Journal of Physics: Conference Series **3224** (2026) 022052 doi:10.1088/1742-6596/3224/2/022052

Uncertainty-Aware Wind Turbine Power Curve Modelling with Rotor Turbulence Intensity Asymmetry Using Two-Stage Gaussian Process Regression

Shubham Baisthakur[1], Nur M.M. Kalimullah[1] and Breiffni Fitzgerald[1]

[1]Department of Civil, Structural and Environmental Engineering, Trinity College Dublin, Dublin, Ireland

E-mail: Shubham.Baisthakur@tcd.ie

Abstract. Accurate power curve modelling under variable atmospheric conditions is essential for wind energy project design and resource assessment. Current practice assumes isotropic turbulence intensity across the rotor, but rotor-scale turbulence intensity asymmetry can significantly amplify power output variability, yet remains absent from standard IEC specifications and power curve modelling literature. This study systematically investigates how rotor turbulence intensity asymmetry contributes to power variance through aeroelastic simulations of a 15 MW reference turbine. A full-factorial experimental design explores the combined influence of wind speed, wind shear, hub-height turbulence intensity, and rotor turbulence intensity asymmetry, generating a dataset of 5,760 simulations across 36 atmospheric conditions and 20 wind speeds. Statistical analysis reveals that while wind speed and hub-height turbulence intensity dominate power variance globally, rotor turbulence intensity asymmetry exhibits pronounced regime-dependent effects. Variance amplification ranges from approximately 37 to 40% in below-rated and near-rated operation, where power output is tightly coupled to inflow conditions, to 24% in above-rated operation where pitch control attenuates but does not eliminate the effect. A two-stage heteroscedastic Gaussian Process Regression framework is developed to capture this regime-dependent structure by separately modelling mean power and input-dependent variance as functions of all four atmospheric parameters. The proposed framework improves prediction accuracy by 27% and reduces prediction uncertainty by 25% relative to univariate wind-speed-only models, demonstrating the value of atmospheric conditioning for probabilistic power curve estimation. The approach provides a practical tool for design-phase uncertainty quantification and wind resource assessment in operating conditions where rotor-scale turbulence variation is pronounced but remains neglected in standard modelling practice.

1. Introduction

Over the past two decades, the levelised cost of electricity from wind power has fallen steadily, driven by larger turbines, advances in control strategies, and optimised aerodynamic designs [1, 2]. In this increasingly competitive sector, accurate estimates of power production are crucial for project financing, grid stability, and energy security. The power curve is a fundamental tool for this purpose, as it establishes the relationship between wind speed and expected power output across the operating wind speed range of the turbine.

Journal of Physics: Conference Series **3224** (2026) 022052 doi:10.1088/1742-6596/3224/2/022052

Current practice for generating power curves, as standardised by IEC 61400-12-1, relies on the method of bins, where power output is expressed as a function of the ten-minute average hub-height wind speed [3]. A single hub-height wind speed, however, fails to capture the influence of higher-order atmospheric variability, particularly wind shear and turbulence intensity (TI), on power production. Studies have demonstrated that these factors can significantly alter the turbine power output [4–6], and their importance grows with rotor size as blade sections at different heights are exposed to increasingly distinct inflow conditions. The IEC standard prescribes correction factors for wind shear and turbulence intensity, and alternate approaches such as rotor-equivalent wind speed methods [7] and data-driven SCADA-based refinement [8] have been proposed to further improve mean power predictions. All of these approaches, however, treat shear and turbulence intensity as rotor-uniform quantities, implicitly assuming that a single representative value adequately characterises the inflow across the entire swept area. The effect of rotor-scale turbulence non-uniformity on power output uncertainty therefore remains unaddressed.

Vertical variation in turbulence intensity across the rotor is a well-established property of the atmospheric boundary layer. Under stable atmospheric stratification, mechanical turbulence generated near the surface is suppressed by buoyancy, creating a strong vertical gradient in turbulence intensity that increases from hub height downward. This gradient intensifies with atmospheric stability and becomes increasingly significant as rotor diameters grow, since larger swept areas span a greater portion of the boundary layer where this gradient is active. While IEC 61400-12-1 acknowledges that turbulence intensity may vary across the rotor, no method for incorporating this spatial variation into power curve modelling is prescribed, and its effect on power output uncertainty has remained largely unevaluated in the literature. To address this, this study introduces the *rotor turbulence intensity asymmetry* factor (gTI), defined as the ratio of turbulence intensity at the lowest to the highest point of the rotor disc, to capture this asymmetry in a physically interpretable form. This dimensionless factor remains independent of the absolute TI level, and directly quantifies the degree of departure from the rotor-uniform assumption embedded in current IEC practice and existing power curve models. The rotor-scale turbulence intensity asymmetry produces asymmetric aerodynamic loading across the blades, creating fluctuations in torque and instantaneous power output.

This study presents a systematic quantification of the impact of rotor turbulence intensity asymmetry on power output uncertainty. A full-factorial design of experiments is chosen to isolate the effect of rotor turbulence intensity asymmetry from other atmospheric parameters, enabling systematic assessment via a two-stage heteroscedastic Gaussian Process Regression (GPR) framework. This approach separately models mean power output and input-dependent variance, allowing changes in power variability to be examined independently of mean performance. The two-stage decomposition allows the variance to vary with atmospheric inputs while avoiding the optimisation complexity of joint heteroscedastic GP models. By separating the effects on the mean from those on variability, the framework identifies which atmospheric parameters have the greatest impact on uncertainty in power estimates, a distinction particularly valuable during the design phase when site characterisation is incomplete.

The remainder of this paper is organised as follows. Section 2 describes the full-factorial experimental design and data generation process. Section 3 presents a statistical characterisation of power variability and identifies the dominant atmospheric parameters influencing power variance, with emphasis on regime-dependent effects of rotor turbulence intensity asymmetry. Section 4 presents the two-stage heteroscedastic GPR methodology. Section 5 evaluates the modelling approach and quantifies the contribution of rotor turbulence intensity asymmetry to power uncertainty. Section 6 concludes with key findings and directions for future work.

Journal of Physics: Conference Series **3224** (2026) 022052 doi:10.1088/1742-6596/3224/2/022052

2. Generation of Training Data

A full-factorial design of experiments is used to generate training data for the heteroscedastic GPR models. The dataset is constructed to isolate the effect of rotor turbulence intensity asymmetry while maintaining systematic coverage of atmospheric conditions relevant to modern large-scale wind turbines. This section summarises the experimental design and simulation configuration.

2.1. Experimental Design

The training dataset is constructed by varying four key atmospheric parameters that influence turbine power production: wind speed (v), shear exponent (α), hub-height turbulence intensity (TI), and rotor turbulence intensity asymmetry (gTI). A full-factorial design ensures systematic coverage of the parameter space and allows the effect of rotor turbulence intensity asymmetry to be examined independently of other atmospheric influences.

The shear exponent α is sampled at three levels: $0.10, 0.20, 0.30$, spanning the range of conditions relevant to normal power production, consistent with the normal wind profile model prescribed in IEC 61400-1 [9]. Hub-height turbulence intensity (TI) is sampled at three levels: $0.08, 0.12, 0.16$, reflecting typical design conditions [10]. Rotor turbulence intensity asymmetry (gTI) is sampled at four levels: $1.0, 1.2, 1.5, 1.8$, where 1.0 represents uniform rotor turbulence and larger values indicate increasing vertical asymmetry.

Wind speed spans the turbine operating range from 3 to 25 m/s, with finer resolution in the below-rated and transition regions to capture critical operating points. Specifically, 12 points are assigned within 3-13 m/s (below-rated and transition at 10.5 m/s) and 8 points within 13-25 m/s (above-rated). These wind speed points are sampled using the Latin hypercube sampling technique [11] to ensure space-filling coverage. Combined with the 36 environmental conditions ($3 \times 3 \times 4$ combinations of α, TI, and gTI), this yields 720 unique environmental parameter-wind speed combinations.

For each combination, eight independent turbulence realisations are generated by varying the TurbSim random seed, capturing the inherent stochastic variability in power at fixed atmospheric conditions. This ensemble approach distinguishes between irreducible turbulence-driven fluctuations and structured variability arising from atmospheric parameter changes. From each 900-second simulation, mean power and variance are computed and aggregated across the eight realisations to yield stable estimates. The complete dataset comprises 36 environmental conditions $\times$ 20 wind speeds $\times$ 8 realisations = 5,760 OpenFAST simulations.

2.2. Simulation Configuration

Aeroelastic simulations employ OpenFAST with the IEA 15 MW reference turbine [12] and ROSCO reference controller [13]. Inflow fields are generated using TurbSim with the von Kármán spectral model and user-defined vertical profiles of wind speed and turbulence intensity, allowing systematic control of vertical turbulence variation while preserving realistic spatial coherence. Each field spans 900 s at 0.05 s temporal resolution on a 285×285 m domain with 25×25 spatial grid (sufficient for the 240 m rotor diameter). OpenFAST integration uses a 0.005 s timestep for numerical stability, and post-processing extracts mean power and variance for each condition, forming the training dataset for the probabilistic framework.

3. Analysis of Power Output Variability and Statistical Characterisation

The full-factorial experimental design described in Section 2 produces a structured dataset in which atmospheric parameters (v, α, TI, and gTI) are systematically varied. This section

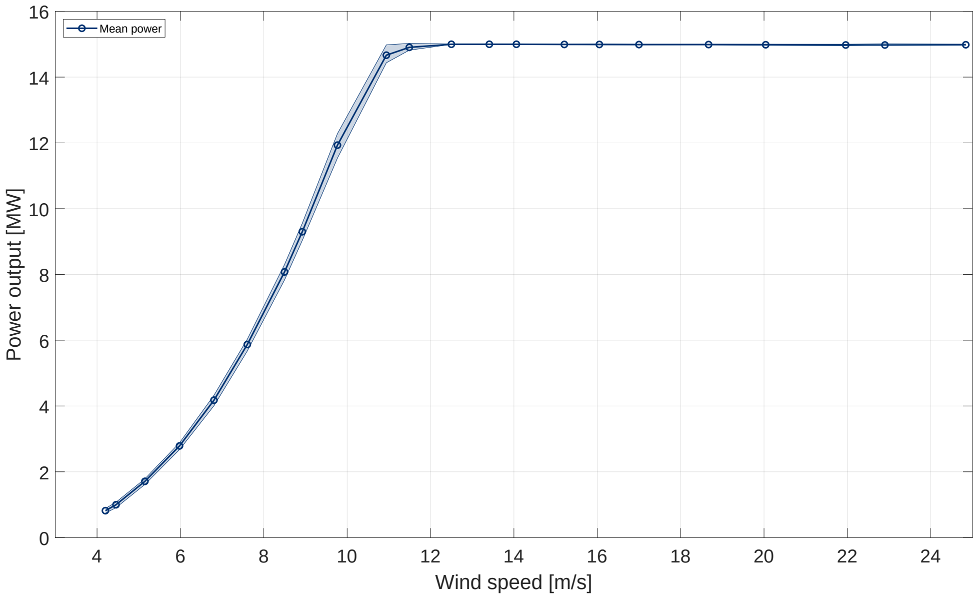

Figure 1. Deterministic mean power curve with interquartile range capturing unresolved atmospheric variability

presents a statistical characterisation of the simulated power response, focusing on how variability manifests across operating conditions and how it deviates from the deterministic representation commonly adopted in power curves.

3.1. Mean power curve and unresolved variability

Figure 1 presents the mean power curve (averaging across all stochastic inflow realisations and atmospheric conditions in each wind speed bin) alongside the interquartile range (IQR; 25th-75th percentiles), representing the deterministic and probabilistic representations, respectively.

Despite the smooth shape of the mean power curve, the shaded region reveals a persistent spread across the below-rated and transition regions. Around rated conditions, the IQR reaches approximately 8% of the mean power, indicating that a single deterministic value does not fully represent the range of power outcomes under realistic atmospheric variability. Importantly, this spread arises at fixed nominal wind speed, showing that power fluctuations are intrinsically linked to higher-order inflow characteristics rather than uncertainty in mean wind speed alone. At multiple operating points, power variance increases significantly even when mean response remains relatively unchanged. This phenomenon, where output variability changes independently of mean power, defines heteroscedastic behaviour and cannot be captured by conventional univariate power curves. Understanding the sources of this variance structure is essential for probabilistic power curve modelling. The following analysis quantifies which atmospheric parameters exert the greatest influence on power variability.

3.2. Parameter Importance Ranking: Relative Contribution to Power Variance

To systematically characterise the influence of atmospheric parameters on power variability, a variance-decomposition approach is employed. This analysis quantifies the marginal contribution of each parameter in $\{v, \alpha, \mathrm{TI}, g\mathrm{TI}\}$ to the overall structure of power variance, without assuming independence or controlling for parameter interactions. For each parameter x_i, the observed dataset is partitioned into groups defined by the unique values of x_i that appear in the simulation ensemble. Within each group, the across-seed variance of mean power σ_p^2 quantifying the 10-

Journal of Physics: Conference Series **3224** (2026) 022052 doi:10.1088/1742-6596/3224/2/022052

minute average power variance across stochastic realisations at each fixed atmospheric condition is averaged across all data points sharing the same value of x_i, yielding a group-wise conditional mean variance estimate. Parameter importance is then quantified as the variance of these group-wise conditional means, normalised to the maximum across all parameters:

$$I_i = \frac{\mathrm{Var}(E[\sigma_p^2 \mid x_i])}{\max_j \mathrm{Var}(E[\sigma_p^2 \mid x_j])} \times 100\%, \tag{1}$$

where $E[\sigma_p^2 \mid x_i]$ denotes the group-wise conditional mean of power variance, computed by averaging σ_p^2 across all observations sharing the same value of x_i. The variance $\mathrm{Var}(E[\sigma_p^2 \mid x_i])$ quantifies how much power variance changes, on average, as x_i ranges across its observed values. The parameter importance score is normalised to wind speed as the reference baseline, assigned a score of 100% by definition, ensuring meaningful comparison of secondary atmospheric effects.

Table 1. Parameter importance ranking for mean power and power variance.

Parameter	**Mean Power (%)**	**Power Variance (%)**
Wind speed	100.0	100.0
Turbulence intensity (TI)	< 0.001	7.4
Rotor turbulence intensity asymmetry (gTI)	< 0.001	0.83
Wind shear (α)	0.015	0.10

The ranking presented in Table 1 reveals a fundamental distinction between parameters controlling mean power and those controlling power variance. Wind speed dominates both dimensions, serving as the primary factor for both expected power output and its variability. For mean power, the other atmospheric parameters contribute negligibly ($< 0.02\%$), confirming that classical binned power curves adequately capture the mean response using wind speed alone. For power variance, however, a distinctly different hierarchy emerges: turbulence intensity contributes significantly with 7.4% relative importance, rotor turbulence intensity asymmetry contributes at 0.83%, and wind shear contributes minimally at 0.10%.

When gTI is interpreted as a fraction of wind speed's effect, its 0.83% ranking might suggest a negligible role in overall power curve uncertainty. However, this aggregate statistic suppresses substantial regime-dependent behaviour. Rotor turbulence intensity asymmetry, while modest in global importance, exhibits pronounced effects in specific atmospheric and operating regimes. This regime-dependence is precisely the heteroscedastic structure that input-dependent surrogate models are designed to capture. The following subsection demonstrates how the influence of gTI varies across the turbine's operating envelope by analysing the temporal variance of power, which provides a direct and physically interpretable measure of turbulence intensity asymmetry on short-term power dispersion across different operating regimes.

3.3. Rotor turbulence intensity asymmetry and heteroscedastic power variability

Although the parameter ranking identifies wind speed and turbulence intensity as the dominant global drivers of power curve uncertainty, this aggregate perspective does not capture regime-dependent behaviour that may be locally significant. To examine this, this section uses the temporal variance of power, which provides a direct and physically interpretable measure of short-term power fluctuation intensity across different operating regimes within a single 10-minute simulation. To isolate and quantify this behaviour, the uncertainty inflation factor is first defined as:

$$\Lambda(v, g\mathrm{TI}) = \frac{\sigma_{P_t}(v, g\mathrm{TI})}{\sigma_{P_t}(v, g\mathrm{TI} = 1.0)}, \tag{2}$$

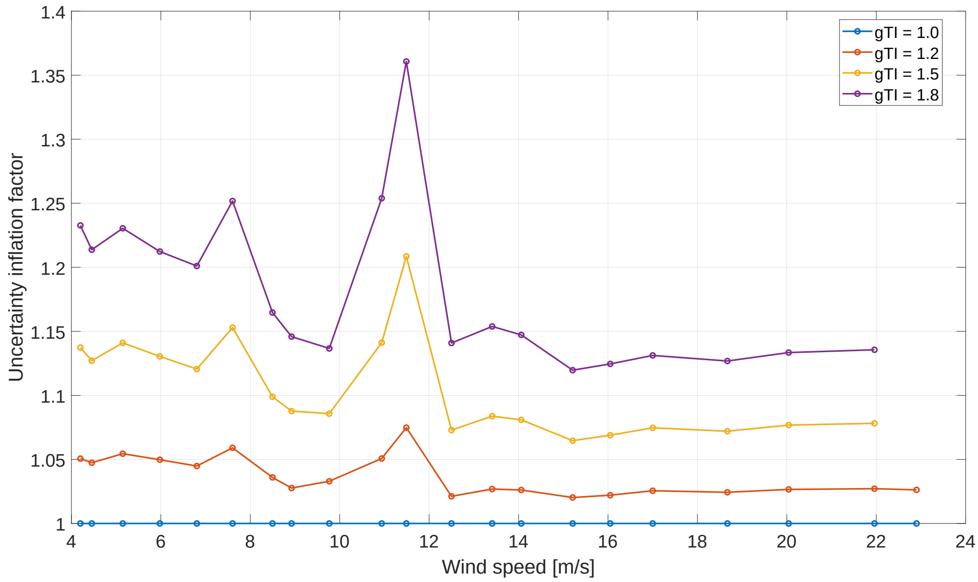

Figure 2. Effect of Rotor Turbulence Intensity Asymmetry on Power Variability

where σ_{P_t} is the standard deviation of instantaneous power, computed as the square root of temporal variance, at wind speed v for a given asymmetry level. Figure 2 shows that this factor increases with asymmetry level, reaching values approaching 1.3-1.35 for the highest asymmetry examined, indicating that rotor turbulence intensity asymmetry alone can increase temporal power variance by 30-35% relative to the symmetric case. Notably, the magnitude and structure of these effects are non-uniform across the operating range, illustrating the regime-dependent nature of gTI's influence.

To quantify the regime-dependence more explicitly, three representative operating cases spanning the turbine's operating envelope are examined: below-rated ($v \approx 6.0$ m/s), near-rated ($v \approx 10.5$ m/s), and above-rated ($v \approx 14.0$ m/s). For each regime, standard deviation of power fluctuation is computed across the entire range of gTI while other atmospheric parameters are held constant at representative stable-atmosphere values ($\alpha = 0.30$, TI $= 0.16$). These parameter combinations represent a stratified, stable turbulent boundary layer where rotor-scale non-uniformity is pronounced and its effects on power variability are most clearly observable, as presented in Table 2.

Table 2. Uncertainty inflation factor across operating regimes and rotor turbulence intensity asymmetry levels. Computed at fixed atmospheric parameters ($\alpha = 0.30$, TI $= 0.16$) to isolate gTI effects.

Regime	**gTI=1.0** (MW^2)	**gTI=1.2**	**gTI=1.5**	**gTI=1.8**	**% Change** (1.0 $\rightarrow$ 1.8)
Below-rated ($v \approx 6.0$ m/s)	1.066	1.155	1.313	1.489	+39.7%
Near-rated ($v \approx 10.5$ m/s)	1.202	1.286	1.428	1.649	+37.2%
Above-rated ($v \approx 14.0$ m/s)	0.266	0.275	0.299	0.331	+24.4%

The scenario analysis demonstrates that power variance increases monotonically with gTI across all operating regimes, indicating that rotor turbulence intensity asymmetry consistently amplifies power output dispersion. The magnitude of this amplification, however, varies substantially across regimes. Below-rated operation exhibits a 39.7% variance increase over the gTI range (1.0 to 1.8), while near-rated operation shows 37.2%, and above-rated operation shows a more modest 24.4% increase. This regime-dependent response reflects the turbine's physical state at each operating point. Below-rated and near-rated regimes are characterised by power production tightly coupled to inflow conditions. The turbine operates below its control saturation point, and power output responds directly and sensitively to variations in wind speed, wind shear, and turbulence. In this operating regime, rotor-scale turbulence non-uniformity translates directly into spatial variation in blade aerodynamic loads, leading to pronounced dispersion in instantaneous power output. Consequently, rotor turbulence intensity asymmetry exerts a strong amplifying effect on power variance. Above-rated operation transitions to a fundamentally different control regime. Active pitch control engages to maintain constant power output despite increasing wind speeds. The control system continuously adjusts blade pitch angles to regulate aerodynamic loads and stabilise power, thereby filtering fluctuations in wind speed and turbulence intensity. The 24.4% variance increase above-rated reflects this active filtering, while pitch control partially attenuates but does not eliminate the effects of rotor-scale turbulence intensity asymmetry on power dispersion.

This regime-dependent variation in gTI sensitivity illustrates a broader characteristic of heteroscedastic behaviour: the influence of any atmospheric parameter on power variance depends fundamentally on the turbine's operating state. Conventional univariate power curves, which apply constant uncertainty bands uniformly across the operating range, cannot capture this state-dependent variability. Similarly, simple parameter-specific adjustments (additive or multiplicative corrections applied to mean power uncertainty) are inadequate because they cannot distinguish between operating regimes. A probabilistic surrogate model that explicitly represents input-dependent variance is necessary to capture the heteroscedastic effects. The following sections present and validate such a framework, demonstrating that a two-stage heteroscedastic Gaussian Process Regression model can automatically learn the regime-dependent patterns identified empirically in this full-factorial analysis.

4. Two-stage Gaussian Process Regression Model

The relationship between wind speed, atmospheric parameters, and power output is non-linear and heteroscedastic, with output variance changing systematically across wind speeds and atmospheric conditions. While joint heteroscedastic GP models can capture this behaviour, their likelihood surfaces are highly complex and difficult to optimise. To address this, we employ a decoupled two-stage framework: the first GP models the mean power curve, and a second GP models the input-dependent variance from the residuals of the first stage.

In the first stage, observed mean power y at environmental conditions $(\mathbf{x}) = (v, \alpha, \mathrm{TI}, g\mathrm{TI})$ is expressed as

$$y = f(\mathbf{x}) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_n^2), \tag{3}$$

where $f(\mathbf{x})$ is a latent function with a GP prior,

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \tag{4}$$

Here, $m(\mathbf{x}) = E[f(\mathbf{x})]$ is the mean function of the GP prior, encoding the prior expectation of the latent function before observing any data, and $k(\mathbf{x}, \mathbf{x}') = E[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]$ is the covariance kernel function, which quantifies the statistical similarity between function values at any two input points $\mathbf{x}$ and $\mathbf{x}'$ and encodes the assumed smoothness and correlation structure of f. Given N training samples $\{\mathbf{X}, \mathbf{y}\}$, the predictive distribution at a new input $\mathbf{x}_*$ is Gaussian,

$$p(y_* \mid \mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \mathcal{N}(y_* \mid \mu_*(\mathbf{x}_*), \sigma_*^2(\mathbf{x}_*)), \tag{5}$$

with mean and variance

$$\mu_*(\mathbf{x}_*) = m(\mathbf{x}_*) + \mathbf{k}_*^T(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}(\mathbf{y} - m(\mathbf{X})), \tag{6a}$$

$$\sigma_*^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{k}_*, \tag{6b}$$

where $\mathbf{K} = k(\mathbf{X}, \mathbf{X})$ and $\mathbf{k}_* = k(\mathbf{X}, \mathbf{x}_*)$. This yields the mean prediction $\hat{\mu}(\mathbf{x})$.

In the second stage, heteroscedastic variance is modelled from the squared residuals of the first stage, $\mathbf{r} = \mathbf{y} - \hat{\mu}(\mathbf{X})$. To enforce positivity and improve numerical stability, their logarithms are used as targets:

$$g(\mathbf{x}) \sim \mathcal{GP}(m_g(\mathbf{x}), k_g(\mathbf{x}, \mathbf{x}')), \quad \mathbf{z} = \log(\mathbf{r}^2 + \epsilon), \tag{7}$$

where $\mathbf{x} = (v, \alpha, \mathrm{TI}, g\mathrm{TI})$ denotes the vector of atmospheric inputs and ϵ is a small regularisation constant. This produces variance predictions $\hat{g}(\mathbf{x}_*)$ at new conditions.

Finally, the outputs of the two stages are combined into a heteroscedastic predictive distribution:

$$p(y_* \mid \mathbf{x}_*) = \mathcal{N}(y_* \mid \hat{\mu}(\mathbf{x}_*), \exp(\hat{g}(\mathbf{x}_*))). \tag{8}$$

This decoupling of mean and variance estimation into two stable regression tasks, both conditioning on the same multivariate input space, provides a flexible and robust framework for capturing both the average power curve and its input-dependent uncertainty. Critically, the second-stage GP receives the same atmospheric parameter inputs as the first stage, enabling the model to learn how variance responds to variations in wind shear, turbulence intensity, and rotor turbulence intensity asymmetry.

5. Heteroscedastic Model Evaluation and Performance Analysis

This section evaluates the heteroscedastic GPR framework using the full-factorial dataset introduced in Section 2. The analysis examines how the univariate and multivariate models capture mean power and input-dependent variance, assessing the value of atmospheric conditioning for reducing prediction uncertainty. The results demonstrate that conventional wind-speed-based models neglect a substantial portion of resolvable variability and justify the need for multivariate heteroscedastic approaches.

5.1. Model Performance: Univariate Versus Multivariate Framework

A univariate heteroscedastic GPR model using wind speed as the only input provides a baseline representing standard deterministic power curve approaches. The multivariate model extends the input space to include wind shear exponent (α), hub-height turbulence intensity (TI), and rotor turbulence intensity asymmetry (gTI). Both models are trained and validated on identical train-test splits to ensure fair comparison.

Performance metrics are summarised in Table 3. The univariate model achieves an RMSE of 184.88 kW and $R^2 = 0.9970$, confirming that wind speed captures the dominant trend in mean power output. However, the mean prediction interval width (MPIW) of 332.83 kW (approximately 2.2% of rated power) reveals substantial unresolved uncertainty: variations in turbulence intensity, wind shear, and rotor-scale non-uniformity appear as residual scatter, forcing the model to adopt conservative bounds across all operating conditions.

The multivariate model reduces RMSE by 27.6% (133.91 kW) and MPIW by 25.3% (248.64 kW). This improvement reflects a restructuring of uncertainty rather than uniform tightening. The model learns to narrow prediction intervals under stable inflow conditions (low TI, uniform gTI) and widen them under disturbed conditions, distinguishing operating states aggregated in

Table 3. Model performance comparison: univariate versus multivariate.

Metric	1D (Univariate)	4D (Multivariate)	Improvement
RMSE (kW)	184.88	133.91	−27.60%
MPIW (kW)	332.83	248.64	−25.30%
R^2	0.9970	0.9994	+0.24%

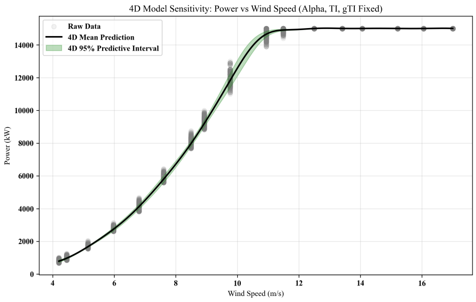

(a) Wind-speed power curve for multivariate GPR model

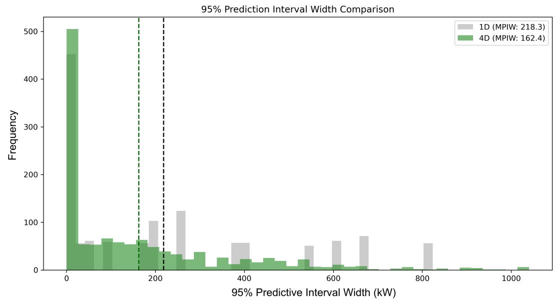

(b) MPIW distributions: univariate vs. multivariate.

Figure 3. Multivariate GPR model performance assessment.

the univariate formulation. The marginal improvement in R^2 (+0.24%) indicates that atmospheric parameters refine local predictions without altering the global variance structure already captured by wind speed alone.

Figure 3 further illustrates these effects. Figure 3(a) shows a representative wind-speed slice with atmospheric parameters held constant. The predicted mean maintains the familiar sigmoidal curve, demonstrating that additional inputs do not disrupt the dominant wind-speed dependence of power production. Instead, their influence manifests in adaptive prediction intervals that vary with operating regime. Figure 3(b) compares the distributions of 95% prediction interval widths. The univariate distribution is broader with higher average width, reflecting unresolved atmospheric variability. The multivariate distribution is shifted toward narrower intervals with fewer extremes, indicating that atmospheric conditioning resolves a portion of the variability treated as irreducible in univariate models.

From a practical perspective, this distinction matters for design assessments. Wind-speed-only models treat unexplained variability as irreducible, inflating uncertainty in downstream calculations. The multivariate framework, by contrast, resolves resolvable variability through physically interpretable inputs, yielding justified uncertainty reduction under favourable conditions and appropriate uncertainty inflation under challenging conditions. These results confirm that incorporating atmospheric parameters, including rotor turbulence intensity asymmetry, into the GPR framework resolves structured variability that wind-speed-only models treat as irreducible.

6. Conclusions

This study presents a systematic assessment of how rotor turbulence intensity asymmetry (gTI) contributes to wind turbine power output variability, addressing a gap in current power curve modelling practice. The parameter importance analysis reveals a fundamental distinction between parameters controlling mean power and those controlling power variance: while wind speed dominates both, gTI contributes more to power variance than wind shear despite receiving considerably less attention in the literature. This global finding is complemented by a regime-dependent analysis showing that gTI amplifies short-term power variance by 37-40% in below-rated and near-rated operation, and by 24% above rated where pitch control partially attenuates but does not eliminate its effect. Together, these findings establish that gTI represents a structured and quantifiable source of power output uncertainty that is absent from current IEC standards and existing power curve models.

The proposed heteroscedastic GPR framework demonstrates that these regime-dependent patterns, masked in global parameter rankings, are learnable by a surrogate model. By separately modelling mean power and input-dependent variance as functions of wind speed, shear, turbulence intensity, and rotor turbulence intensity asymmetry, the framework achieves a 27.6% reduction in RMSE and 25.3% reduction in prediction interval width relative to univariate baselines. The restructuring of uncertainty reflects the empirically identified regime-dependent structure, validating the physical relevance of the surrogate model. However, the current framework is limited to the atmospheric conditions explored in the full-factorial design; extending it to additional parameters would enhance applicability across diverse site conditions. A natural extension would be a systematic comparison against IEC-prescribed wind shear and turbulence intensity correction procedures, using a continuously sampled experimental design that permits direct benchmarking against correction-adjusted power curve baselines. These developments would further strengthen the practical utility of the proposed framework for energy yield assessment, where rotor-scale turbulence variation is pronounced and design-phase uncertainty quantification relies heavily on simulation-based assessments.

Acknowledgements

This work is supported by Science Foundation Ireland (SFI) grant no. 20/FFP-P/8702.

References

[1] Veers P, Dykes K, Basu S, Bianchini A, Clifton A, Green P, Holttinen H, Kitzing L, Kosovic B, Lundquist J K *et al.* 2022 Grand Challenges: wind energy research needs for a global energy transition *Wind Energy Science Discussions* **2022** 1–8

[2] Veers P, Bottasso C L, Manuel L, Naughton J, Pao L, Paquette J, Robertson A, Robinson M, Ananthan S, Barlas T *et al.* 2023 Grand challenges in the design, manufacture, and operation of future wind turbine systems *Wind Energy Science* **8** 1071–1131

[3] International Electrotechnical Commission 2022 Wind energy generation systems — Part 12-1: Power performance measurement of electricity producing wind turbines Tech. Rep. IEC 61400-12-1:2022 International Electrotechnical Commission

[4] Elliott D L and Cadogan J B 1990 Effects of wind shear and turbulence on wind turbine power curves Tech. rep. Pacific Northwest Lab., Richland, WA (USA)

[5] Clifton A and Wagner R 2014 Accounting for the effect of turbulence on wind turbine power curves *Journal of Physics: Conference Series* vol 524 p 012109

[6] Sebastiani A, Angelou N and Peña A 2024 Wind turbine power curve modelling under wake conditions using measurements from a spinner-mounted lidar *Applied Energy* **364** 122985

[7] Wharton S and Lundquist J K 2012 Atmospheric stability affects wind turbine power collection *Environmental Research Letters* **7** 014005

[8] Olaofe Z O and Folly K A 2013 Wind energy analysis based on turbine and developed site power curves: A case-study of darling city *Renewable Energy* **53** 306–318

[9] International Electrotechnical Commission 2019 Wind energy generation systems — Part 1: Design requirements Tech. Rep. IEC 61400-1:2019 International Electrotechnical Commission

[10] IEA Wind Power Curve Working Group 2020 The Power Curve Working Group's assessment of wind turbine power performance models *Wind Energy Science* **5** 199–250

[11] McKay M D, Beckman R J and Conover W J 1979 A comparison of three methods for selecting values of input variables in the analysis of output from a computer code *Technometrics* **21** 239–245

[12] Gaertner E, Ballow F, Bortolotti J, Fleming R, Folk N, Goossels S, Hall M, Hart E, Jonas B, Kretschmer T *et al.* 2020 Definition of the IEA 15-Megawatt offshore reference wind turbine Tech. Rep. NREL/TP-5000-75698 National Renewable Energy Laboratory Golden, CO

[13] Abbas N J, Zalkind D S, Pao L and Wright A 2022 A reference open-source controller for fixed and floating offshore wind turbines *Wind Energy Science* **7** 53–73