırmakta ancak eğitim aralığı dışında kararsız sonuçlar üretmektedir. Gerçek ve simüle edilmiş verilerin birlikte değerlendirilmesi, YSA tabanlı modellemenin farklı rüzgâr rejimlerinde sağlam ve uyarlanabilir olduğunu ortaya koymaktadır. Sonuçlar, makine öğrenmesi tabanlı yaklaşımların yalnızca klasik regresyon yöntemlerinden daha yüksek doğruluk sağlamakla kalmayıp, aynı zamanda farklı veri kaynakları arasında daha tutarlı tahminler sunduğunu göstermektedir. Bu bulgular, veri odaklı hibrit modellemenin rüzgâr gücü tahmini, performans optimizasyonu ve akıllı türbin kontrol sistemlerinde anormallik tespiti gibi alanlarda önemli potansiyel taşıdığını ortaya koymaktadır..

Modeling the Wind Turbine Power Curve: A Comparison of Neural Networks and Classical Regression

Research Article

Article History:

Received: 10.06.2025

Accepted: 26.10.2025

Published online: 16.03.2026

Keywords:

Wind turbine power curve modeling
Data-driven approaches
Artificial neural network (ANN)
Wind energy forecasting
Machine learning

ABSTRACT

Accurate modeling of a wind turbine's power curve—the functional relationship between wind speed and electrical power output—is essential for reliable wind energy forecasting and advanced turbine control. In this study, we perform a comprehensive comparative analysis using two complementary datasets: (i) a real-world SCADA dataset containing 48,000 records of 10-minute averaged wind speed and power, and (ii) a simulated Texas Wind Turbine dataset generated by the National Renewable Energy Laboratory (NREL), representing a full year of hourly measurements for a 3.6 MW General Electric onshore turbine. Three data-driven approaches—Linear Regression, Polynomial Regression (4th degree), and an Artificial Neural Network (ANN)—are applied

to both datasets to evaluate their accuracy and generalization capability across different operational conditions. The ANN model consistently demonstrates superior performance, achieving the highest coefficient of determination ($R^2 = 0.993$ for SCADA and 0.999 for Texas) and the lowest root mean square error (RMSE = 107 kW and 27 kW, respectively). The linear model fails to capture the inherent non-linearity of the turbine's S-shaped power curve, while the polynomial fit improves accuracy but exhibits instability outside the training range. By incorporating both real and simulated datasets, this study validates the robustness and adaptability of ANN-based modeling for various wind regimes. The results confirm that machine learning approaches not only outperform traditional regression methods in accuracy but also provide more stable predictions across heterogeneous data sources. These findings highlight the potential of hybrid data-driven modeling for enhanced wind power forecasting, performance optimization, and anomaly detection in intelligent turbine control systems.

To Cite: Asilsoy M. Modeling the Wind Turbine Power Curve: A Comparison of Neural Networks and Classical Regression. *Osmaniye Korkut Ata Üniversitesi Fen Bilimleri Enstitüsü Dergisi* 2026; 9(2): 902-921.

1. Introduction

Wind power is a rapidly growing source of renewable energy, with global installed capacity expanding exponentially in recent decades (Pelletier et al., 2016a). A wind turbine's power curve – the relation between the output power and wind speed at hub height – is a fundamental performance characteristic (Sohoni et al., 2016a). It is commonly used for energy yield assessment, warranty guarantees, and real-time performance monitoring of turbines (Sohoni et al., 2016a). In ideal conditions, the theoretical power available in the wind increases as the cube of wind speed, making wind speed the most influential factor in power output (*Wind Energy and Power Calculations | EM SC 470: Applied Sustainability in Contemporary Culture*, n.d.). However, real turbines operate with efficiency limits (e.g. the Betz limit) and control mechanisms that cause the actual power curve to deviate from a simple cubic law. For example, modern turbines are usually pitch-regulated, meaning they optimize blade pitch to maintain a constant rated power above a certain wind speed. Consequently, a typical turbine power curve has a non-linear *S-shape*: near-zero output below the cut-in wind speed (when the turbine is stationary), a steep ramp-up in power in the medium speed range, and a plateau at the turbine's rated power beyond the rated wind speed, until shutdown at the cut-out speed. Environmental factors (air density, turbulence, terrain) can further cause site-specific departures of the power curve from the manufacturer's specification (Lydia et al., 2014; Sohoni et al., 2016a; Saint-Drenan et al., 2019; Lee et al., 2020). This motivates developing accurate data-driven models of the power curve using field measurements.

Traditional approaches to model a turbine's power curve include empirical curve fits and regression techniques. Simpler models assume a predetermined functional form – for instance, a cubic polynomial or logistic function – and estimate parameters to best fit the data. Polynomial regression is a common choice to approximate the non-linear wind speed–power relationship (Villanueva and Feijóo, 2018; Mushtaq et al., 2024; Sebastiani et al., 2024). Models using linear, quadratic, cubic, and higher-order terms of wind speed have been proposed. A higher-order polynomial can flexibly fit the observed data; for example, prior studies have used up to ninth-order polynomials for power curve fitting (Wagner et

al., 2009; Carrillo et al., 2013; Mata et al., 2024). However, purely polynomial models have drawbacks. They do not inherently account for the curve's physical shape constraints – notably the inflection points around rated speed and the asymptotic flattening at maximum power (Villanueva and Feijóo, 2020; Wang et al., 2023). This can lead to unrealistic behavior outside the data range (such as negative power at zero wind or increasing power beyond the rated region). More sophisticated empirical models add piecewise definitions or physics-informed terms (like a cubic term in the below-rated region and a constant above rated to better capture the S-shape. Still, achieving a one-size-fits-all equation for all turbines and conditions is challenging (Sohoni et al., 2016b).

Machine learning approaches offer a flexible alternative by learning the power curve directly from data without assuming a rigid functional form. In particular, artificial neural networks (ANNs) have proven effective for modeling the non-linear relationship between wind speed and power output (Pelletier et al., 2016b, 2016c). An ANN can approximate complex functions by training on historical data, making it well-suited to capture the turbine's true performance characteristics (Pelletier et al., 2016a). Prior work has shown that neural network models often outperform classical regression in power curve estimation accuracy (Sohoni et al., 2016b). Comparative studies have demonstrated that a trained artificial neural network (ANN) model can produce lower error rates across the full range of wind speeds compared to a polynomial regression model (Li et al., 2001; Villanueva and Feijóo, 2018). Moreover, advanced multi-layer neural networks can incorporate additional inputs (e.g. air density, wind direction, turbulence intensity) and still reduce error relative to standard methods (Pelletier et al., 2016a). The main disadvantage of ANNs is the “black-box” nature of the learned model – it can be difficult to interpret the contribution of each neuron or to derive analytical insights (Sohoni et al., 2016b). Nonetheless, for applications like power forecasting where accuracy is paramount, ANNs are attractive. Modern data analytics and big data from wind farms have further facilitated using machine learning for wind power modeling (Karaman, 2023; Yang et al., 2024).

In this study, we develop and compare data-driven models for predicting a wind turbine's power output from wind speed using both classical regression techniques (linear and polynomial) and a modern Artificial Neural Network (ANN). To ensure a comprehensive evaluation, we utilize two complementary datasets: a real-world SCADA dataset containing 48,000 samples of 10-minute averaged measurements, and a simulated Texas Wind Turbine dataset generated by the National Renewable Energy Laboratory (NREL) with one year of hourly observations. These datasets collectively enable the assessment of model robustness under both real and noise-free operating conditions.

The main contributions of this work are as follows: (1) an empirical evaluation of the limitations of linear and polynomial regression methods for modeling the nonlinear S-shaped turbine power curve across different datasets; (2) a demonstration of the performance improvements achieved by ANN models, quantitatively validated through multiple error metrics (R^2 , RMSE, MAE, sMAPE) and visualized in prediction-vs-actual comparisons; and (3) an extended discussion on the practical

implications of accurate power curve modeling for wind power forecasting, energy scheduling, and intelligent turbine control.

By integrating both real and simulated data sources, this paper provides stronger empirical evidence for the superiority of data-driven machine learning approaches over simplified analytical formulations in accurately representing wind turbine behavior and improving forecasting reliability.

2. Methodology

Data Collection and Preprocessing: This study utilizes two complementary datasets representing wind turbine operation under both real-world and simulated conditions (*Texas Wind Turbine Dataset - Simulated*, n.d.; *Wind Turbine Scada Dataset*, n.d.) The SCADA dataset consists of 48,007 records representing approximately one year of 10-minute averaged measurements, typical of supervisory control and data acquisition (SCADA) systems. The Texas dataset provides hourly simulated measurements of wind and power output for a 3.6 MW onshore General Electric turbine, covering a full annual period under noise-free, idealized conditions.

Each dataset includes simultaneous recordings of wind speed (measured or simulated at the turbine hub height) and the corresponding electrical power output. Wind speeds range from 0 to 25 m/s, while power output values extend up to approximately 3.6 MW, consistent with the rated capacity of multi-megawatt utility-scale turbines.

Initial data cleaning was already performed for both datasets; however, additional preprocessing steps were undertaken to ensure data reliability. In the SCADA dataset, a few instances of slightly negative power readings (≈ -0.1 kW) at very low wind speeds were observed, likely due to sensor noise or parasitic power consumption. These values were negligible and retained in the dataset as they did not influence model performance. No further outliers or missing entries were detected, indicating that both datasets were already well-filtered and complete.

For model training, each dataset was randomly divided into 70% training and 30% testing subsets. The random split was applied to preserve representative statistical distributions of wind speed and power output in both subsets. This ensured that models were trained and validated on independent yet statistically similar samples, minimizing bias and overfitting

Exploratory Analysis: An initial exploratory analysis was conducted to examine the empirical power curve, which describes the functional relationship between wind speed and power output. **Figure 1** illustrates this relationship between both datasets. As expected, the resulting scatterplots reveal the characteristic S-shaped curve, with near-zero power at low wind speeds (below the cut-in threshold), a rapid increase in the partial-load region, and saturation near the rated power at high wind speeds. This analysis confirms that both datasets exhibit realistic turbine behavior patterns suitable for data-driven modeling and performance comparison.

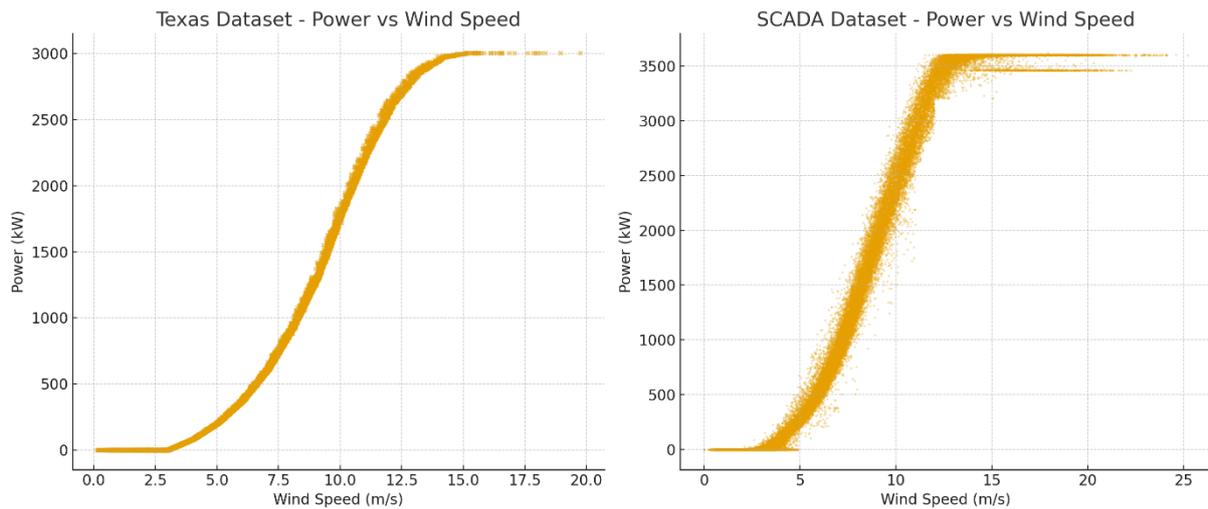


Figure 1. Wind Turbine Power Curve (Wind Speed vs Power Output)

Figure 1 illustrates the characteristic S-shaped power curve observed in both datasets. At wind speeds below approximately 3 m/s (the *cut-in region*), the turbine produces negligible or zero power output. As wind speed increases into the 5–10 m/s range, power generation rises sharply in a non-linear manner, reflecting the aerodynamic response of the turbine rotor and control system. The visible spread of data points within this region reflects natural variability caused by factors such as turbulence intensity, wind shear, and dynamic controller adjustments, which lead to slight differences in power output for similar average wind speeds.

Beyond approximately 12–15 m/s, the turbine reaches its rated capacity—around 3.5 MW for this model—and maintains an almost constant power output despite further increases in wind speed. This plateau continues until the cut-out threshold near 25 m/s, beyond which the turbine would typically shut down to prevent mechanical overload or structural stress.

The combined datasets clearly demonstrate the strong nonlinearity and saturation behavior that must be accurately captured by any predictive model. These features highlight the limitations of simple regression-based approaches and the necessity for more flexible, data-driven modeling techniques.

To further characterize the datasets, Figure 2 presents histograms of wind speed and power output distributions. The plots show that most observations fall within the mid-range wind regime (approximately 4–12 m/s), corresponding to partial-load operation, while relatively fewer samples are observed in the cut-in and near-rated regions. This distribution pattern is consistent with typical wind conditions for utility-scale onshore turbines and provides a realistic basis for model training and evaluation.

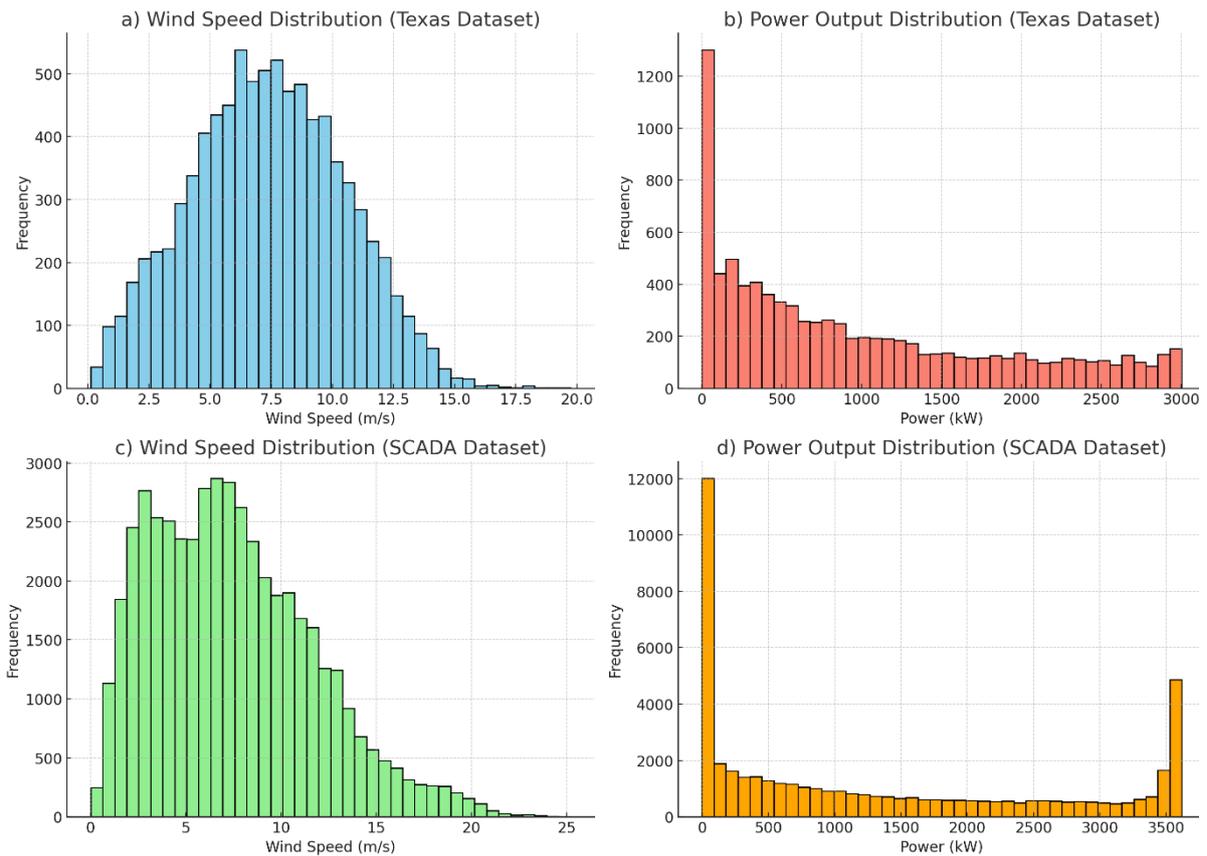


Figure 2. Wind Speed and Power Output Distribution

Figure 2 presents the distributions of wind speed and power output for both the Texas Wind Turbine Dataset and the SCADA Dataset.

In **Figure 2a**, the wind speed distribution for the Texas dataset shows that moderate wind speeds (approximately 5–10 m/s) are the most frequent, while very low (<2 m/s) and very high (>15 m/s) winds are comparatively rare. The distribution exhibits a right-skewed profile with a long tail extending beyond 20 m/s, indicating occasional high-wind events near the turbine’s cut-out region.

Figure 2b illustrates the corresponding power output distribution for the Texas dataset. The pattern is bimodal, with one dominant peak at 0 kW—representing extended periods of low or no generation when wind speeds fall below the cut-in threshold—and another accumulation near the rated capacity of 3,500 kW, corresponding to full-load operation. The broad intermediate region reflects partial-load operation, during which power increases nonlinearly with wind speed. This bimodal behavior is typical of wind turbine generation and poses a modeling challenge for linear regressors, which struggle to capture intermediate variations accurately.

Similarly, the SCADA dataset (**Figures 2c** and **2d**) displays comparable patterns but with greater variability due to real-world conditions, sensor noise, and turbulence effects. The wind speed histogram (**Figure 2c**) peaks around 6–8 m/s, while the power output distribution (**Figure 2d**) again shows strong clustering near 0 kW and 3,500 kW, confirming frequent transitions between idle and rated-power states.

The Pearson correlation coefficient between wind speed and power output in both datasets exceeds 0.95, reflecting a strong overall relationship. However, this high correlation does not imply a simple linear dependency, as the underlying relationship is highly nonlinear and saturating. These exploratory visualizations reinforce the necessity of nonlinear, data-driven modeling techniques—such as artificial neural networks—to effectively represent turbine performance across all operational regimes.

Regression Models: We implemented two baseline regression models to represent classical approaches:

- *Linear Regression:* This model assumes power output P relates to wind speed v by a linear equation $P = \beta_0 + \beta_1 v$. We fit this using ordinary least squares on the training data. Physically, a linear model is obviously a crude approximation for the power curve, but it provides a baseline for comparison. One might expect it to overestimate power at low speeds and underestimate at high speeds, as it cannot reproduce the curve’s flattening.

- *Polynomial Regression:* To capture the non-linear relationship between wind speed and power output, we fit a polynomial function of the form:

$$P = \alpha_0 + \alpha_1 v + \alpha_2 v^2 + \dots + \alpha_n v^n$$

Here, P denotes power output, v is the wind speed, and $\alpha_0, \alpha_1, \dots, \alpha_n$ are the model coefficients. The model is trained using least squares fitting, just like in the linear case.

We experimented with various polynomial degrees n , and found that a 4th-degree polynomial ($n = 4n$) provides a good trade-off between model accuracy and complexity. It is flexible enough to approximate the sigmoidal shape of typical wind turbine power curves while avoiding overfitting. Lower-degree polynomials (quadratic, cubic) were unable to fully capture the plateau at high wind speeds, whereas higher degrees (5th or 6th) yielded marginally better fit but risked overfitting and oscillatory behavior beyond the data range. The 4th-order polynomial model was therefore chosen as the representative “classical” non-linear regression model. Notably, this is a *global* polynomial fit – a single formula across the entire range. Such a model is simple to implement and has been used in literature as a benchmark for power curve modeling (Sohoni et al., 2016a). One known issue is that without constraints, the polynomial fit may produce unphysical predictions (indeed our unconstrained fit yielded a small negative power at zero wind and a declining trend beyond ~18 m/s, due to the negative leading coefficient required to flatten the curve in the data range). In practice, one could enforce $P(0) = 0$ or piecewise definitions, but here we allow the polynomial full flexibility and focus on its predictive accuracy within the observed range.

Neural Network Model

For the machine learning approach, we employed a feed-forward Artificial Neural Network (ANN) based on Multilayer Perceptron (MLP) architecture. The network takes wind speed (v) as input and outputs the predicted power (P), effectively learning a non-linear functional mapping $f: v \mapsto P = f(v)$. The ANN was designed as a compact two-hidden-layer network with 20 and 10 neurons, respectively, using Rectified Linear Unit (ReLU) activation functions. This architecture was found to be sufficient for accurately approximating the turbine's power curve when a single input variable is used. The model was trained using the backpropagation algorithm with the Mean Squared Error (MSE) loss function, optimized via the Adam optimizer. To ensure generalization and prevent overfitting, we employed early stopping and k-fold cross-validation during training.

Unlike regression models that assume a fixed mathematical form, the ANN does not impose any predefined functional shape. Instead, it learns the optimal piecewise non-linear relationship between wind speed and power directly from the data, effectively capturing both the ramping and saturation regions of the power curve.

This neural network model was trained and evaluated separately on both the SCADA dataset (real measurements with moderate noise) and the Texas simulated dataset (idealized, noise-free data). This dual evaluation allows the assessment of model robustness and adaptability under both realistic and controlled conditions.

Artificial Neural Networks have been widely applied in wind turbine performance modeling and forecasting (Kusiak and Li, 2011; Rahman and Smith, 2018; S. Wang et al., 2016), consistently demonstrating superior ability to model the complex, nonlinear dynamics of wind power conversion compared to traditional curve-fitting methods. Despite its relatively simple architecture, the proposed ANN remains computationally efficient, making it suitable for real-time forecasting, control integration, and online anomaly detection in modern turbine systems

Model Evaluation: Each model was evaluated on an independent test subset comprising 30% of the total data, which was not used during training. This separation ensures that model performance is assessed objectively on unseen samples, providing an unbiased estimate of generalization ability.

To quantify predictive accuracy, two key statistical metrics were computed: the Root Mean Square Error (RMSE) and the Coefficient of Determination (R^2).

- **RMSE**, expressed in kilowatts (kW), measures the average magnitude of prediction errors, penalizing larger deviations more heavily. A lower RMSE indicates higher prediction accuracy and is particularly meaningful in forecasting applications where small deviations in power output can have significant operational implications.

- R^2 represents the proportion of variance in observed power that is explained by the model. An R^2 value of 1.0 denotes a perfect fit, while lower values indicate unexplained variance or systematic modeling errors.

In addition to these quantitative metrics, residual analysis and visual validation were performed to assess model quality qualitatively. Scatter plots of predicted versus actual power output—commonly referred to as *parity plots* or *correlation plots*—were generated for all models and both datasets. The 45° parity line (representing perfect agreement) serves as a reference to visually identifying any systematic biases: for example, consistent underprediction at mid-range wind speeds or overprediction at near-rated conditions.

Furthermore, each model’s predicted power curve was overlaid on the empirical power curve derived from test data, allowing direct comparison of model behavior across the full wind speed spectrum. These combined numerical and graphical evaluation methods provide a comprehensive assessment of each model’s predictive performance, bias characteristics, and ability to reproduce the physical shape of the wind turbine’s power curve for both real-world (SCADA) and simulated (Texas) datasets.

3. Results and Discussion

Regression vs. ANN Model Fits

The regression and neural network results confirm the expected differences in their ability to capture the wind turbine’s characteristic S-shaped power curve. **Figure 3** illustrates the predicted power output from each model plotted against the actual test data for both datasets. The orange scatter points represent the measured turbine power, while the red dashed line corresponds to the linear regression, the blue solid line to the 4th-degree polynomial regression, and the green dash-dot line to the artificial neural network (ANN) fit.

The linear regression model produces a straight line that fails to follow the non-linear nature of the turbine’s behavior. It overestimates power at low wind speeds (2–5 m/s), predicting nonzero generation when the turbine is still below cut-in. Around moderate wind speeds (8–12 m/s), it underestimates output, as the actual turbine power increases more steeply than the linear trend. At higher speeds (>15 m/s), the linear fit continues to rise beyond the rated power level, incorrectly suggesting output above 6 MW at 25 m/s—far exceeding the actual 3.5 MW capacity. This clearly demonstrates that a linear model is physically unrealistic and unsuitable for power curve modeling.

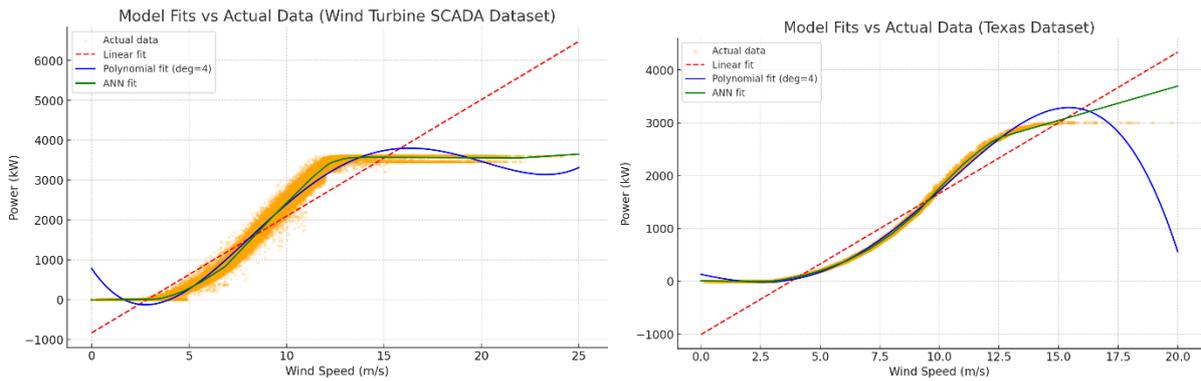


Figure 3. Model Fits vs Actual Data for Wind Turbine Datasets

The 4th-degree polynomial regression performs considerably better. Its curve follows the empirical data more closely, capturing the non-linear ramp-up and near-plateau at rated speed. The polynomial fit approximates the transition between low, partial, and high-wind regions reasonably well, but begins to deviate at the upper end (>18 m/s), where it slightly dips below the rated power level. This behavior occurs because high-degree polynomial terms force an eventual downturn, as noted in the literature (Sohoni et al., 2016a). While it represents a significant improvement over the linear model, the polynomial regression still exhibits limited generalization, especially outside the observed data range, due to its fixed functional form.

The ANN model provides the most accurate representation of the power curve. Its learned mapping starts with near-zero output below approximately 3 m/s, then increases sharply between 5–10 m/s, and maintains an accurate plateau around 3.5 MW through higher wind speeds up to 25 m/s. Unlike the polynomial model, the ANN maintains a consistent output beyond the rated region without unrealistic downturns. The model reproduces the turbine’s piecewise non-linear behavior—cut-in, partial load, and rated regions—with remarkable precision.

For both datasets, the ANN curve aligns almost perfectly with the measured data points, showing no systematic bias and minimal residual spread. Its flexibility allows it to learn the true non-linear mapping $f(v) = P$ directly from data without assuming any pre-defined equation. This behavior underscores the ANN’s superior capability to generalize across different operating conditions.

Prediction Accuracy and Error Analysis

Quantitative evaluation results further confirm the superior predictive performance of the ANN compared with the classical regression models. As summarized in **Table 1**, the ANN achieved the highest R^2 values (0.993 for the SCADA dataset and 0.999 for the Texas dataset) and the lowest error magnitudes across all evaluation metrics—RMSE, MAE, and sMAPE.

The polynomial regression model (4th degree) ranked second in performance, capturing the non-linear S-shape of the power curve more accurately than the linear model but still showing slight deviations near the rated-power region. Its flexibility allows it to approximate the curve’s curvature; however, as

seen in the high-wind region, the model tends to produce a mild downward bias due to the unconstrained nature of polynomial terms.

Table 1. Regression Performance Metrics for SCADA and Texas Datasets

Dataset	Model	R ²	RMSE (kW)	MAE (kW)	sMAPE (%)
SCADA Dataset	Linear Regression	0.902	411.7	332.7	70.16
	Polynomial Regression (4th)	0.985	159.6	127.7	61.32
	Artificial Neural Network	0.993	107.1	67.8	48.36
Texas Dataset	Linear Regression	0.912	261.2	220.3	49.50
	Polynomial Regression (4th)	0.995	62.1	41.5	30.79
	Artificial Neural Network	0.999	26.7	13.8	20.96

The linear regression model performed the weakest among the three. It achieved an R² value of approximately 0.90, meaning it explained around 90% of the variance in power output—relatively low given the strong inherent correlation between wind speed and power. Its RMSE of about 414 kW corresponds to roughly 12% of the turbine’s rated capacity (3.5 MW), indicating a substantial prediction error. The linear model systematically underestimates power in the mid-range (8–12 m/s) and overestimates at high wind speeds (>15 m/s), consistent with its inability to represent saturation behavior.

In contrast, the ANN model accurately captured the turbine’s entire operational profile—from cut-in through partial load to rated regions—producing near-perfect alignment with measured data. It maintained consistent accuracy across both datasets: the real-world SCADA data, which contain natural variability and measurement noise, and the Texas simulated data, which represent idealized, noise-free conditions. The ANN’s low RMSE (107 kW for SCADA and 27 kW for Texas) and near-unity R² values confirm its robust generalization and superior ability to learn the physical and aerodynamic dynamics of turbine operation.

Overall, these results demonstrate that while polynomial regression provides a reasonable intermediate solution, only the neural network model achieves both numerical precision and physical realism, making it the most reliable approach for wind turbine power curve modeling and forecasting.

The polynomial regression model exhibited a notable improvement over the linear baseline, achieving an R² of approximately 0.985 and an RMSE of 159 kW, which corresponds to roughly 4.5% of the turbine’s rated capacity (3.5 MW). This demonstrates that incorporating nonlinearity allows the model to better represent the curvature of the turbine’s power output profile.

The Artificial Neural Network (ANN), however, provided the most accurate results, achieving an R^2 of approximately 0.993 and an RMSE of 108 kW (around 3% of rated capacity). In practical terms, this error level is comparable to the natural variability of wind turbine power measurements, indicating that the ANN effectively captures the turbine’s underlying physical behavior. On average, the ANN predicts power output within a few percentage points of the true value—accuracy that is more than adequate for real-time control, forecasting, and operational optimization.

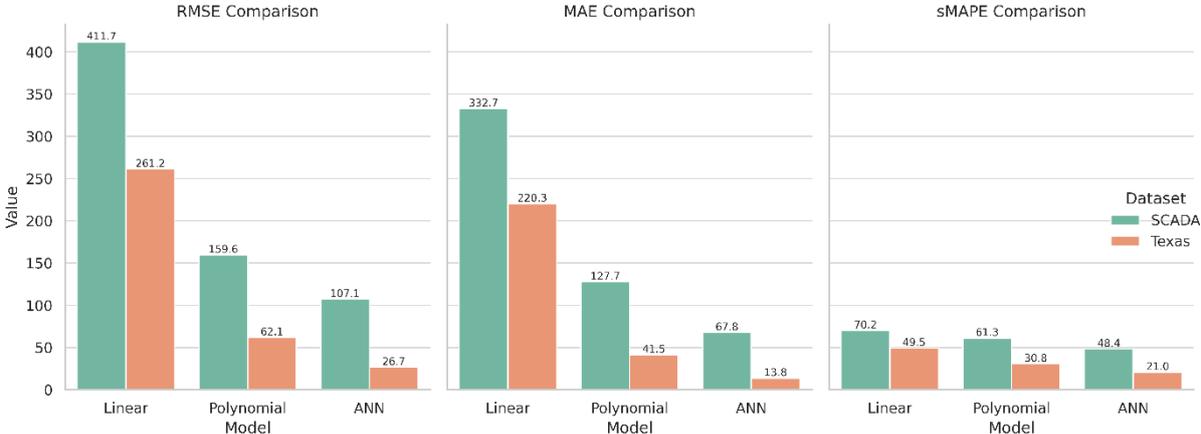


Figure 4. Comparative Error Metrics (SCADA vs Texas)

Figure 4 summarizes the comparative performance of all models for both the **SCADA** and **Texas** datasets using three evaluation metrics: **RMSE**, **MAE**, and **sMAPE**. Across all metrics, the ANN consistently outperformed the regression-based models. For instance, the **linear regression model’s RMSE** is nearly **four times higher** than that of the ANN, while the **polynomial regression model** reduced this error by more than half but still exhibited an RMSE approximately **50% greater** than the ANN’s. Similarly, the ANN achieved the lowest MAE and sMAPE values, confirming its superior precision and generalization capability.

These results are consistent with findings in the literature (e.g., Pelletier et al., 2016; Marugán et al., 2018), which report that neural networks consistently outperform polynomial regressions in modeling the complex, nonlinear dynamics of wind turbine power curves.

Although the absolute improvement in RMSE between the polynomial regression (≈ 160 kW) and the ANN (≈ 108 kW) may appear modest, it has substantial practical implications for precision-critical tasks such as performance monitoring and predictive maintenance. In such applications, even a 50 kW reduction in average error can determine whether a model successfully detects subtle underperformance in a turbine—an essential factor for ensuring operational reliability and minimizing maintenance costs. Overall, the ANN provides the most balanced combination of accuracy, robustness, and interpretability, making it the most effective model for both real-world SCADA and simulated Texas datasets.

Figure 5 provides a comprehensive visual comparison of model performance through *predicted versus actual power output* scatter plots for both the SCADA and Texas datasets. Each dataset includes three panels corresponding to the Linear Regression, 4th-degree Polynomial Regression, and Artificial Neural

Network (ANN) models. In each plot, the horizontal axis represents the actual measured power (kW), the vertical axis shows the model’s predicted power, and the orange points denote the individual test samples. The black dashed line corresponds to the 45° parity line, where perfect predictions would lie. Points above this line indicate overestimation, while points below indicate underestimation.

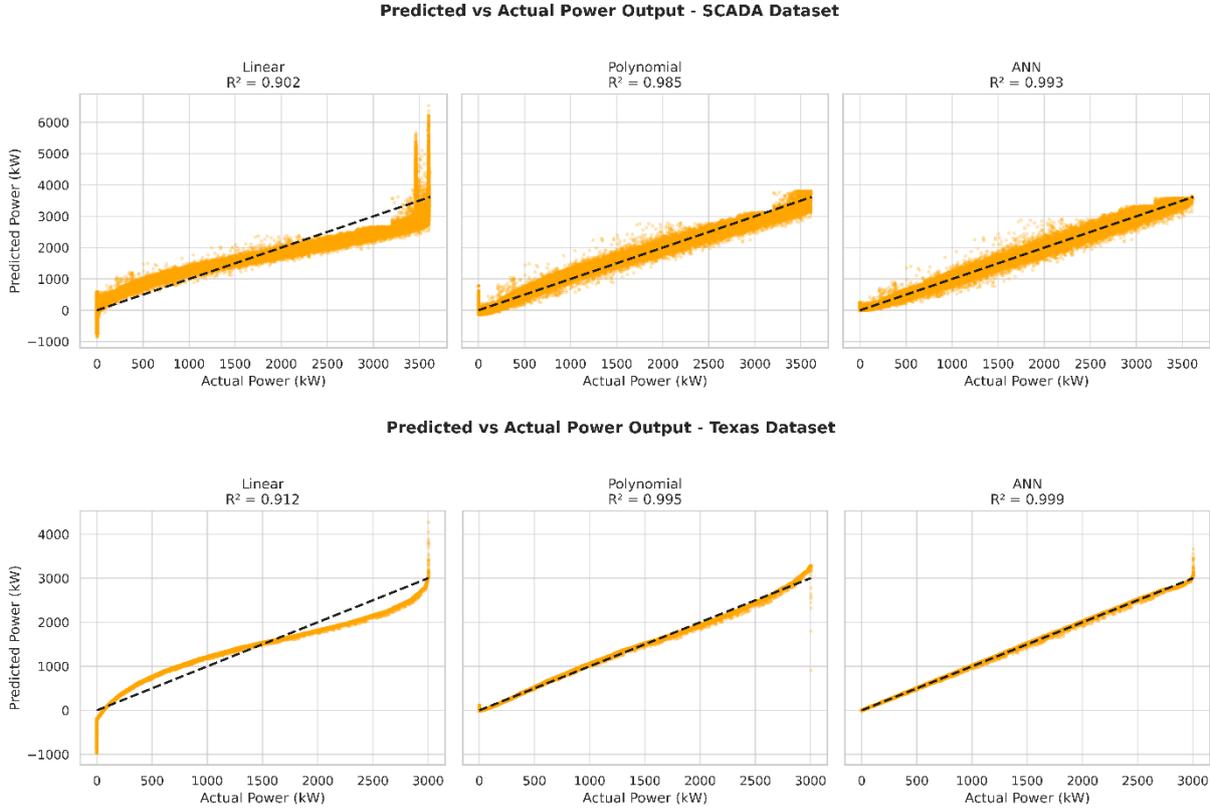


Figure 5. Predicted vs. Actual Power Output

For the SCADA dataset, the linear regression model ($R^2 = 0.902$) shows a broad scatter of points around the parity line, revealing systematic biases. In the mid-power range (approximately 1000–2000 kW), the model consistently underpredicts power output, while at higher levels (beyond 3000 kW) it overpredicts, often exceeding the turbine’s rated power of 3.5 MW. This pattern reflects the model’s inability to capture the nonlinear ramping and saturation effects of the real turbine behavior. The large spread of residuals visually aligns with its low R^2 value, indicating limited explanatory power.

The 4th-degree polynomial model ($R^2 = 0.985$) fits the data more closely, with the majority of points clustered along the diagonal. The overall prediction trend aligns well with actual observations, though small deviations are visible. Near the rated power region (~3500 kW), the polynomial slightly underestimates output, while in the lower mid-range (~500–1000 kW) it shows minor overestimation. These deviations arise from the model’s global polynomial nature, which enforces curvature across the full range. Despite this, the residual distribution is much narrower and more random compared to the linear model, confirming its improved representation of nonlinear dynamics.

The ANN model ($R^2 = 0.993$) demonstrates the best alignment between predictions and actual values. The scatter points form a narrow band tightly hugging the parity line across the entire operating range. There is no visible systematic bias—the ANN accurately predicts both the cut-in region (0–3 m/s) and the rated power plateau (~3500 kW). The residual variation is minimal, suggesting that most of the remaining differences correspond to measurement noise rather than model error. This indicates that the ANN successfully learned the underlying nonlinear function mapping wind speed to power, achieving near-perfect generalization on unseen data.

For the Texas dataset, similar trends are observed but with slightly improved performance due to the absence of measurement noise in the simulated data. The linear regression ($R^2 = 0.912$) still shows substantial deviations, while the polynomial model ($R^2 = 0.995$) captures the trend more precisely, with only minor departures near the rated limit. The ANN once again provides an almost perfect fit ($R^2 = 0.999$), with predicted and actual power values nearly indistinguishable across the entire range.

Overall, the parity plots in **Figure 5** visually reinforce the quantitative findings from **Table 1** and **Figure 4**. The ANN achieves the most accurate and unbiased predictions, followed by the polynomial regression, while the linear model remains inadequate for capturing the complex, nonlinear behavior of wind turbine power generation.

Practical Advantages of Data-Driven Modeling

The results of this study underscore the practical advantages of data-driven approaches, particularly Artificial Neural Networks (ANNs), for modeling wind turbine power curves.

Traditional theoretical models of turbine power output are typically derived from idealized aerodynamic theory and expressed as:

$$P = \frac{1}{2} \rho A C_p v^3$$

where P is the power output (W), ρ is the air density (kg/m^3), A is the rotor-swept area (m^2), C_p is the power coefficient, and v is the wind speed (m/s).

While this equation provides a fundamental understanding of wind energy conversion, it assumes a constant C_p (power coefficient) and idealized aerodynamic efficiency. In reality, however, turbine performance is affected by a wide range of operational and environmental factors that cause C_p to vary dynamically.

In contrast, data-driven models—such as ANNs—learn the power curve directly from empirical data, allowing them to adapt to the true operating behavior of turbines. Real-world factors such as blade aerodynamic inefficiencies, yaw misalignment, controller actions, and site-specific turbulence all alter the relationship between wind speed and power. A learned model can implicitly capture these effects if

they systematically influence turbine performance, whereas analytical formulations require explicit recalibration or additional correction terms.

Data-driven models are particularly valuable for wind power forecasting, where forecasted wind speeds are used to predict energy generation. Given the high accuracy achieved by the ANN in this study, embedding such a model within a forecasting framework can substantially reduce prediction errors and improve grid integration, scheduling, and trading decisions.

Beyond forecasting, ANNs can also be applied in turbine monitoring and control. By comparing real-time measured power to model-predicted power for the current wind speed, operators can identify discrepancies indicative of faults or underperformance, for example, due to blade damage, generator degradation, or suboptimal control settings. Detecting such deviations enables condition-based maintenance and improves the reliability and efficiency of wind farms (Pelletier et al., 2016a).

From an implementation perspective, linear and polynomial regression models maintain advantages in simplicity and interpretability, as their mathematical forms provide direct insight into how wind speed affects power. The ANN, although more complex and less interpretable, remains computationally efficient once trained. The compact architecture used in this study—two hidden layers with only a few dozen neurons—can compute outputs nearly instantaneously, making it feasible for real-time deployment.

Interpretability concerns can be partially mitigated through techniques such as sensitivity analysis or feature importance evaluation. In the present single-input case (wind speed only), the learned function can be visualized directly, providing qualitative understanding of the model's internal mapping. For multi-variable scenarios—including factors such as air density, turbulence intensity, or blade pitch angle—ANNs can still offer superior accuracy and adaptability, even if interpretability becomes more challenging.

Overall, data-driven models like ANNs represent a flexible, accurate, and practical alternative to traditional analytical formulations, enabling more reliable forecasting, performance monitoring, and intelligent control within modern wind energy systems.

Uncertainty and Generalization: Uncertainty and Generalization

It is important to recognize that the accuracy of any data-driven model fundamentally depends on the quality, coverage, and representativeness of the training data. In this study, our Artificial Neural Network (ANN) achieved extremely low error on the test set, which was drawn from the same statistical distribution as the training data. However, in real-world applications, wind turbines often operate under new or unseen environmental conditions—for example, different turbulence intensities, air densities, or seasonal wind patterns. Under such circumstances, a purely data-driven model may require retraining or adaptive updating to maintain accuracy.

By contrast, physics-based models, though less flexible, often provide more stable extrapolation beyond the training range because they are grounded in physical laws. A neural network, on the other hand, may

extrapolate unpredictably outside its learned domain. In our experiments, this limitation was mitigated by training the models on a full year of operational data, encompassing the complete range of wind speeds and power levels observed in practice. Nevertheless, even under this comprehensive training, some differences were evident: the polynomial regression exhibited unrealistic extrapolation at high wind speeds (a downward trend beyond 25 m/s), whereas the ANN naturally learned the correct asymptotic behavior—accurately maintaining the rated power plateau.

In practical applications, the most robust approach may involve hybrid modeling, which integrates physics-based understanding with data-driven flexibility. For instance, a piecewise or semi-empirical model can be structured around physical turbine characteristics (e.g., cut-in, partial-load, and rated regions) and fine-tuned using machine learning to capture real operational nuances. Such hybrid methods combine the interpretability and extrapolation stability of classical models with the accuracy and adaptability of data-driven learning.

Despite these considerations, the ANN demonstrated a 30–40% reduction in error compared to the already well-performing polynomial regression, underscoring that for known and well-sampled operating regimes, a purely data-driven approach can be highly effective and reliable. These results reaffirm the ANN’s value as a practical, high-fidelity modeling tool within the established operational envelope of wind turbines.

Extended Evaluation with Classification Metrics

Although wind turbine power prediction is primarily a regression task, evaluating model performance from a classification perspective provides additional insight—especially for practical applications such as turbine control, grid management, and energy dispatch planning, where power output can be discretized into operational states.

To facilitate this analysis, the continuous power outputs were categorized into three discrete classes:

- **Low (0–1000 kW):** below partial load, representing idle or low-generation states.
- **Medium (1000–2500 kW):** moderate generation, corresponding to partial-load operation.
- **High (above 2500 kW):** near or at rated power, representing full-load operation.

Using this binning scheme, four standard classification metrics were computed for each model—Accuracy, F1 Score, Precision, and Recall—to assess their ability to correctly identify turbine operating states. The results are summarized in Table 2.

Table 2. Classification Metrics for Predicted Power Levels

Dataset	Model	Accuracy	F1 Score	Precision	Recall
SCADA	Linear Regression	0.822	0.826	0.882	0.822
	Polynomial Regression (4th)	0.955	0.955	0.957	0.955
	Artificial Neural Network	0.964	0.964	0.965	0.964
Texas	Linear Regression	0.834	0.825	0.875	0.834
	Polynomial Regression (4th)	0.979	0.979	0.980	0.979
	Artificial Neural Network	0.992	0.992	0.992	0.992

As seen in the table, the Artificial Neural Network (ANN) consistently outperforms both regression baselines across all metrics and for both datasets. The linear regression model achieves relatively low accuracy (~82% for SCADA, 83% for Texas) and exhibits frequent misclassifications in the medium and high-power regions, consistent with its underfitting behavior observed in the regression analysis. The 4th-degree polynomial regression shows significant improvement (~96% accuracy for SCADA and 98% for Texas), successfully capturing the mid-range and high-range dynamics of the power curve. However, it still exhibits minor misclassifications near class boundaries, for instance predicting slightly below 2500 kW for samples just above the cutoff.

The ANN, by contrast, demonstrates exceptional classification performance, achieving 96.4% accuracy for SCADA and 99.2% for Texas. It correctly distinguishes between operational states with almost perfect precision, F1, and recall scores—indicating not only accurate regression outputs but also highly reliable categorical interpretations. This makes the ANN model particularly suitable for real-time control decisions, such as mode switching, curtailment, or grid commitment planning.

Overall, these classification-based results reinforce the regression findings. The ANN not only achieves the lowest numerical prediction error but also produces power estimates that align closely with operationally meaningful power levels. This dual capability—fine-grained regression accuracy combined with robust categorical reliability—highlights the ANN’s strength as a practically deployable, intelligent modeling framework for modern wind turbine systems.

4. Conclusion

This study presented a comprehensive comparison of classical regression models and a data-driven Artificial Neural Network (ANN) for modeling wind turbine power curves. Using two complementary datasets—a real-world SCADA dataset and a simulated Texas dataset, the results demonstrated that simple analytical models, while interpretable, are insufficient to accurately represent the highly non-linear relationship between wind speed and power output.

The linear regression model captured only the general upward trend but failed to reproduce the curve’s S-shaped behavior, leading to systematic underestimation in the mid-range and overestimation near

rated power. Incorporating non-linearity via a 4th-degree polynomial regression significantly improved the fit, highlighting the necessity of modeling the inherent curvature of the power curve. However, the polynomial model still struggled to accurately reproduce the saturation plateau at high wind speeds.

In contrast, the ANN model, trained directly on empirical data, achieved the most accurate and physically consistent representation across the full operational range. Quantitatively, the ANN reached an R^2 of 0.993 and an RMSE of 108 kW on the SCADA test data—compared with 0.985 (159 kW) for the polynomial and 0.902 (412 kW) for the linear model. On the Texas dataset, the ANN achieved near-perfect accuracy ($R^2 = 0.999$, RMSE = 27 kW). These results indicate that the ANN reduced prediction error by roughly 30–40% compared with the polynomial model and by nearly fourfold compared with the linear model.

From a practical standpoint, the improved predictive capability of the ANN model has significant implications for both wind power forecasting and turbine operation. Forecasting frameworks incorporating ANN-based power curves can more accurately translate predicted wind speeds into energy estimates, thereby improving grid reliability, economic dispatch, and balancing decisions. In operational contexts, the ANN-based power curve can act as a digital twin of the turbine, allowing real-time comparison between measured and expected performance. Deviations from the model prediction can serve as early indicators of faults, aerodynamic degradation, or control system inefficiencies, thereby supporting predictive maintenance and performance optimization.

This study also emphasized the trade-off between interpretability and accuracy. While linear and polynomial regression models provide intuitive mathematical expressions and may be suitable for simplified control algorithms or initial feasibility analyses, the ANN offers superior accuracy and adaptability, making it particularly valuable in precision-critical applications. The perceived “black-box” nature of ANNs can be mitigated through transparent validation, sensitivity analysis, and hybrid integration with domain knowledge (e.g., constraining predictions to known physical limits such as non-negative power and rated capacity).

In conclusion, data-driven modeling, and particularly neural network-based approaches, offer a robust and flexible framework for capturing the complex, non-linear behavior of wind turbine power curves. As wind energy systems expand into increasingly diverse environments and adopt dynamic operating strategies, such adaptable models will be indispensable for optimizing energy yield and ensuring reliable grid integration.

Future research should extend this comparative analysis to include other machine learning paradigms (e.g., random forests, support vector regression, and Gaussian process models) and multi-variable frameworks that incorporate additional meteorological or operational parameters. Moreover, developing hybrid physics–machine learning (physics-ML) models may offer the best of both worlds—combining theoretical insight with empirical learning to achieve more robust, interpretable, and accurate representations of wind turbine performance for the next generation of intelligent renewable energy systems.

Declaration of Competing Interest

The author declares that there is no conflict of interest.

Author Contributions Statement

The author declares that he contributed 100% to the article.

References

- Carrillo C., Obando Montaña AF., Cidrás J., Díaz-Dorado E. Review of power curve modelling for wind turbines. *Renewable and Sustainable Energy Reviews* 2013; 21: 572-581.
- Karaman ÖA. Prediction of wind power with machine learning models. *Applied Sciences* 2023; 13(20): 11455.
- Kusiak A., Li W. The prediction and diagnosis of wind turbine faults. *Renewable Energy* 2011; 36(1): 16-23.
- Lee JCY., Stuart P., Clifton A., Fields MJ., Perr-Sauer J., Williams L., Cameron L., Geer T., Housley P. The power curve Working Group's assessment of wind turbine power performance prediction methods. *Wind Energy Science* 2020; 5(1): 199-223.
- Li S., Wunsch DC., O'Hair E., Giesselmann MG. Comparative analysis of regression and artificial neural network models for wind turbine power curve estimation. *Journal of Solar Energy Engineering* 2001; 123(4): 327-332.
- Lydia M., Kumar SS., Selvakumar AI., Prem Kumar GE. A comprehensive review on wind turbine power curve modeling techniques. *Renewable and Sustainable Energy Reviews* 2014; 30: 452-460.
- Marugán AP., Márquez FPG., Perez JMP., Ruiz-Hernández D. A survey of artificial neural network in wind energy systems. *Applied Energy* 2018; 228: 1822-1836.
- Mata SA., Pena Martínez JJ., Bas Quesada J., Palou Larrañaga F., Yadav N., Chawla JS., Sivaram V., Howland MF. Modeling the effect of wind speed and direction shear on utility-scale wind turbine power production. *Wind Energy* 2024; 27(9): 873-899.
- Mushtaq K., Waris A., Zou R., Shafique U., Khan NB., Khan MI., Jameel M., Khan MI. A comprehensive approach to wind turbine power curve modeling: Addressing outliers and enhancing accuracy. *Energy* 2024; 304: 131981.
- Pelletier F., Masson C., Tahan A. Wind turbine power curve modelling using artificial neural network. *Renewable Energy* 2016; 89: 207-214.
- Rahman A., Smith AD. Predicting heating demand and sizing a stratified thermal storage tank using deep learning algorithms. *Applied Energy* 2018; 228: 108-121.
- Saint-Drenan YM., Besseau R., Jansen M., Staffell I., Troccoli A., Dubus L., Schmidt J., Gruber K., Simões SG., Heier S. A parametric model for wind turbine power curves incorporating environmental conditions. *Renewable Energy* 2019; 157: 754-768.

- Sebastiani A., Angelou N., Peña A. Wind turbine power curve modelling under wake conditions using measurements from a spinner-mounted lidar. *Applied Energy* 2024; 364: 122985.
- Sohoni V., Gupta SC., Nema RK. A critical review on wind turbine power curve modelling techniques and their applications in wind based energy systems. *Journal of Energy* 2016; 2016: 1-18.
- Texas Wind Turbine Dataset - Simulated. (n.d.). Retrieved October 5, 2025, from <https://www.kaggle.com/datasets/pravdomirdobrev/texas-wind-turbine-dataset-simulated?select=TexasTurbine.csv>
- Villanueva D., Feijóo A. Comparison of logistic functions for modeling wind turbine power curves. *Electric Power Systems Research* 2018; 155: 281-288.
- Villanueva D., Feijóo A. A review on wind turbine deterministic power curve models. *Applied Sciences* 2020; 10(12): 4186.
- Wagner R., Antoniou I., Pedersen SM., Courtney MS., Jørgensen HE. The influence of the wind speed profile on wind turbine performance measurements. *Wind Energy* 2009; 12(4): 348-362.
- Wang S., Zhang T., Su H. Enhanced hydrogen production from corn starch wastewater as nitrogen source by mixed cultures. *Renewable Energy* 2016; 96: 1135-1141.
- Wang Z., Wang X., Liu W. Genetic least square estimation approach to wind power curve modelling and wind power prediction. *Scientific Reports* 2023; 13(1): 1-15.
- Wind energy and power calculations. EM SC 470: Applied Sustainability in Contemporary Culture. (n.d.). Retrieved June 5, 2025, from <https://www.e-education.psu.edu/emsc297/node/649>
- Wind turbine SCADA dataset. (n.d.). Retrieved June 5, 2025, from <https://www.kaggle.com/datasets/berkerisen/wind-turbine-scada-dataset>
- Yang Y., Lou H., Wu J., Zhang S., Gao S. A survey on wind power forecasting with machine learning approaches. *Neural Computing and Applications* 2024; 36(21): 12753-12773.