Reliable Wind Power Forecasts with Conformal Prediction

Carl Somers, F. Javier Rubio, Domna Ladopoulou
Department of Statistical Science, University College London

UCL

Motivation

The grid depends on balancing supply and demand. As wind replaces dispatchable sources like gas, operators rely on accurate forecasts for unit commitments. Without reliable uncertainty estimates, they must risk blackouts or keep costly, carbon-heavy plants on standby. Point forecasts ignore this uncertainty, creating demand for probabilistic outputs. Yet common methods (MCMC, ensembles, or Bayesian approaches) often lack coverage guarantees, impose strong assumptions, or require expensive training.

> **Goal:** We want to develop **predictive intervals** that have valid **coverage guarantees**, with **minimal assumptions**.

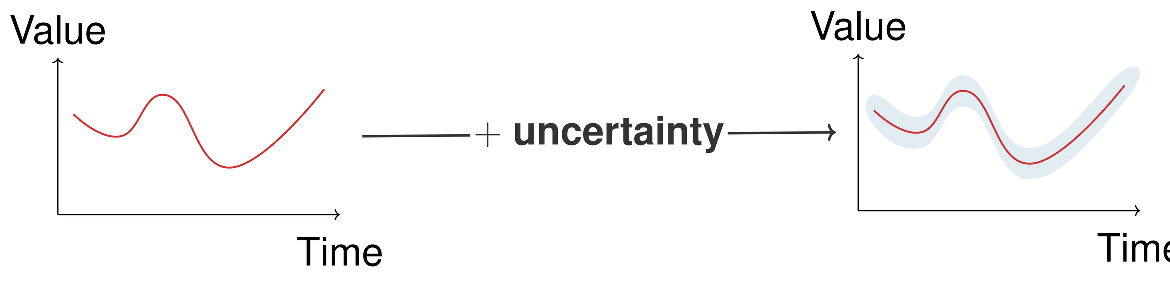

Point Predictions — **Probabilistic Output**

Conformal Prediction (Vovk et al., 2005)

Conformal prediction is a model-agnostic framework that converts any point predictor into a prediction interval with guaranteed coverage. It has seen wide success, from LLMs to image classification, with split conformal prediction (SCP) being the most common and practical variant.

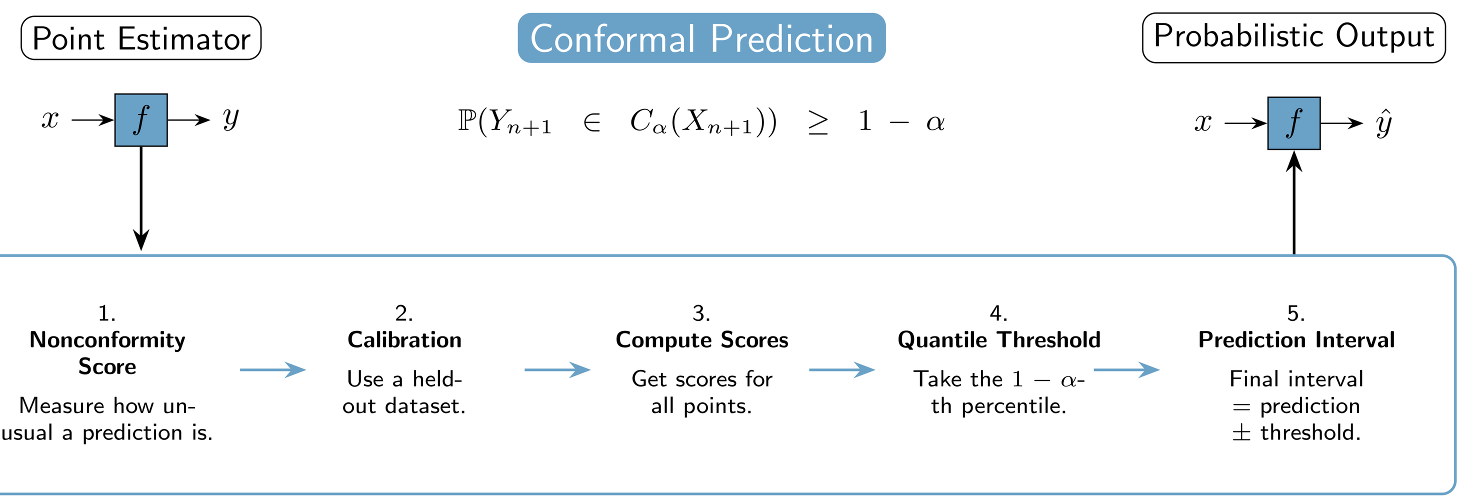

This involves gathering residual non-conformity scores (often absolute errors) from a calibration set and forming a predictive set via the empirical quantile. Under the assumption of exchangeability (data points are order invariant) this guarantees coverage of at least $1-\alpha$. But wind power time series are volatile and temporally dependent, breaking this assumption.

> **But:** Classic CP assumes data points are independent or *exchangeable*. Wind is not - it's volatile and correlated over time.

Adaptive Conformal Inference (Gibbs & Candes, 2021)

This is especially apt for grid operators in **online forecasting**, where conditions change constantly and retraining models isn't always feasible. The method adapts the α-quantile through online optimization:

$$\alpha_{t+1} = \alpha_t + \gamma\left(\alpha - \mathbb{I}_{y\notin\hat{C}(x_t)}\right),$$

where α is the user specified **miscoverage** (e.g 10%), α_t is the adaptive quantile based on previous miscoverage, and γ the **learning rate**, governing the speed of adaption.

Reducing sensitivity to learning rate

- **Expert based aggregation ACI (AgACI)**, runs a grid of γs, and runs an optimisation algorithm to select the optimal weighting for both upper and lower bound individually. Empirically strong, no guarantees.
- **Dynamically tuned ACI (DtACI)** uses a sub-model to *meta-learn* the best γ from a grid. Guarantees over some window to select the retrospective best performing γ.

Point Forecasting Models

While conformal wrappers are model agnostic, their width is only as good as the point forecasts it contains. We demonstrate using three distinct forecasting models.

- **Spatio-Temporal Graph Convolutional Network (ST-GCN)** - A state-of-the-art model chosen to capture the complex spatio-temporal dynamics of the wind farm.
- **Quantile Regression Light Gradient Boosting Machine (LGBM-QR)** - A commonly used, yet extremely powerful probabilistic machine learning approach.
- **Auto-Regressive Integrated Moving Average (ARIMA)** - Traditional univariate statistical time series model benchmark.

Evaluation

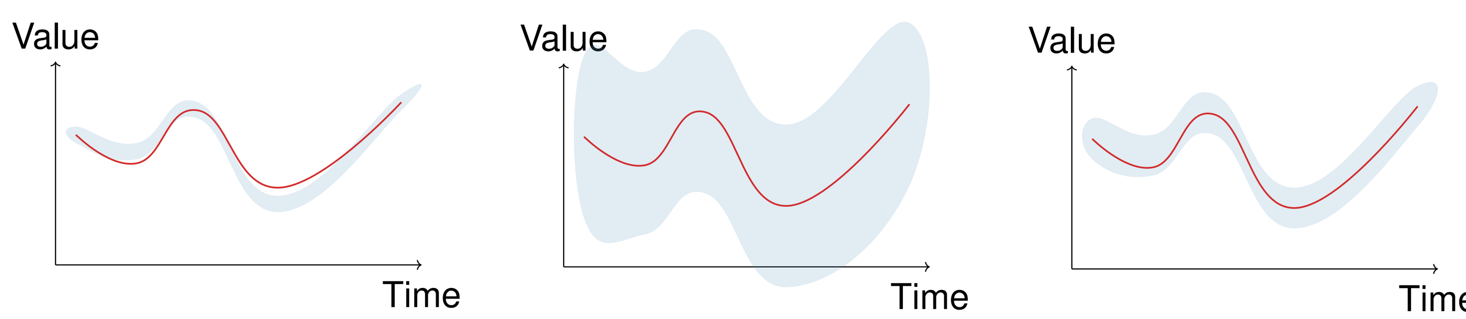

Precise but Unreliable (Low Coverage) — **Reliable but Useless (Wide Width)** — **The Goal: Reliable & Precise**

We observe both **empirical coverage** (proportion of time the true value is within the prediction interval, with 90% target) and the interval **width**. This is an inherent trade-off between operationally useful intervals and reliable ones.

Case Study: Kelmarsh Wind Farm

- **Location:** 12.3MW Kelmarsh Wind Farm, Northamptonshire, UK.
- **Data:** 2 years (2022–2024) of on-site SCADA & NWP forecasts.
- **Evaluation:** online, expanding-window evaluation on a 1-year test set.

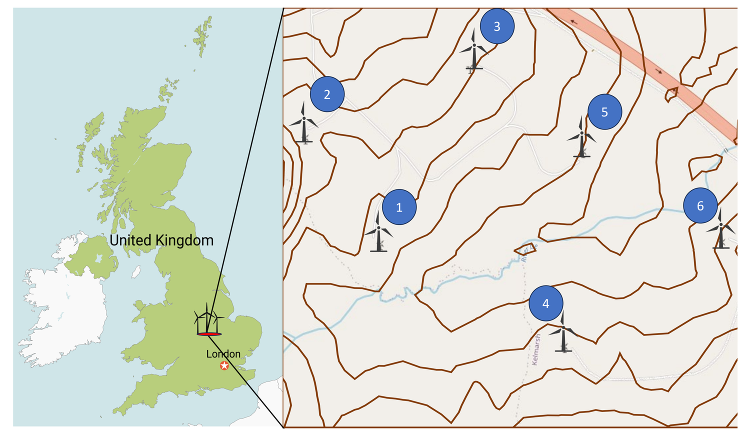

Results

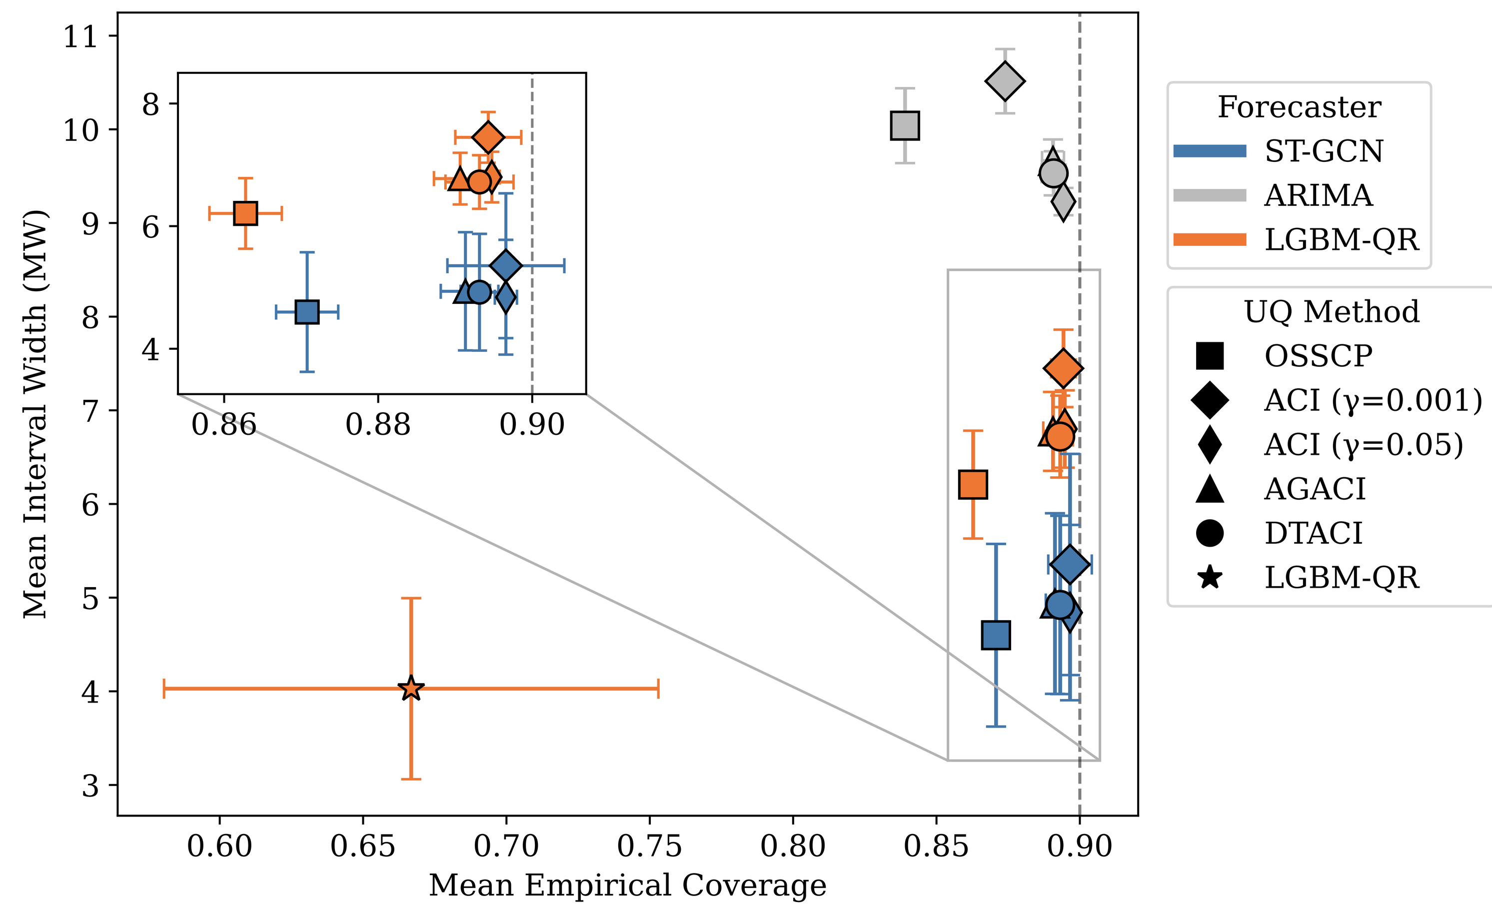

- ACI-based methods (●, ▲, ◆) hit the 90% coverage target, unlike the non-adaptive online SCP baseline (■) and Quantile Regression (⋆), which under-cover.
- The stronger ST-GCN forecaster yields sharper intervals than ARIMA, showing the importance of the underlying point forecast.

Key Points

> - **Adaptive Conformal Inference (ACI)** delivers reliable 90% intervals where standard methods fail.
> - **Better base forecasts** (e.g., ST-GCN) yield sharper, more useful intervals.
> - **Trustworthy uncertainty** improves grid stability and speeds renewable integration.

Code ⟶
Paper ⟶
Presentation ⟶