DONAU LABS

Exploratory Wind-Turbine Modelling Study

Original Problem 5 - Scope, interim results and industrial direction

Interim report · 20 November 2026 · Version 1.7

Original Problem 5 asks whether turbines produce systematically different power under comparable wind and operating conditions. We used public Kelmarsh SCADA to explore that question and build a model of expected ten-minute power. A separate external test examined whether that fitted model could predict power at Penmanshiel. The data sources are listed in [1] and [2].

These are two linked assessments: prediction error measures how closely the model estimates recorded power; a peer comparison measures relative differences among turbines. Good prediction alone cannot establish a turbine fault or a recoverable loss. This exploratory study provides methods and findings to discuss with site specialists before defining an industrial application.

Operational value remains unproven

This phase is an exploratory modelling exercise. It produced working prediction models and identified changing relative power patterns. We have yet to connect these findings to a confirmed fault, a maintenance decision, an avoidable production loss or measurable benefit to an operator. The Penmanshiel test establishes a limitation of the fitted model; its direct operational value remains unproven.

How the study developed

Stage	Result	What we learned
2016 Kelmarsh prototypes	Matched-record RMSE: wind-only 87.5 kW; seven inputs 48.03 kW; five inputs 48.01 kW.	The nonlinear model estimates power more closely when operating signals are included. Removing the two wind-direction terms changed error very little.
2017 Kelmarsh time test	Frozen five-input RMSE 62.5 kW, or 9.0% of mean power, after status screening.	The model retained predictive value on a later period at the same site. Subsequent exploratory refinements reused 2017; future model changes need a fresh test period.
Penmanshiel transfer test	2017 pitch data were empty. In 2024, 151,803 screened, range-compatible readings from WT11-WT15 gave RMSE 157.5 kW (78.5% of their mean power).	Correlation was 0.968, but average overprediction was 133.7 kW. The frozen Kelmarsh model transferred poorly to this tested MM82 configuration.

The Penmanshiel score uses the fixed mean prediction of four frozen Kelmarsh peer models; individual models also showed substantial errors. Its range-restricted subset had mean power of 200.7 kW and favoured lower rotor speeds. Scores across the two sites have different operating coverage and denominators.

Limits of the present evidence

The work uses public observational data and a limited family of regression models. Sensor mounting and channel equivalence remain partly unresolved; status records have incomplete durations. Rotational-speed and pitch inputs may absorb a performance difference. Correlation, low error and narrow conditional intervals do not establish a physical fault, recoverable energy, commercial benefit or deployment readiness.

Wind-Turbine Prediction Study

Original Problem 5 - Multi-year findings and Vietnam deployment

Interim report · 20 November 2026 · Version 1.7

The follow-up compares Kelmarsh WT1, WT2, WT4 and WT5 across every published year from 2016 to 2024. These four turbines share the MM92 model and 78.5 m hub height. Values below are each turbine’s power difference from the other three after matching reported wind speed and direction.

Year	Comparable intervals	WT1	WT2	WT4	WT5
2016	32,609	-5.90%	+1.96%	+0.55%	+3.51%
2017	49,330	-4.84%	+1.39%	+0.08%	+3.46%
2018	47,936	-0.90%	+0.97%	-2.07%	+2.03%
2019	50,858	+0.01%	+0.83%	-2.94%	+2.14%
2020	49,544	+0.10%	+0.90%	-2.40%	+1.42%
2021	45,765	+1.82%	+4.31%	-6.14%	+0.16%
2022	42,647	+1.56%	-2.54%	-0.73%	+1.74%
2023	39,718	-1.20%	-1.32%	+1.46%	+1.07%
2024	46,313	-1.22%	-3.26%	+3.89%	+0.67%

Interpretation

The initial 5-6% WT1 shortfall is specific to the early period. It narrows in 2018, is close to zero in 2019-2020, becomes positive in 2021-2022 and is about -1.2% in 2023-2024. WT4 is relatively lower in 2018-2022, while WT2 is lower in 2022-2024. The common-condition check on page 3 preserves these broad changes. A changing gap can reflect the target turbine, its peers, exposure, controls or measurement calibration.

The results support time-dependent differences among comparable turbines. They provide a basis for examining maintenance, control and sensor-change histories around changes in relative performance. Each percentage uses weighted peer power as denominator; these estimates do not identify a physical cause or quantify annual recoverable energy.

2016 starts on 15 April. The other years cover their available records within each calendar year. Counts refer to simultaneous timestamps across the four turbines before excluding cells with fewer than 30 readings per turbine.

Persistence, uncertainty and limitations

Year	WT1 difference	Approximate 95% interval	WT1 at fixed common conditions
2016	-5.90%	-6.18% to -5.61%	-6.15%
2017	-4.84%	-5.17% to -4.44%	-5.39%
2018	-0.90%	-1.39% to -0.44%	-1.24%
2019	+0.01%	-0.17% to +0.22%	-0.02%
2020	+0.10%	-0.10% to +0.33%	-0.00%
2021	+1.82%	+0.90% to +2.77%	+2.20%
2022	+1.56%	-0.02% to +3.67%	+1.77%
2023	-1.20%	-1.61% to -0.74%	-1.64%
2024	-1.22%	-1.44% to -0.94%	-1.42%

The confidence intervals use 3,000 positive block-weight bootstrap draws with seven-day blocks shared across turbines. Fourteen-day blocks provide a sensitivity check. Intervals are pointwise and conditional on fixed supported cells, measured signals and screening; they exclude systematic sensor uncertainty.

The fixed-conditions check uses the same 73 wind-speed/direction cells and weights across all years. This helps distinguish a changing annual mix of wind conditions from changes in the relative power relationship.

Method

Compare mean power in 1 m/s reported wind-speed bins and 45-degree direction sectors. Require 30 readings per turbine in each cell. Weight cells by their minimum turbine count, using the same weights for all four turbines. Exclude intervals overlapping documented restrictions, with a ten-minute margin for uncertain timestamp alignment. Repeat with stricter ambiguity screens.

Source integrity

Publisher MD5 checksums were verified for all seven newly downloaded annual archives. Later-year loading records missing timestamps and duplicate rows. Repeated readings are consolidated only where populated values agree. Conflicting timestamps are excluded. Across 2018-2024, the loader recorded 16,977,600 duplicate timestamp rows and 0 conflicting turbine timestamps.

The earlier Penmanshiel test showed substantial direct-transfer error for the five-input Kelmarsh model. Local calibration is recommended for a new turbine configuration and site. The environmental peer comparison and the operational prediction model serve complementary diagnostic purposes.

From exploration to a site pilot

How the evidence informs the next approach

Transfer to Penmanshiel gave substantial error. Turbine type, geometry, site and measurement year changed together, so the test cannot isolate location as the cause. It establishes a limitation of the current fitted model.

We propose local calibration for comparable turbines, followed by separate local validation. Astolfi et al. support site- and turbine-specific SCADA models [3]. DNV describes comparisons with normal turbine behaviour and their measurement limitations [4]. These sources support this next experiment; operational usefulness still requires demonstration.

The sequence should begin with defining the site's operational problem with its experts. Local calibration remains a proposed next experiment. It has not been performed for Penmanshiel or a Vietnamese site.

Information needed for a Vietnamese pilot

What is needed	Purpose
Configuration and measurements	Turbine IDs, model, rating, rotor and hub dimensions; power/wind units, averaging, timestamps and sensor/correction definitions.
Operating history	Status meanings and intervals, restrictions/setpoints, commissioning, maintenance and sensor/control changes.
Qualified inputs	Wind and power for the baseline; direction and temperature for environmental checks; rotational speeds and pitch for operational interpretation.
Calibration and test coverage	Represent local seasonal and operating conditions. Reserve separate test records before tuning; a full annual cycle is desirable.
Pilot monitoring	Report coverage, error, bias and uncertainty; review sensor/control changes and input drift.

Define the problem and intended decision

The next phase should begin with a concrete operator problem and supporting site evidence. Donau Labs should discuss the suspected issue with operations, maintenance and performance specialists and the intended decision maker. Agree the action a result would trigger, the smallest useful effect and the maintenance or field evidence needed to judge it. Then specify calibration, validation and pilot acceptance criteria. Further generic model fitting alone is unlikely to establish operational value.

Immediate input

Actual turbine-level SCADA and metadata from a selected Vietnamese site are required. The regional weather packages previously reviewed provide insufficient information for this turbine-power comparison. Donau Labs would lead qualification, modelling and internal validation.

References

[1] Cubico. Kelmarsh wind farm data, v4 (2025). Static metadata and annual SCADA.

[2] Cubico. Penmanshiel wind farm data, v3 (2025). Static metadata and SCADA.

[3] Astolfi, Castellani, Lombardi & Terzi (2021). Multivariate SCADA Data Analysis Methods for Real-World Wind Turbine Power Curve Monitoring. Energies 14(4), 1105. See Introduction.

[4] DNV. Rectifying underperforming wind farms. See ten-minute SCADA analysis and nacelle-wind limitations.