An Interpretable Probabilistic Model for Short-Term Solar Power Forecasting

Georgios Mitrentsis

Outline

1. Introduction
2. Our approach
3. Natural gradient boosting (NGBoost)
4. SHAP method
5. Power and meteorological data
6. Results
7. Conclusion

Introduction

Introduction

Motivation

The ongoing transition of the power system towards a fossil-fuel-free system has led to a wide **integration of renewable energy sources** (RES) worldwide.

The **stochastic nature** of RES power induced by volatile weather conditions hinders the **reliable electricity supply**.

RES are typically connected to the grid through **power electronics** leading to declining system inertia. The system operates close to **stability margins**.

Accurate and reliable power **forecasting** can alleviate those challenges allowing for a **large-scale RES integration.**

Introduction

Current issues - Forecasting

- It is based on complex machine learning models which are considered **black-boxes**.
- The **more complex** become those models, the **harder** it gets to **understand** them and explain their results.

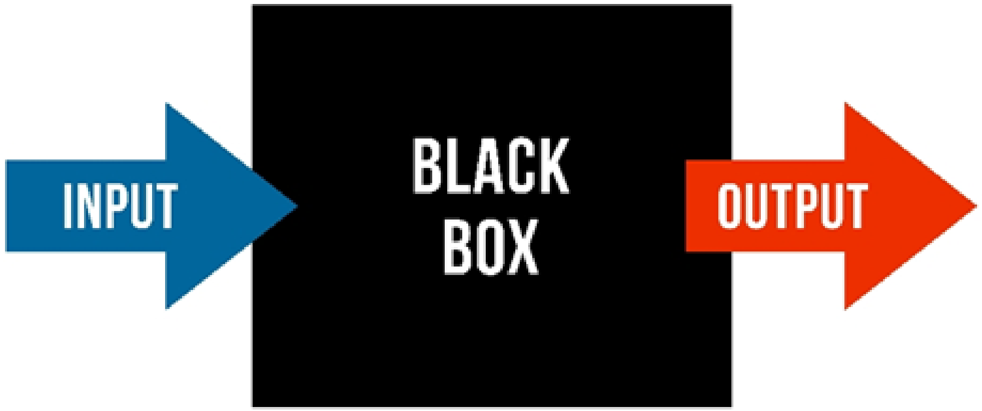

Introduction

Current issues - Applications

System Operators

Power Traders

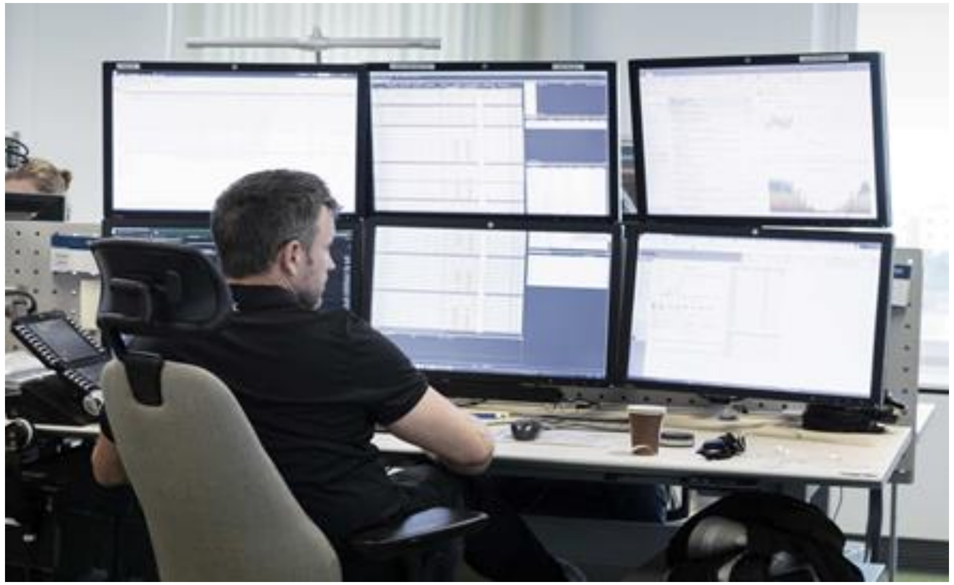

Even at **big technology companies**, many bugs in machine learning pipelines **may not be discovered**[1]!

[1] M. Zinkevich, “*Rules of machine learning: Best practices for ML engineering*,” 2017.

Our approach

Our approach

Schematic representation

- We propose the application of the natural gradient boosting (**NGBoost**) algorithm for yielding **probabilistic** PV power **forecasts**.

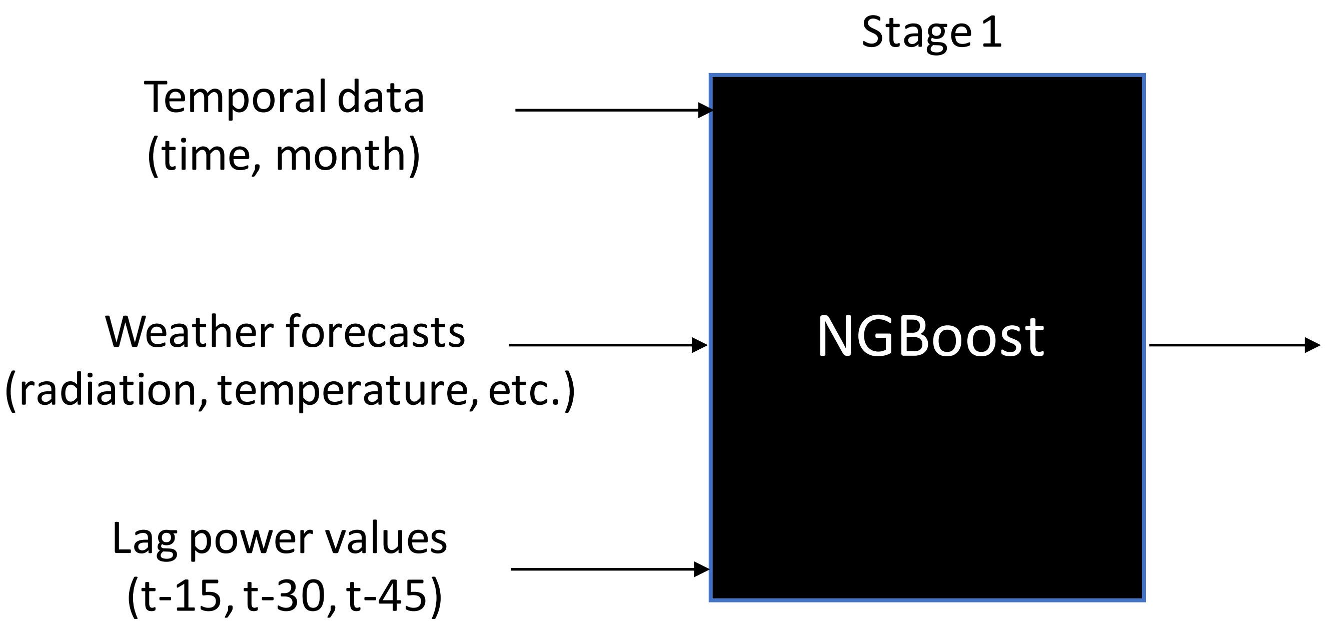

Our approach

Schematic representation

- We propose the application of the natural gradient boosting (**NGBoost**) algorithm for yielding **probabilistic** PV power **forecasts**.

- We combine it with Shapley additive explanation (**SHAP**) values, which is the unique consistent attribution method that can provide theoretical optimal **explanations** about the predictions of a model.

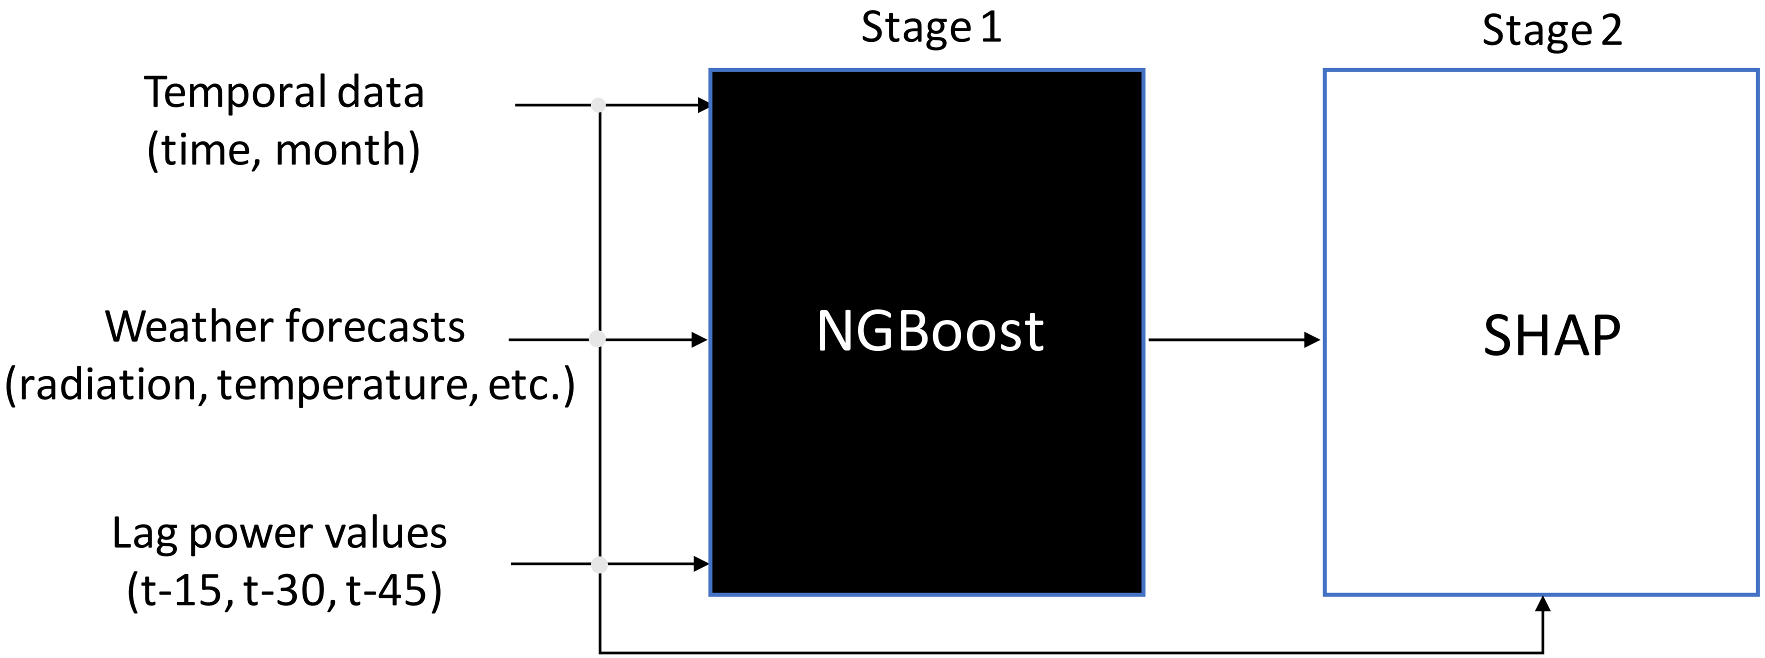

Our approach

Schematic representation

- We propose the application of the natural gradient boosting (**NGBoost**) algorithm for yielding **probabilistic** PV power **forecasts**.

- We combine it with Shapley additive explanation (**SHAP**) values, which is the unique consistent attribution method that can provide theoretical optimal **explanations** about the predictions of a model.

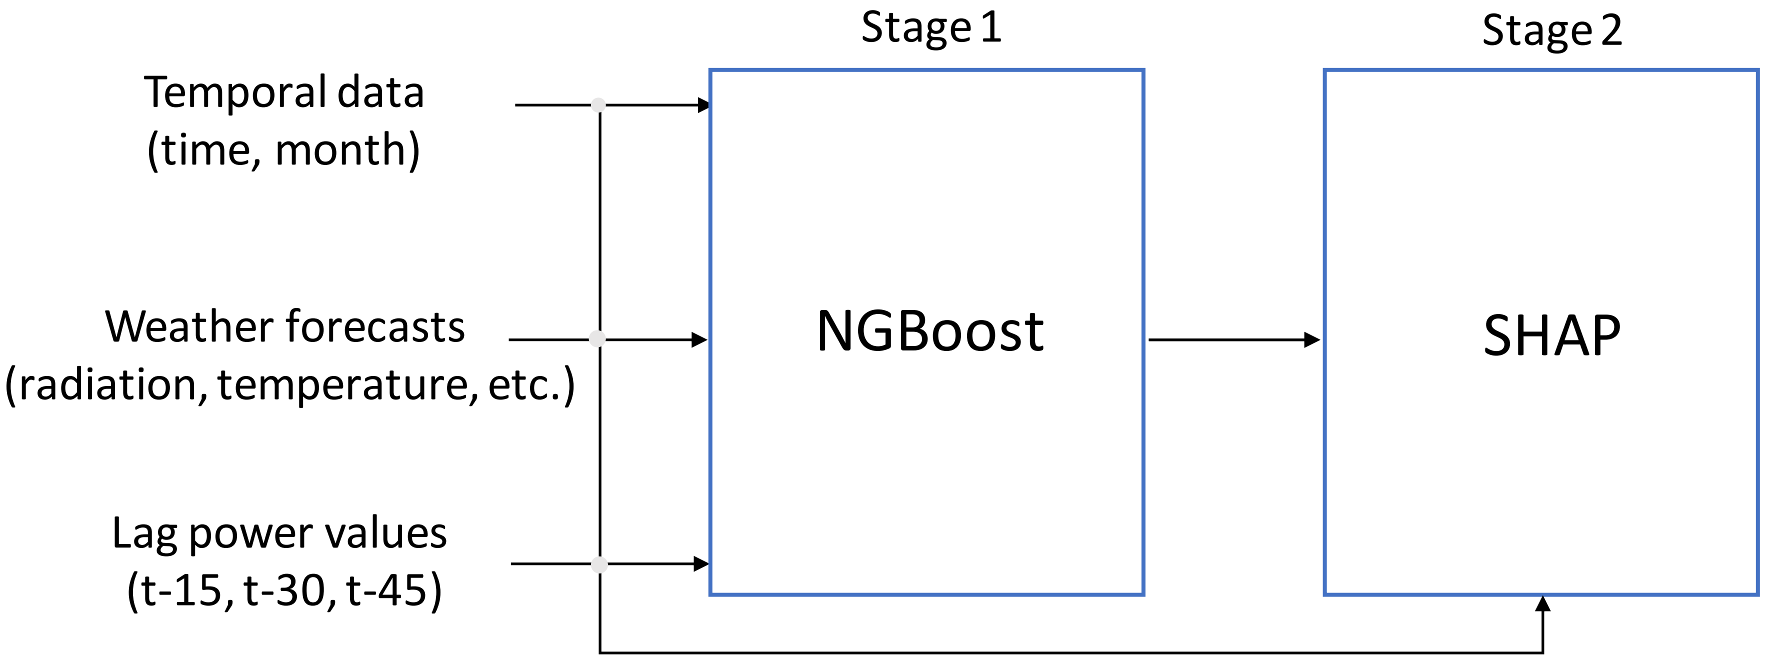

Our approach

Schematic representation

- We propose the application of the natural gradient boosting (**NGBoost**) algorithm for yielding **probabilistic** PV power **forecasts**.
- We combine it with Shapley additive explanation (**SHAP**) values, which is the unique consistent attribution method that can provide theoretical optimal **explanations** about the predictions of a model.

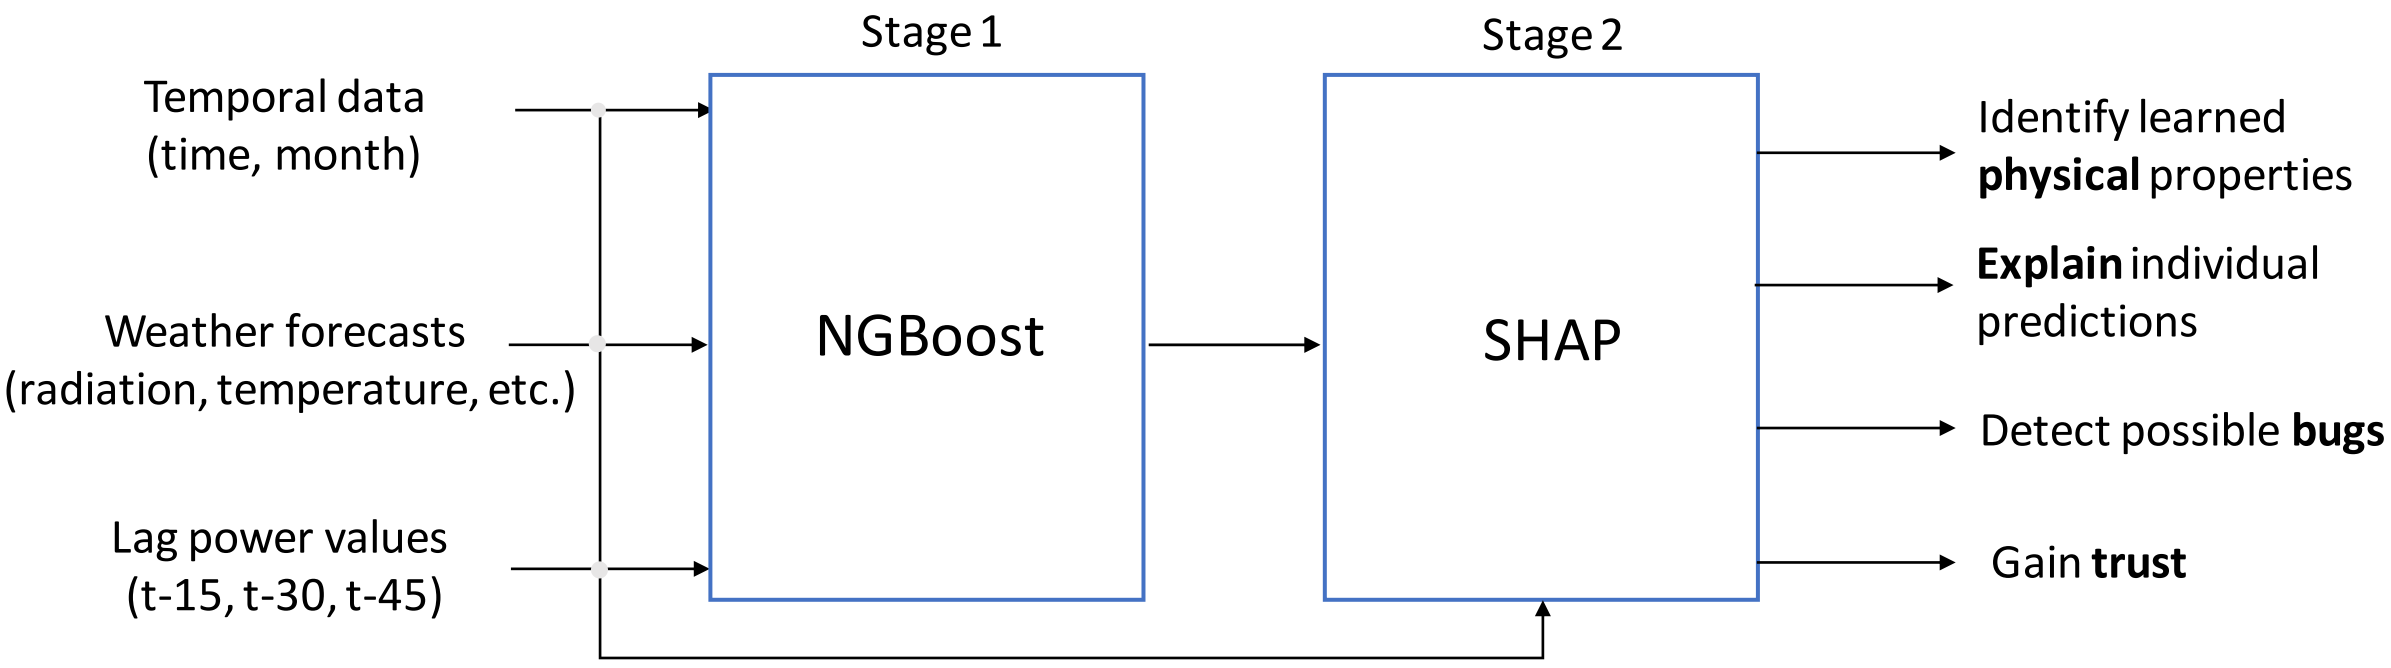

Natural Gradient Boosting (NGBoost)

NGBoost

Schematic representation

- NGBoost is a **gradient boosting algorithm** for solving **probabilistic** regression problems.
- Gradient boosting algorithms are based on the **sequential training** of several base learners.
- Each learner is optimized by minimizing the current **residual** as estimated by the ensemble of the previous learners.

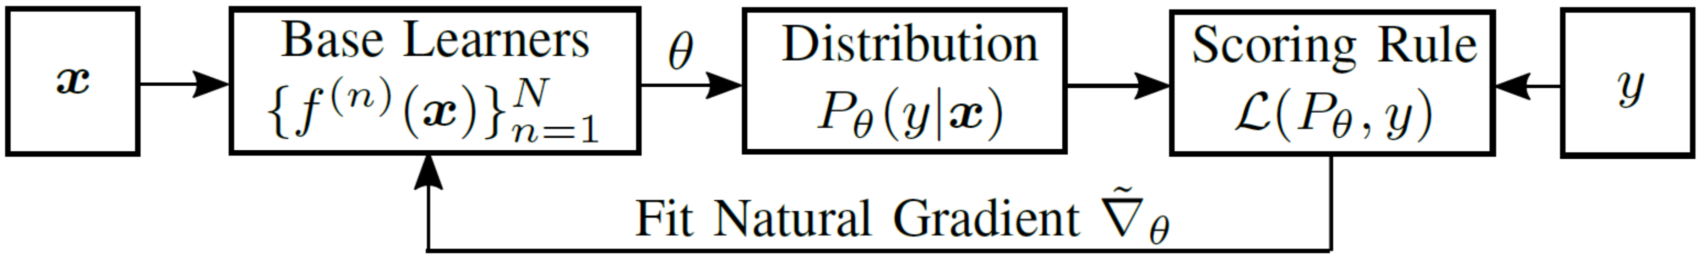

- ***Base learners***: Shallow decision trees have been proved to be an effective choice.
- ***Distribution***: Normal, Exponential, Laplace, …
- ***Scoring rule***: Logarithmic score, CRPS

SHAP

Model interpretability

SHAP

- **Model interpretability** is usually defined by a set of **feature attribution** values that quantify the influence of each input feature on the model output.

- The estimation of the feature attributions can be seen as a **cooperative game theory** problem, where each **feature** (player) contributes **differently** to the game.

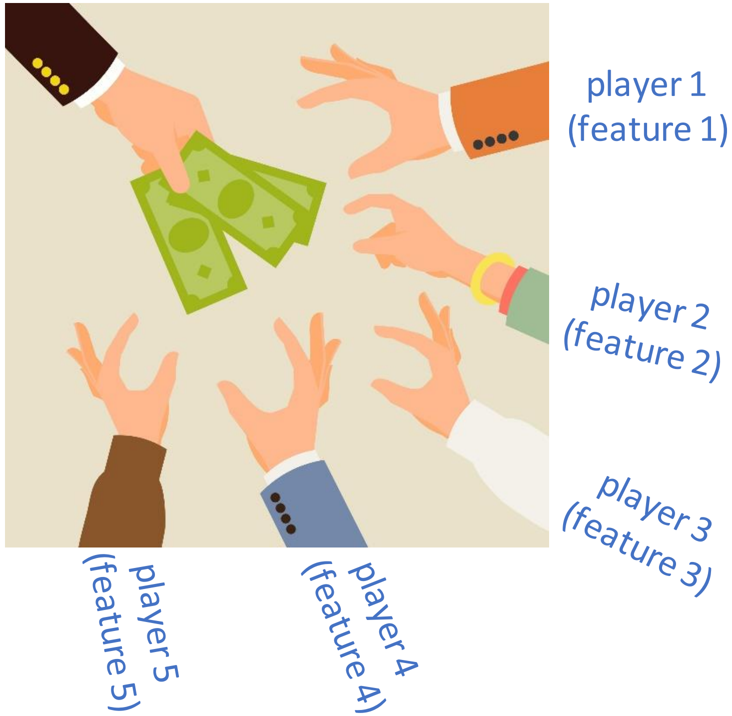

Fairness properties:
- Additivity (amounts sums up to the final reward)
- Consistency (more contribution → not less money)

SHAP

Model interpretability

SHAP

- The SHAP method estimates the influence of a feature by observing how the model behaves **with** and **without** that feature.

- Those **different combinations** result in a computationally challenging calculation:

$$\phi_d = \underbrace{\sum_{\mathcal{S} \subseteq \mathcal{Z} \setminus \{d\}}}_{\text{different combinations of features}} \frac{|\mathcal{S}|!(|\mathcal{Z}| - |\mathcal{S}| - 1)!}{|\mathcal{Z}|!} [\underbrace{f_{\boldsymbol{x}}(\mathcal{S} \cup \{d\})}_{\text{model output with feature } d} - \underbrace{f_{\boldsymbol{x}}(\mathcal{S})}_{\text{model output without feature } d}]$$

- **Tree based models** enjoy the exact calculation of SHAP values in **polynomial** instead of exponential time[1].

[1] Lundberg, Scott M., et al. "From local explanations to global understanding with explainable AI for trees." *Nature machine intelligence* 2.1 (2020): 56-67.

Power and Meteorogical Data

Power and Meteorological Data

- Time series data from two PV parks (**3.2MW** and **1.8MW**) located in Southern Germany (BW) are employed.

- **Time resolution**: 15 min
- **Forecasting horizon**: 36 hours (recursive multi-step predictions – 15min. intervals)
- **Training set**: 1 year, **Test set**: following month

- Data from 22:00 to 06:00 (**night hours**) are **discarded** from the recorded dataset.

- Input features:
  - ***Weather data (forecasts)***: temperature, relative humidity, precipitation, wind speed, ground level solar radiation
  - ***Temporal data***: month, time of the day (hour)
  - ***Lagged power values***: t-15, t-30, t-45

- We map the cyclical month variable onto a unit circle:

$$\text{month\_sin} = \sin(2\pi \cdot \text{month}/12)$$
$$\text{month\_cos} = \cos(2\pi \cdot \text{month}/12)$$

Results

Results

Comparative results

- We compare the proposed **NGBoost** with **lower upper bound estimation** (**LUBE**), **Gaussian process** (**GP**), and **persistence** (only for point forecasts).

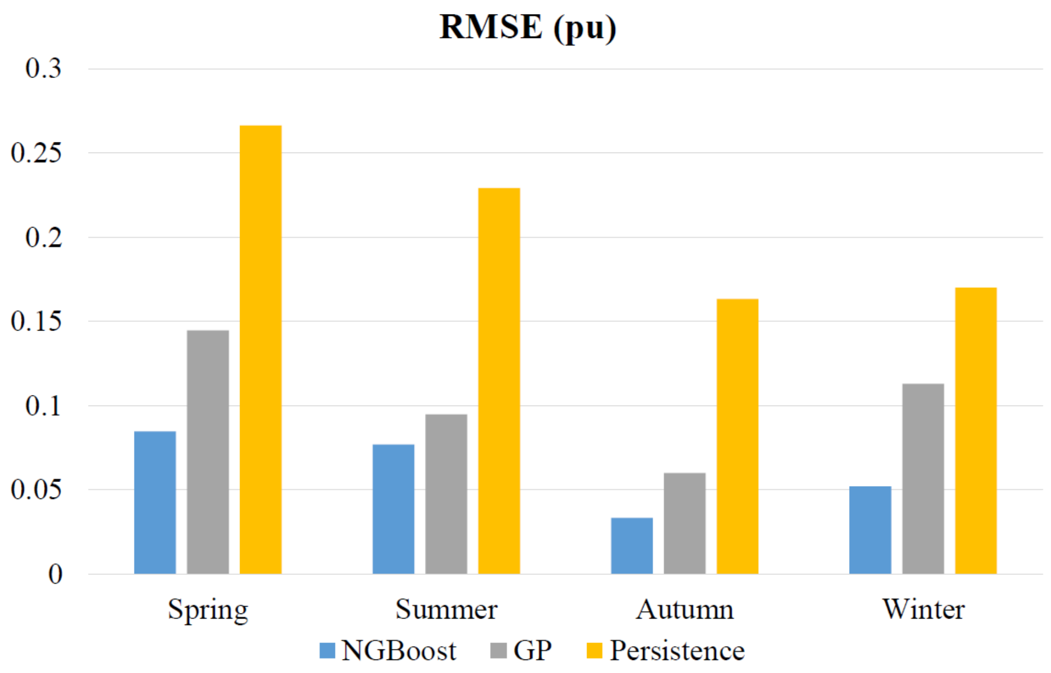

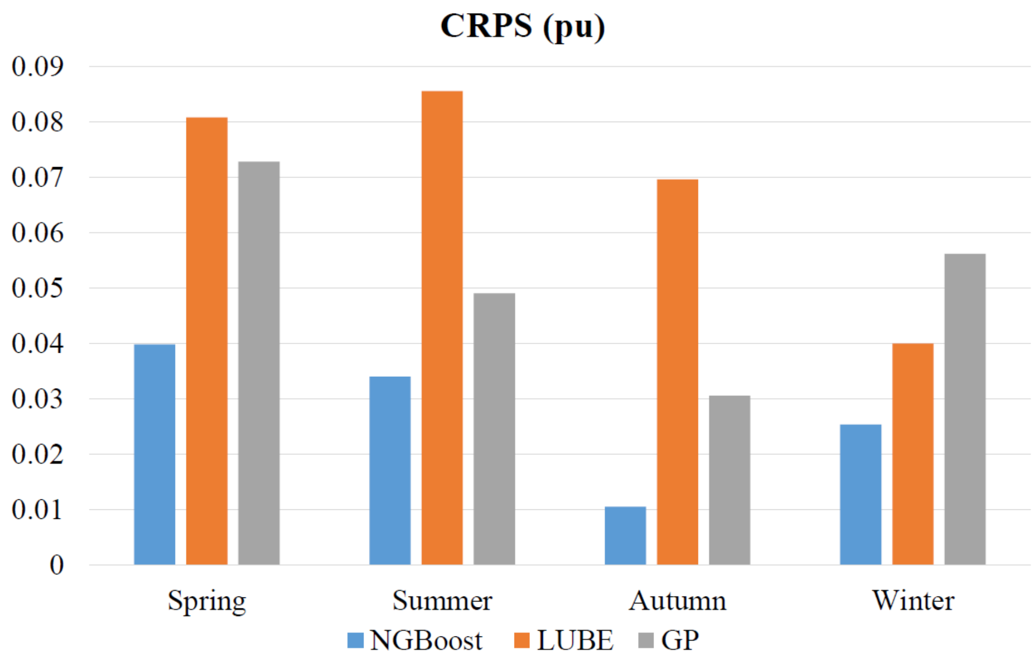

Results

Individual predictions

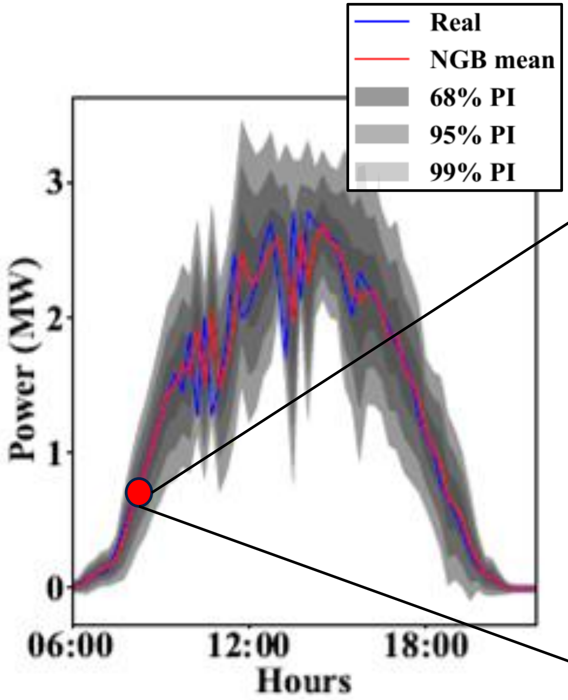

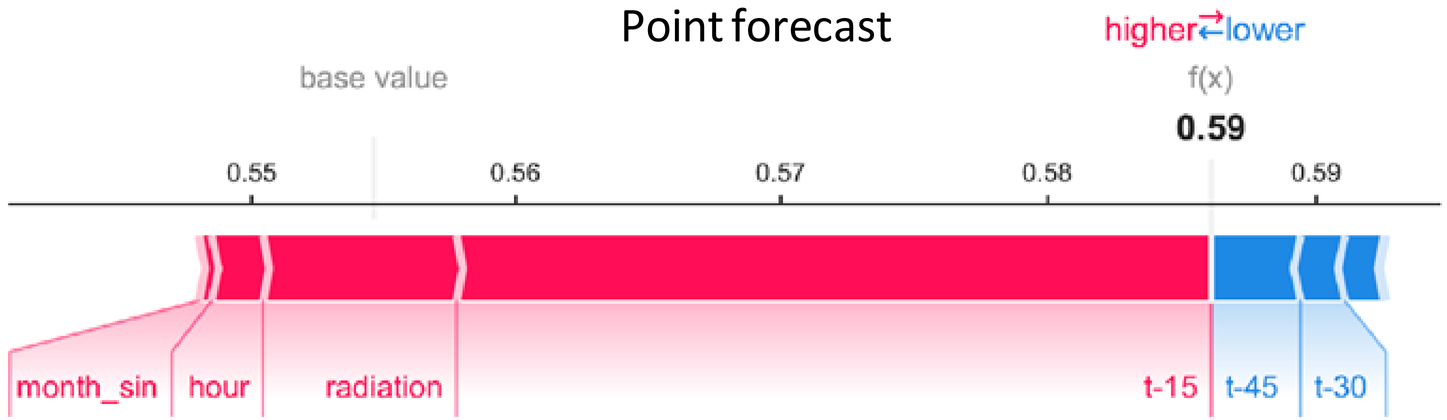

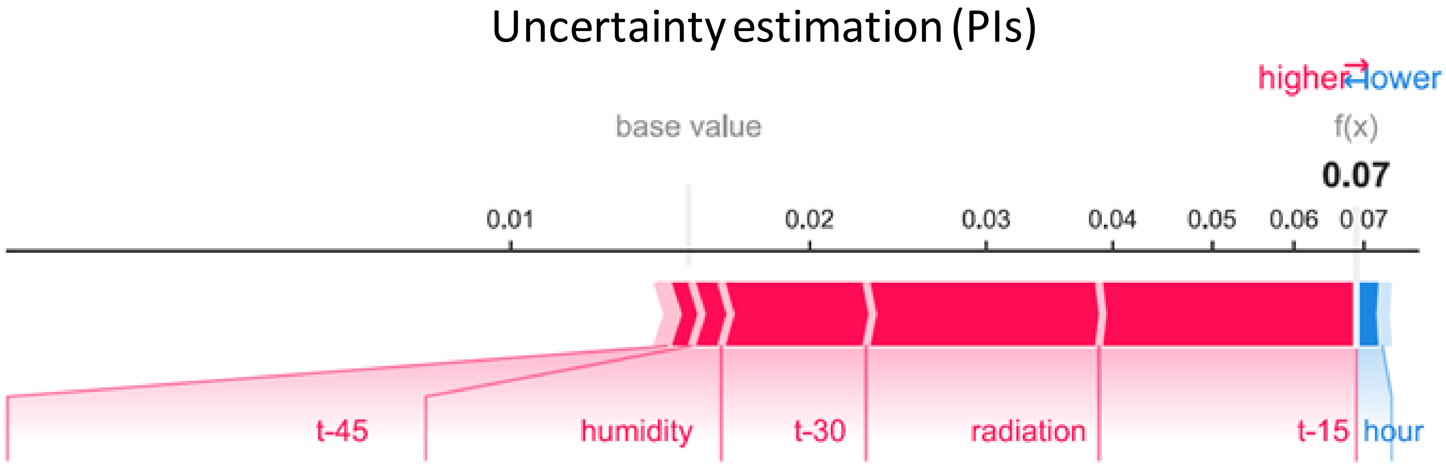

Results

Interpretation of point forecasts

- SHAP summary plots provide a **detailed insight** about the model.
- Each **dot** corresponds to **a training example**, and it is **colored** based on **how big is the value** of the respective feature.

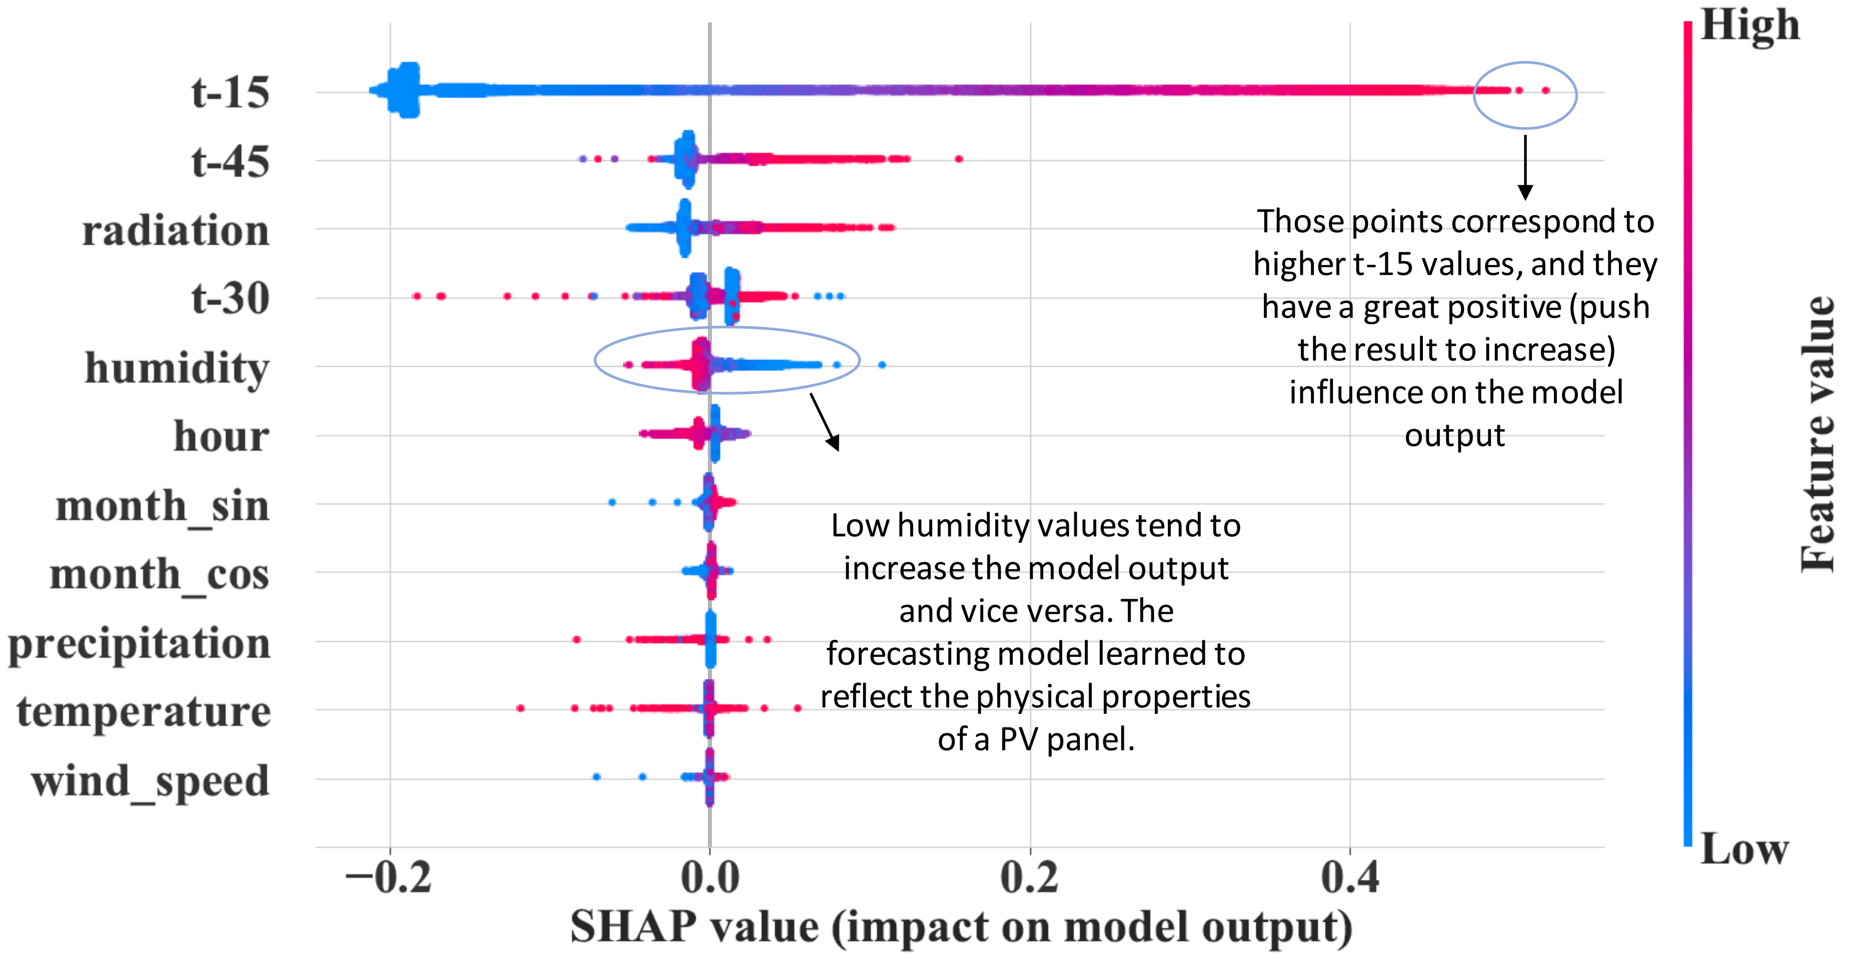

Those points correspond to higher t-15 values, and they have a great positive (push the result to increase) influence on the model output

Low humidity values tend to increase the model output and vice versa. The forecasting model learned to reflect the physical properties of a PV panel.

Results

Interaction plots

- SHAP interaction plots capture pairwise interactions between features.

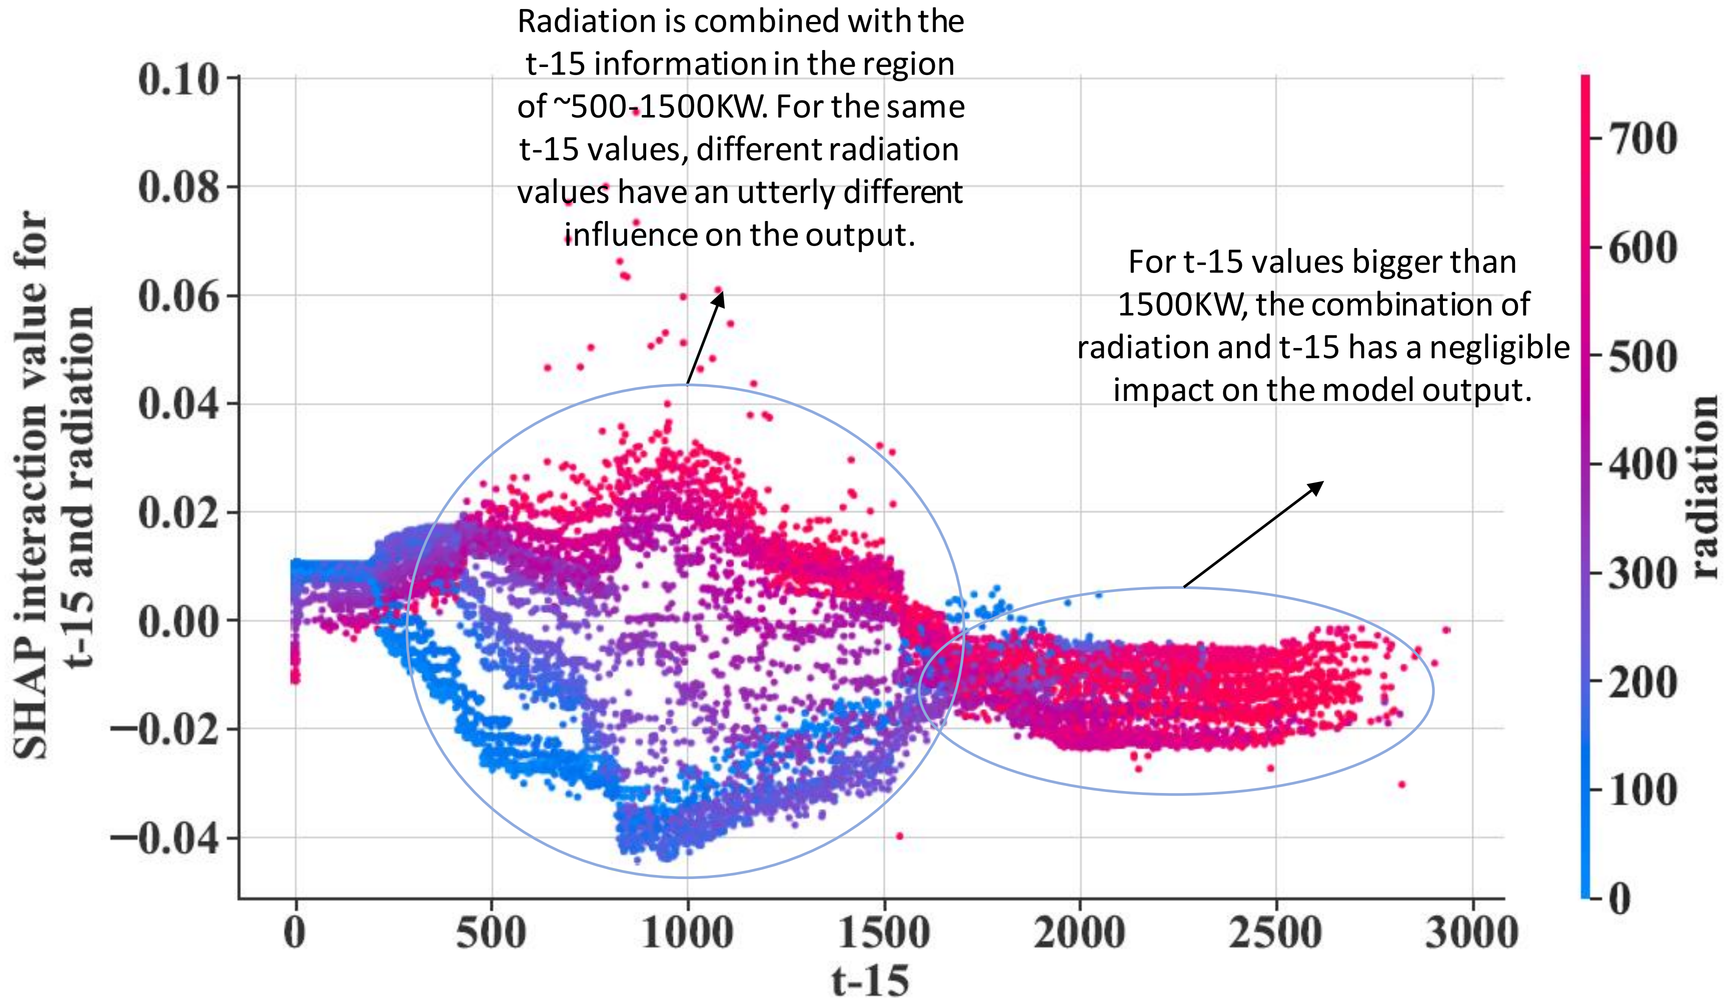

Results

Interpretation of uncertainty (PIs)

- SHAP summary plots provide a detailed insight about the model.
- Each dot corresponds to a training example, and it is colored based on how big is the value of the respective feature.

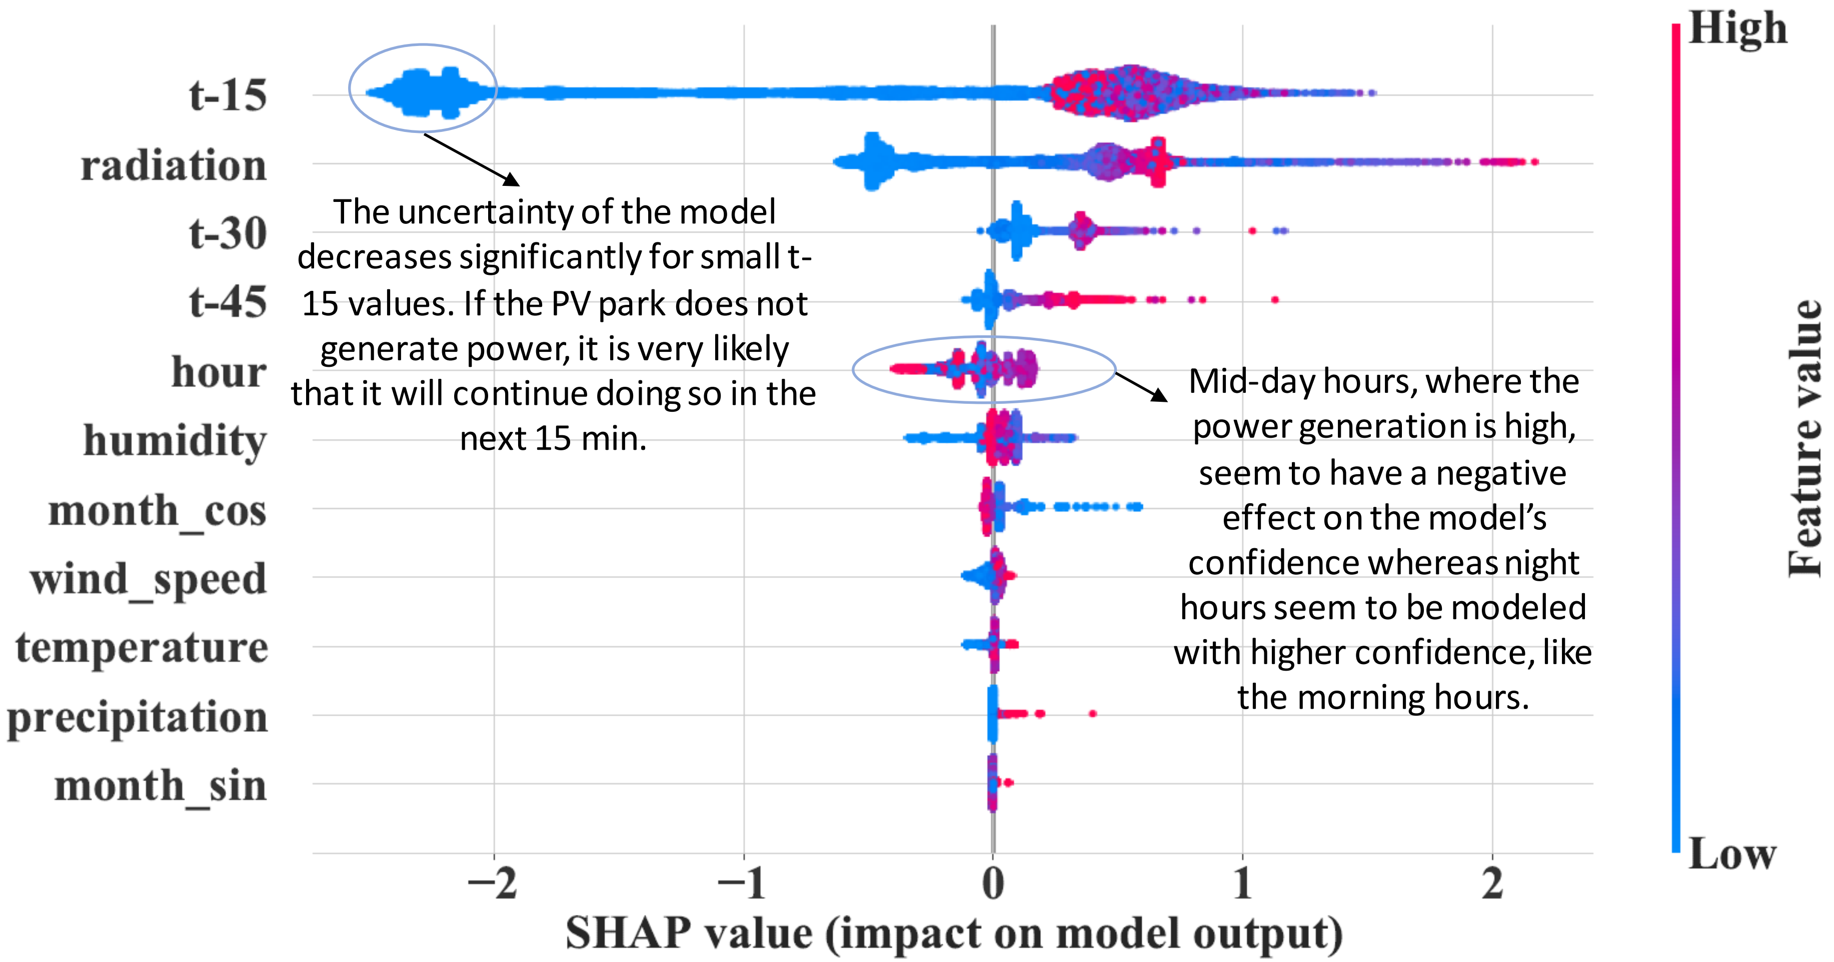

Results

Interaction plots

- SHAP interaction plots capture pairwise interactions between features.

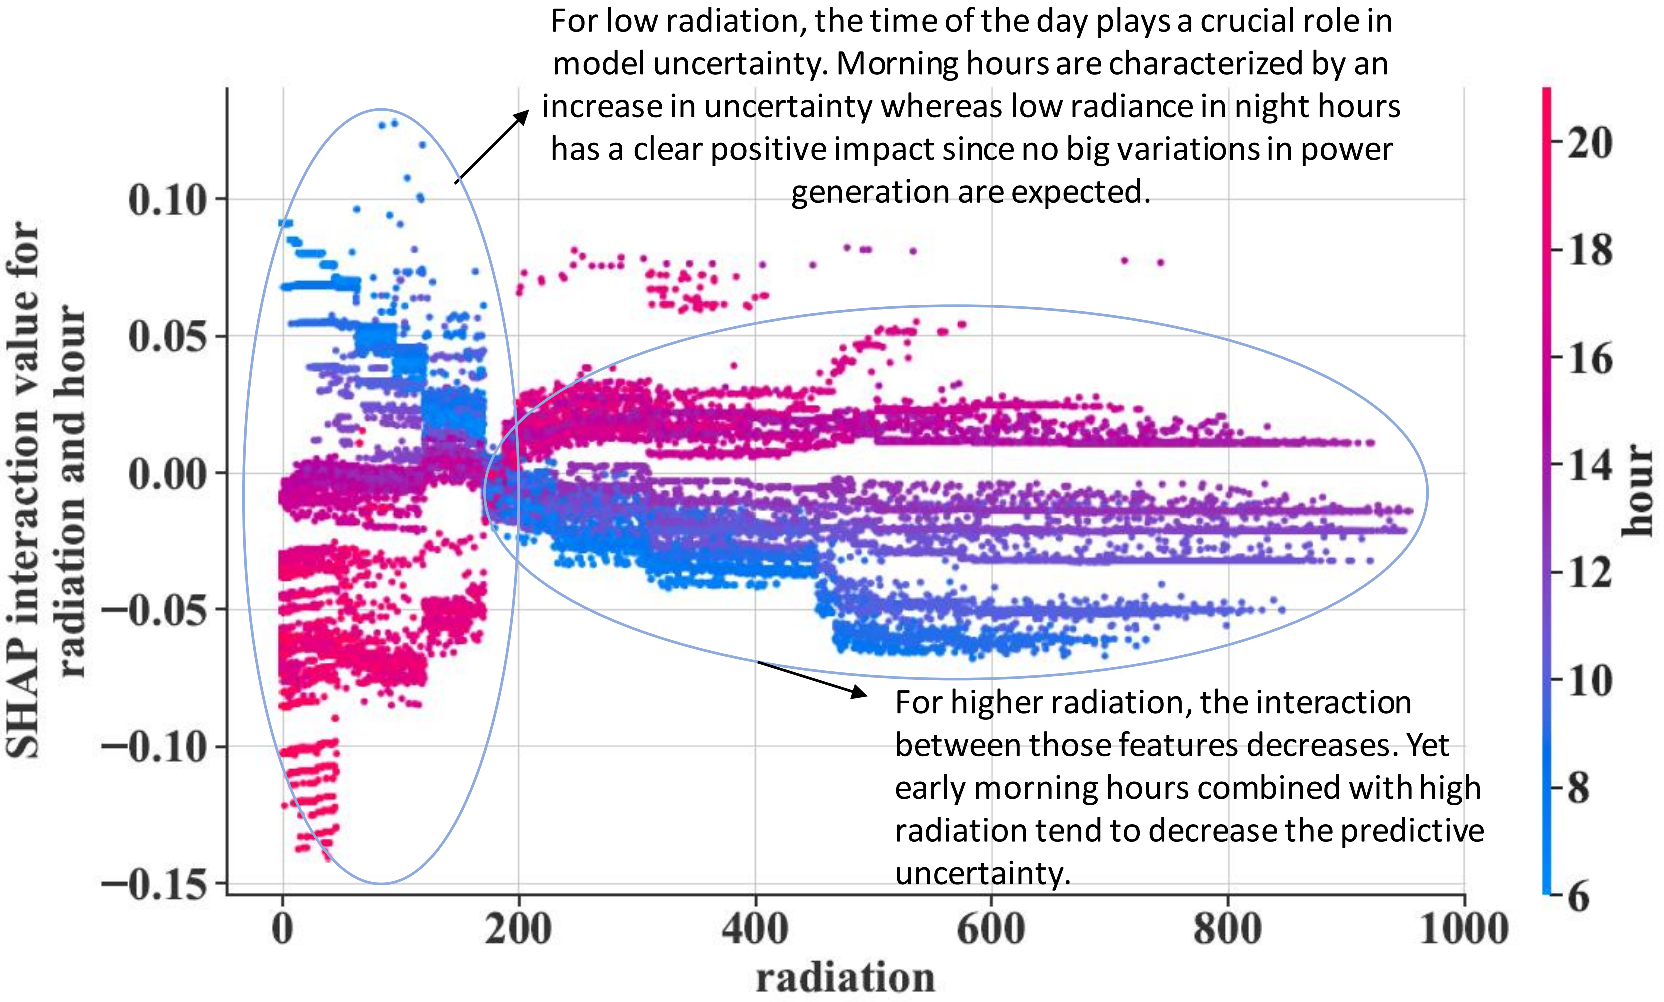

Results

Model interpretation

- A detailed analysis of the derived SHAP values revealed that the forecasting models learned some **nonlinear feature relationships** that follow **known physical properties** and **human logic and intuition!**

- This outcome may have a significant impact on **tackling the missing trust** in machine learning models and thus, help them become widespread.

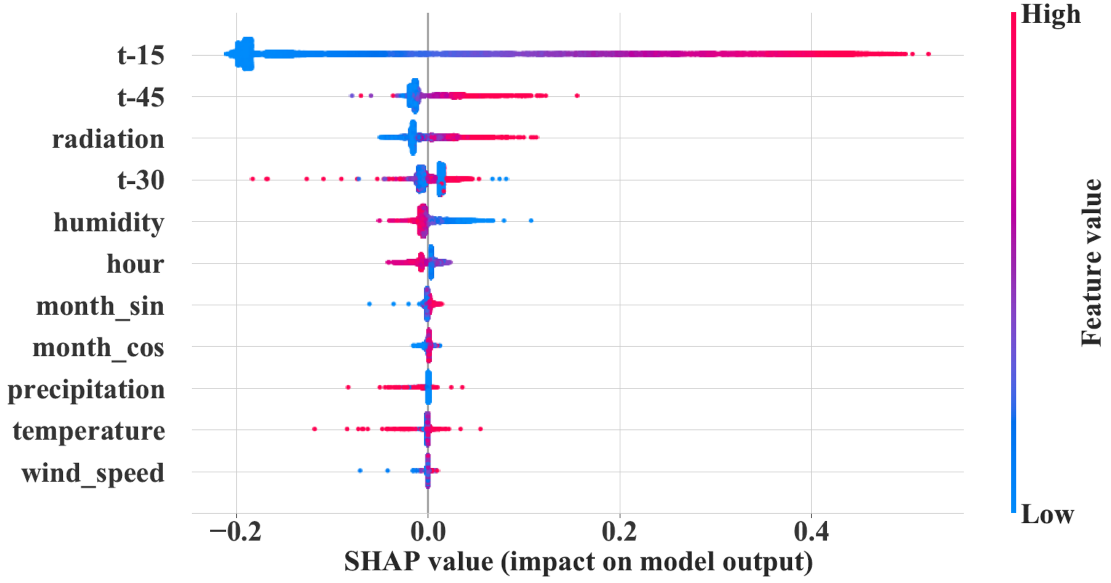

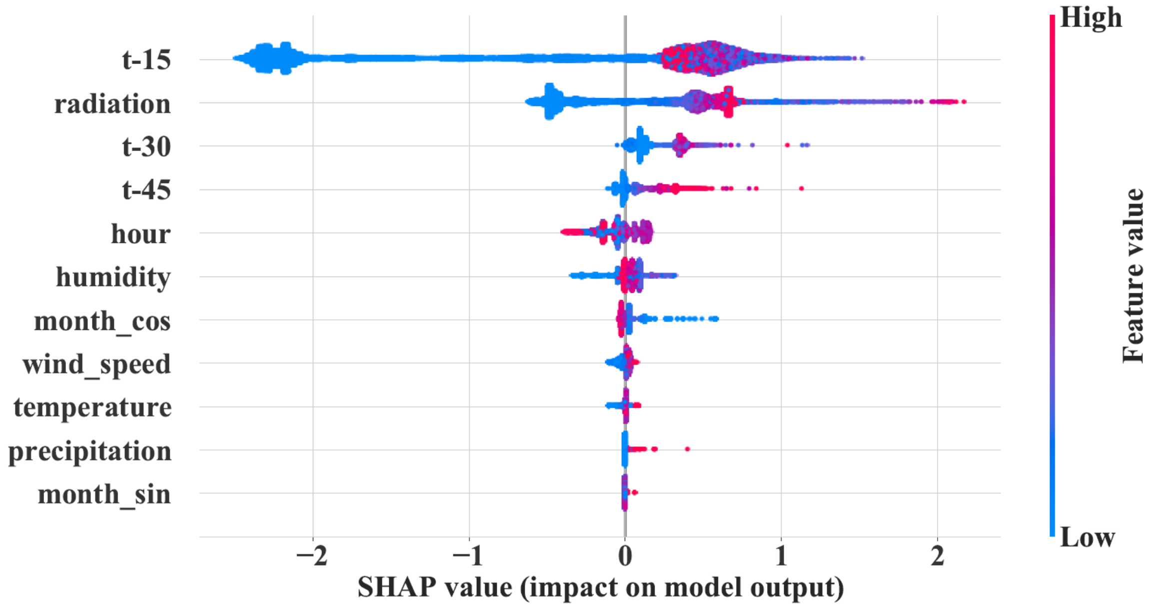

Results

Feature selection

- Based on model interpretation, **precipitation, temperature**, and **wind speed** are deployed only for a small number of observations by the point forecasting model while their contribution to estimating the uncertainty related to a prediction is rather negative.

- Precipitation, temperature, and wind speed were **discarded**, and the model was retrained.

- There was an **increase** in accuracy of around **6%** for RMSE and around **10%** for CRPS! This result may be caused by a local optimum solution due to the higher dimensionality of the first models.

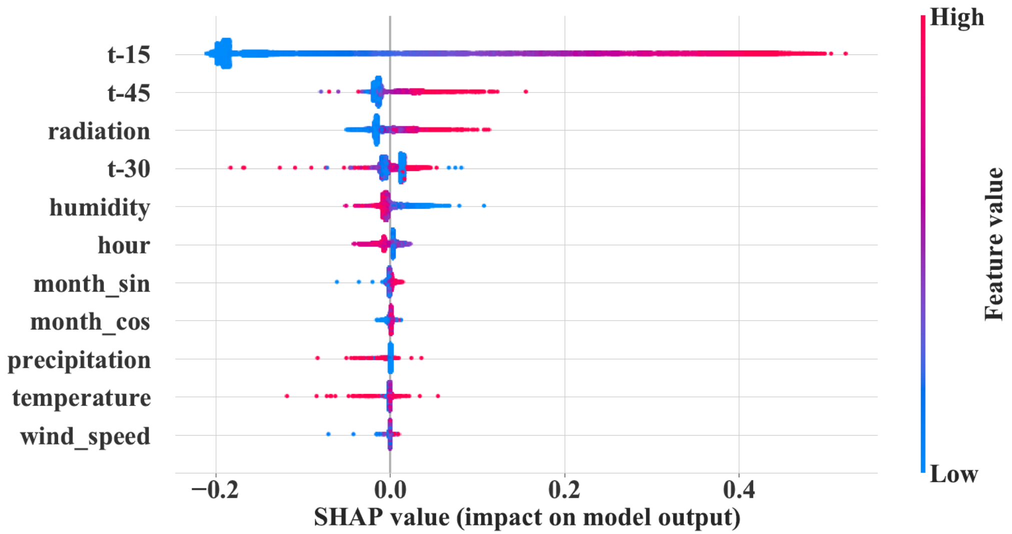

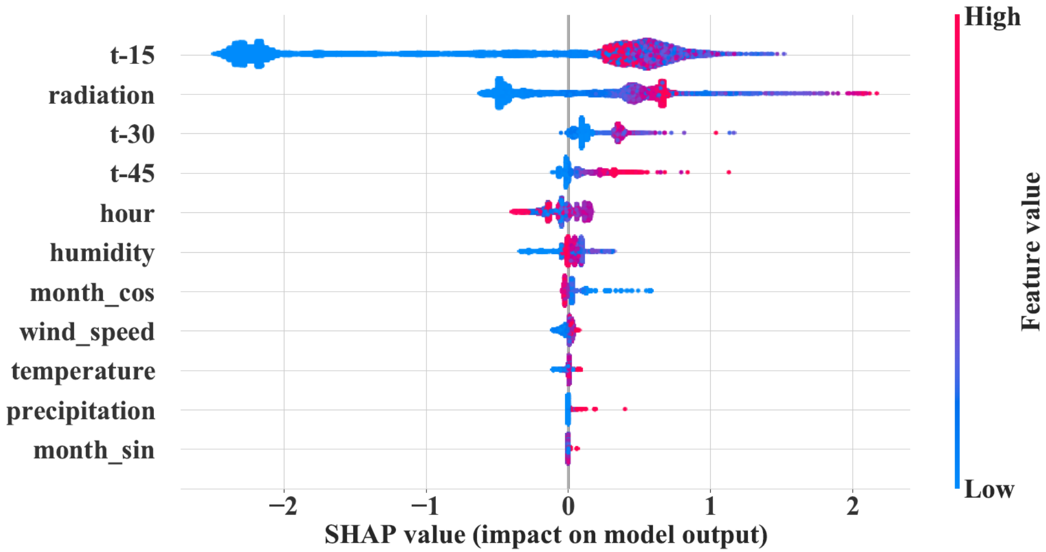

Conclusion

Conclusion

- We propose a probabilistic forecasting method that yields highly **accurate** and reliable forecasts while providing **full transparency** on its predictions.
- Based on a thorough comparison using both deterministic and probabilistic metrics, we showed that NGBoost can achieve **better performance** than GP and LUBE, regardless of the seasonal weather and power variations.
- A detailed analysis of the derived SHAP values revealed that the forecasting models came up with some nonlinear feature relationships that follow **known physical properties** and **human logic and intuition**.
- Apart from explaining individual predictions, SHAP values were employed for the **optimal feature selection**.
- No unusual or surprising relationship was developed by the proposed model. This finding is of utmost importance considering that **debugging machine learning** models is an extremely challenging task.

Thank you.

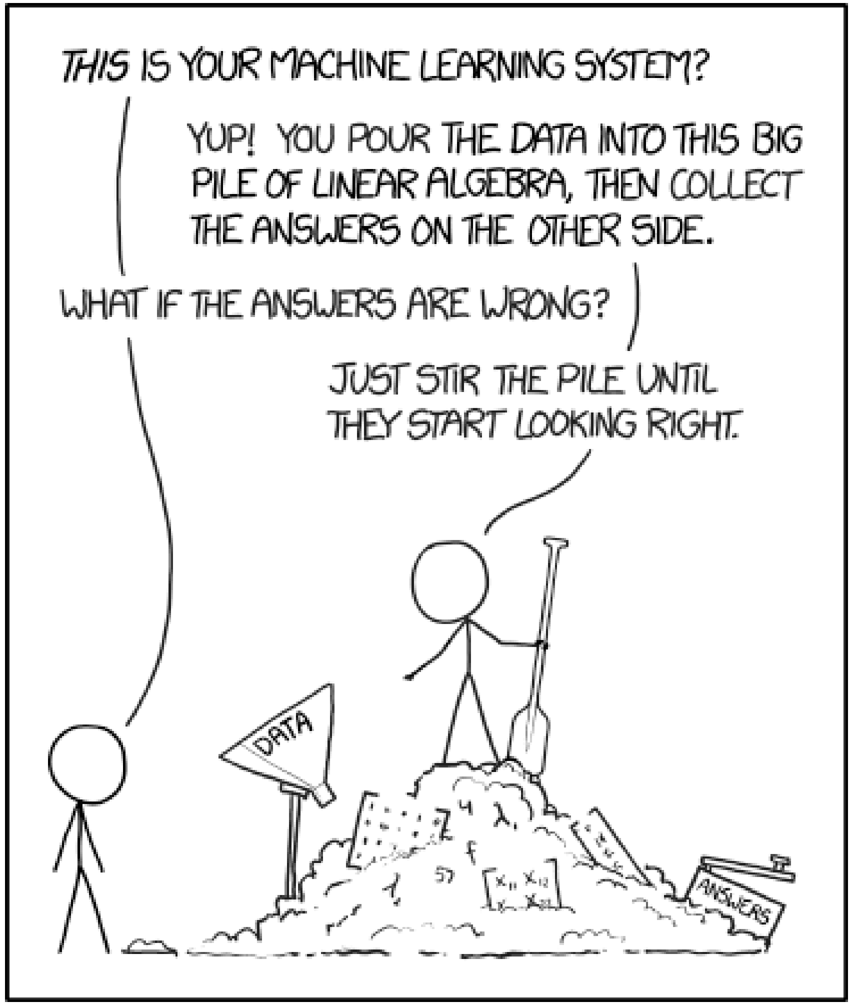

THIS IS YOUR MACHINE LEARNING SYSTEM?
YUP! YOU POUR THE DATA INTO THIS BIG PILE OF LINEAR ALGEBRA, THEN COLLECT THE ANSWERS ON THE OTHER SIDE.
WHAT IF THE ANSWERS ARE WRONG?
JUST STIR THE PILE UNTIL THEY START LOOKING RIGHT.
DATA
ANSWERS

Introduction

Current approaches

Forecasting

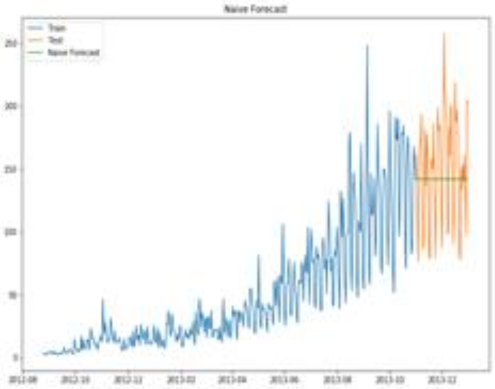

Deterministic

Has been widely developed.

- Yields point forecasts at each time step.
- No information about predictive uncertainty.

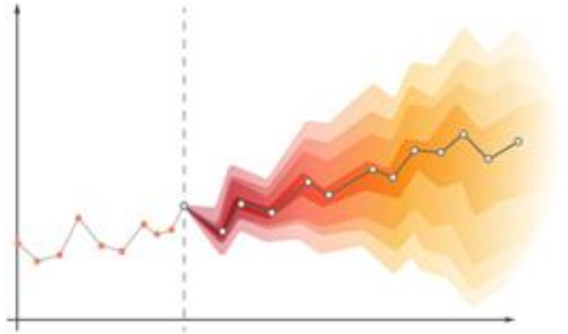

Probabilistic

Has not been studied to a great extend.

- Yields probability distributions at each time steps
- Uncertainty information about future outcomes.

Results

Probabilistic forecasting

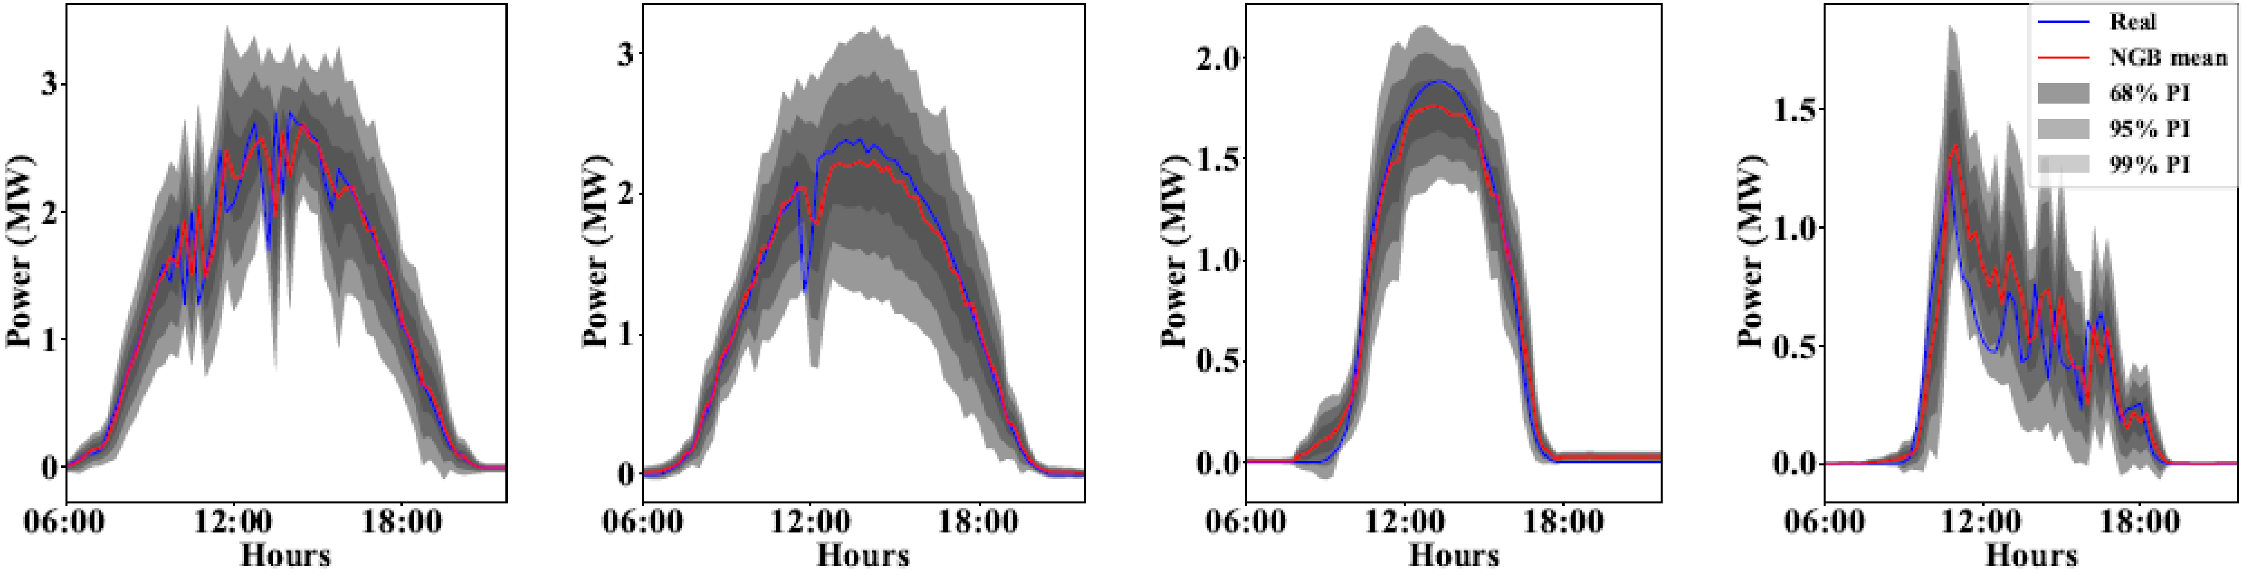

(a) Spring day. (b) Summer day. (c) Autumn day. (d) Winter day.

PVP1: Day-ahead PV power forecasts using NGBoost.

- The NGBoost algorithm was able to yield highly **accurate** and **sharp** probabilistic forecasts.

Results

Interaction plots

- SHAP interaction plots capture pairwise interactions between features.