

University of Stuttgart

Institute of Combustion and
Power Plant Technology



Pfaffenwaldring 23 • 70569 Stuttgart • Germany • E-mail: georgios.mitrentsis@ifk.uni-stuttgart.de

Stuttgart, 27.09.2021

Enabling Energy Transition: An Interpretable AI Model for Probabilistic Solar Power Forecasting

Dipl.-Ing. Georgios Mitrentsis

Student's address / contact details:

Georgios Mitrentsis
Institute of Combustion and Power Plant Technology – IFK
Department Power Generation and Automatic Control

Universität Stuttgart
Pfaffenwaldring 23
D-70569 Stuttgart

E-Mail: georgios.mitrentsis@ifk.uni-stuttgart.de
Tel.: +49 176 87851752

Enabling Energy Transition: An Interpretable AI Model for Probabilistic Solar Power Forecasting

1 Introduction

The ongoing transition of the power system towards a more fossil-fuel-free system has led to a wide integration of renewable energy sources (RES) worldwide. Photovoltaic (PV) power forms one of the most widespread and promising alternatives to conventional power plants [1]. However, the stochastic nature of PV power induced by volatile weather conditions hinders the reliable electricity supply which can result in an increase in the reserve capacity of the system. At the same time, PV parks are typically connected to the grid through power electronics leading to a decline of system inertia and thus, pushing the system closer to its stability margins [2]. Against this background, accurate and reliable PV power forecasting can alleviate the aforementioned challenges allowing for a cost-effective and large-scale PV integration.

PV power forecasting can be split into two general categories, namely deterministic and probabilistic approaches. The former yield single point forecasts at each time step whereas in the latter, the PV output is considered as a random variable and the model generates the respective probability density functions or prediction intervals (PIs). While deterministic forecasting has been extensively studied in the literature, probabilistic approaches have only recently gained attention [3]. As expected, probabilistic forecasts are more valuable compared to deterministic ones since they provide uncertainty information about the future power outcomes [3]. This information can be leveraged by the involved parties, e.g., TSOs, DSOs, traders, legislators, for a lower risk and a more beneficial decision making [4]. Indicative examples can be found in [5] and [6], which exploit probabilistic load and wind power forecasts in order to estimate the optimal reserve requirements.

Yet, despite the inspiring research work in probabilistic RES forecasting, it has still not been studied to a great extent. At the same time, we also lack a systematic approach to integrate it into the system operation [3, 7]. Furthermore, apart probably from the bootstrapping techniques, the rest of the approaches can be considered as complex black box models that require high computational resources for training. Those characteristics may be crucial for the widespread integration of those algorithms in practice, especially in the case of TSOs and DSOs that are responsible for the safe and reliable operation of the power system, since the more complex become those models, the harder it gets to understand them and explain their results. Therefore, in applications where safety critical decisions are to be made, e.g., power system operation and control, the trust in complex black box models may be questioned. At the same time, due to the role of TSOs and DSOs in society, one of their main duties is to be transparent and share information about the forecasting quantities with the general public. Hence, their trust in those models is of utmost importance. It is also worth mentioning that the forecasting models are broadly employed by energy traders. Therefore, identifying learned patterns or rules, detecting possible bias in the predictions, and exploring regions where the used features are not sufficient for an accurate prediction may be crucial for the respective bidding strategies and consequently, for the maximization of their profit. Moreover, the ability to interpret the generated forecasts has been characterized as “*vital*” in the very recent work of [1].

To address those issues, a two stage interpretable probabilistic framework is proposed, which can generate highly accurate, reliable, and sharp probabilistic forecasts while providing full transparency on its predictions. In the first stage, we propose the application of the natural gradient boosting (NGBoost) algorithm for yielding probabilistic PV power forecasts. Due to its ability to deploy decision trees as base learners, it is combined with Shapley additive explanation (SHAP) values in the second stage. The estimation of SHAP values is the unique consistent attribution method that is able to provide theoretical optimal explanations about the predictions of a model [8].

2 Interpretable Forecasting model

As mentioned, the NGBoost algorithm is deployed in the first stage of the proposed framework. NGBoost is a gradient boosting algorithm for solving probabilistic regression problems [9]. In general, gradient boosting algorithms are based on the sequential training of several base learners, which all together form an additive ensemble. Each learner is optimized by minimizing the current residual as estimated by the ensemble of the previous learners. Then, the output of that learner is scaled by a learning rate and it is appended to the current ensemble.

To benchmark the performance of NGBoost in PV power forecasting, its output was compared with the one generated by two state-of-the-art probabilistic approaches, the Gaussian process (GP) and the lower upper bound estimation (LUBE), respectively. To do so, a set of validation metrics assessing the accuracy, bias, reliability, sharpness, and overall probabilistic performance of the forecasting models has been deployed.

Nevertheless, probabilistic forecasting models are predominantly complex black box models that lack transparency. In the second stage of the proposed framework, the goal is to understand the model and its learned

feature relationships. It is worth mentioning that the term interpretability refers to the ability of the model to communicate why and how a prediction was made using feature attribution values, which basically describe the contribution of each feature to the final output. To interpret the predictions, the SHAP method is exploited in the second stage. It can estimate the influence of a feature by observing how the model behaves with and without that feature. To do so, it casts the problem of feature attribution to a cooperative game theory problem, where each player (feature) contributes differently to the game. As proved in [10], the SHAP method is the unique approach able to deliver local explanations while providing theoretical guarantees about the method’s consistency based on game theory [8].

3 Field and Weather Data

The power data employed in this study were acquired from two PV parks located in Southern Germany, and in particular in the federal state of Baden-Württemberg. Those two PV parks have a nominal power of around 3.2 MVA and 1.8 MVA, respectively, and they are located in two different regions of this federal state. We will refer to them as *PVP1* and *PVP2*, respectively. The recorded data corresponding to quarterly hour values were obtained from February 2018 to October 2019. As for the deployed features, we use a set populated by the most commonly used meteorological variables, time variables, and past power values as input for the model. These include temperature, relative humidity, precipitation, wind speed, ground level solar radiation, month, time of the day, and lagged power values [1]. Based on the autocorrelation values of our dataset, we deploy the three past power values corresponding to the power generation 15 min, 30 min, and 45 min before the desired time step. We will refer to them as $t-15$, $t-30$, $t-45$, respectively. Note that in order to ensure that December and January are seen as two consecutive months by the model and not as two completely different feature values, i.e., $\text{month} = 1$ or $\text{month} = 12$, we map the cyclical month variable onto a unit circle, where for each month we compute its sin and cos. Furthermore, the meteorological data were obtained from the German Weather Agency (Deutscher Wetterdienst - DWD). Those data were acquired from weather stations located a few kilometers away from the respective PV parks and thus, they do not match the exact weather values at the desired locations. In this regard, the deployed weather data are assumed as the corresponding weather forecasts at those two locations.

In this work, we are interested in short-term forecasting for a time horizon of 36 h (day-ahead) and a time resolution of 15 min. To do so, we perform recursive multi-step predictions based on the predicted values of the previous time steps. In addition, the training set comprises the data from one whole year while the test set comprises the data from the respective following month. Four different pairs of training-test sets are constructed in total, where each test set corresponds to a different season, in order to assess the performance of the algorithm under different weather and generation scenarios.

4 Results

4.1 Probabilistic forecasting: Comparative results

In this subsection, we thoroughly examine the performance of NGBoost in probabilistic PV forecasting while taking into account the influence of the seasons. To do so, we compare the proposed NGBoost with LUBE (only for PIs since the algorithm does not yield point forecasts), GP, and persistence (predicted power = power generated the previous day at the exact same time). Different models of the same approach are developed based on different sets of hyperparameters. Specifically, we vary the depth of the trees in NGBoost, the number of neurons (nr) and the penalty factor η in LUBE, and the kernel function in GP. The comparative plots of root mean square error (RMSE) and the continuous ranked probability score (CRPS), which is a widely used metric for probabilistic forecasts, are illustrated in Fig. 1.

As observed in the RMSE plot, NGBoost outperforms GP and persistence with respect to the point forecasts, regardless of the time of the year. Interesting are the results of summer and winter, which are characterized by utterly different weather conditions. During summer, the PV output is usually a clear bell shaped curve and thus, both NGBoost and GP can generate accurate point forecasts. On the contrary, winter shows sudden power deviations that make forecasting more challenging. Regarding the probabilistic forecasts, the overall performance of NGBoost in probabilistic forecasting is remarkably better than that of GP and LUBE, in particular in autumn and winter, which are characterized by more volatile weather conditions and intermittent generation. It is also worth emphasizing that the NGBoost models can be trained within less than 2 minutes due to its gradient boosting architecture whereas the other two algorithms require around 25 minutes each. The fast training time is of significant importance for the development of a machine learning model in practice, as initially, different features and hyperparameters should be tried out in order to find the optimal combination of them.

Fig. 2 shows four indicative results as generated by the NGBoost algorithm using data from PVP1. As depicted in Fig. 2, each plot illustrates a completely different generation pattern due to seasonal variations. Yet the NGBoost

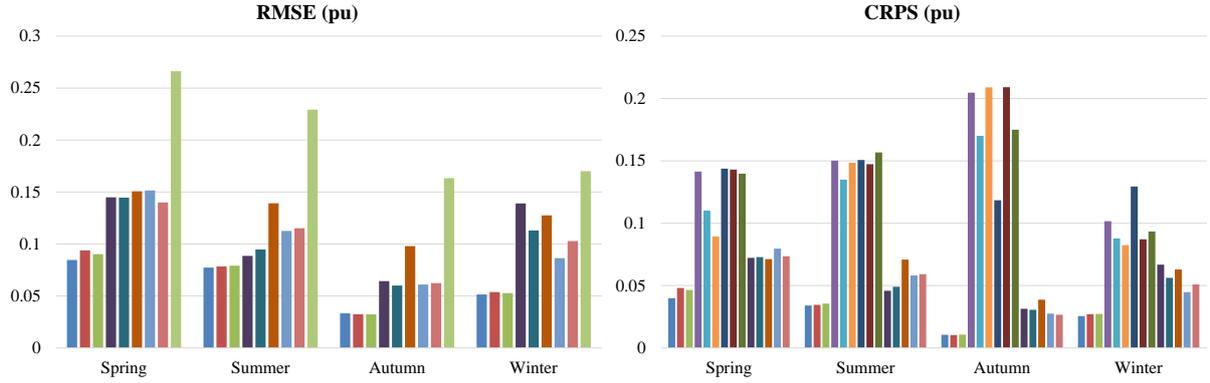
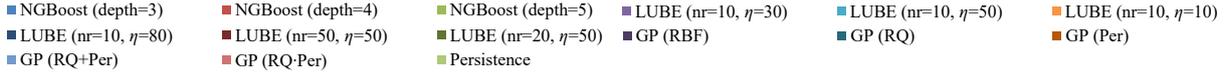


Figure 1: Average RMSE and CRPS of the validation set in pu.

algorithm was able to yield highly accurate and sharp probabilistic forecasts, where the point forecasts lie always within the boundaries of the predicted distributions.

4.2 Interpretation of the model

In this subsection, we study the results of the second stage of the proposed forecasting framework which aims at interpreting and understanding the output of the probabilistic forecasting model. To do so, we deploy the SHAP summary plot that summarizes the local explanations in a compact yet information-rich way, as shown in Fig. 3. In those plots, each dot corresponds to an individual prediction while its position along the x-axis corresponds to the impact of the respective feature on the model output. Additionally, the color of the dot denotes the feature value in order to highlight to what extent different feature values affect the result. For instance, high t-15 values tend to push the model output to higher values whereas close to zero t-15 values push the model output to decrease (Fig. 3a). Hence, the higher the t-15 values are, the bigger will be their impact on the model output. The vertical dispersion of the dots implies a higher number of observations with a similar impact. In this context, those SHAP summary plots reveal information about both the magnitude and the direction of each feature's impact as well as the number of observations with those characteristics. Note that all models, i.e., both for PVP1 and PVP2, yielded almost identical SHAP plots. Therefore, in order to avoid redundancy, we present the respective SHAP plots of a randomly selected model from PVP1.

Starting with the point forecasts (Fig. 3a), a few indicative observations are briefly presented in this paragraph. High radiation and t-45 values seem to have a bigger impact on the model output than their corresponding lower values. This observation may be attributed to the fact that for low radiation and lagged power values, the information about t-15 seems to suffice for the model to yield a prediction. On the contrary, for higher radiation and lagged power values, where the PV generates power, the model requires to employ more features in order to make an ac-

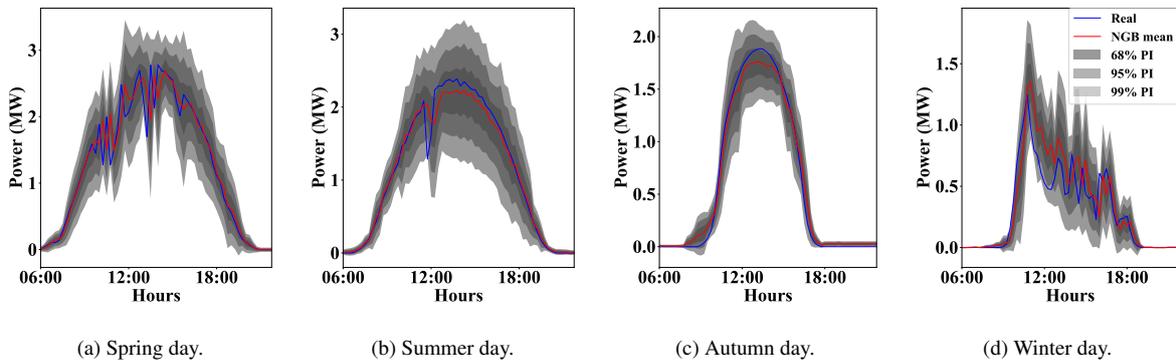


Figure 2: PVP1: Day-ahead PV power forecasts using NGBost.

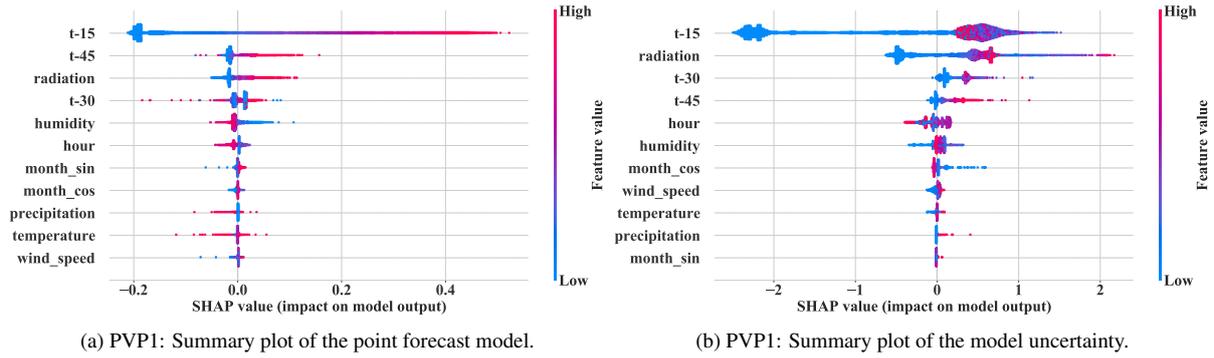


Figure 3: SHAP summary plots of PVP1.

curate prediction. Moreover, bigger humidity values lead to smaller power predictions and vice versa, following the physical properties of a PV cell, whose efficiency drops in high humidity. Similar common sense explanations apply to the time of the day feature (hour), where mid-day hours have a positive impact on the model output since the bell-shape curve reaches its peak during that time period, late hours push the result towards smaller values, and early morning hours towards positive.

As for how the model yields the PIs associated with the predictions (Fig. 3b), the uncertainty of the model decreases significantly for small t-15 values. For instance, if the PV park does not generate power, it is very likely that it will continue doing so in the next 15 min. For higher power values, t-15 has a comparable impact for most of the observations indicating the stochastic nature of PV generation. Consequently, higher radiation usually means lower confidence about the model prediction. This behavior can be also observed in the t-30 and t-45 features. Furthermore, mid-day hours, where the power generation is high, seem to have a negative effect on the model's confidence whereas night hours seem to be modeled with higher confidence, like the morning hours. This can be readily explained by the higher variance of power values during mid-day hours throughout a year. In addition, low month_cos values correspond to spring and summer, where the generation is high and thus, it is subject to many variations. During those two seasons, the model is less confident whereas during autumn and winter, where generation is not expected to be high, the model is more confident.

In addition to SHAP summary plots, two indicative examples of the derived SHAP interaction plots are illustrated in Fig. 4. The SHAP interaction plots express the pairwise interactions between features, as they learned by the model. In our case, high interaction values have been developed between pairs of features with high-high importance and high-low importance whereas low ranked features do not show remarkable interactions with each other. In this context, a representative SHAP interaction plot between t-15 and radiation is depicted in Fig. 4a for the case of point forecast. As derived by this figure, t-15 values within the range of 0.5 MW-1.5 MW can have an utterly different influence on the model predictions depending on the radiation values. As expected, high radiation leads to higher and low radiation to lower power predictions. For higher lagged power values (> 1.5MW), the interaction between t-15 and radiation tends to decrease. This can be attributed to the fact that the information about radiation is redundant when the PV park generates high power. As for the uncertainty estimation, an in-

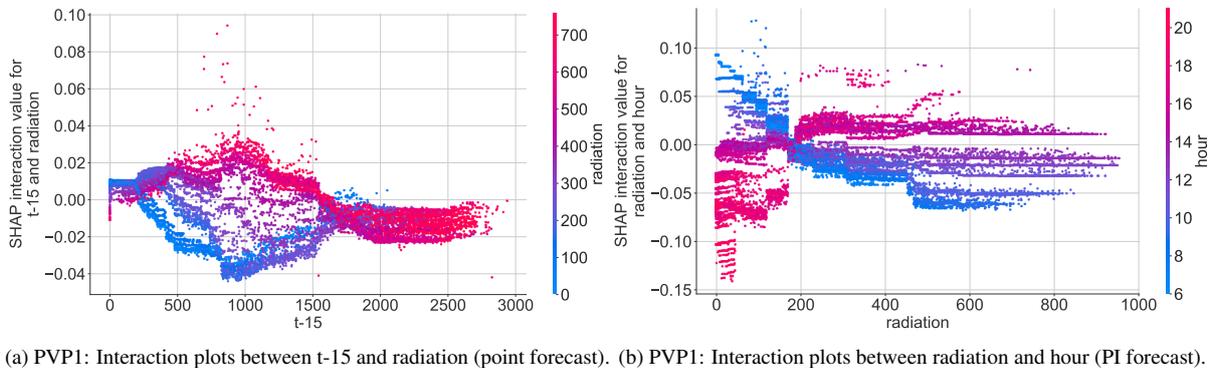


Figure 4: Two indicative SHAP interaction plots of PVP1.

dicative interaction plot between the time of the day (hour) and radiation is shown in Fig. 4b, where the nonlinear interactions between those two features are illustrated. For low radiation, the time of the day plays a crucial role in model’s uncertainty. Morning hours are characterized by an increase in uncertainty whereas low radiance in night hours has a clear positive impact since no big variations in power generation are expected compared to the morning hours. For radiation values above 200 W/m^2 , the interaction between those features decreases. Yet early morning hours combined with high radiation tend to decrease the predictive uncertainty.

Overall, it was demonstrated that the model can successfully encode the stochastic nature of PV generation while learning properties that follow physical laws, human intuition, and common sense. Finally, based on the model interpretation, it was identified that the weather information of 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. Thus, those features were discarded from our datasets and the models were retrained. As a result, there was an increase in accuracy of around 6% for RMSE and around 10% for CRPS. Therefore, following the proposed methodology, one can fully understand how a trained model works, determine the features leading to an increase in uncertainty, and possibly discard them. It is worth pointing out that there may be specific examples where those features are useful for the model and lead to lower prediction errors. In the context of the overall performance however, they may have a rather negative impact.

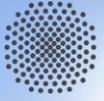
5 Conclusions

The optimal PV integration is a challenging problem due to volatile weather conditions leading to intermittent generation. The latter has a direct impact on system stability and reliable electricity supply and thus, accurate and reliable PV forecasting techniques are of utmost importance. In this regard, a small number of probabilistic models have been recently introduced in the literature aiming at providing more information about the derived forecasts compared to deterministic methods. However, most of the existing probabilistic approaches are complex black box models characterized by high training requirements and extensive hyperparameter tuning. Furthermore, the lack of interpretability may prevent their practical implementation in a safety critical system as the power system. To this end, we propose a probabilistic forecasting method that can be readily implemented without demanding expert knowledge, does not require extensive training times and hyperparameter tuning, yields highly accurate and reliable forecasts, and finally, provides full transparency about its predictions.

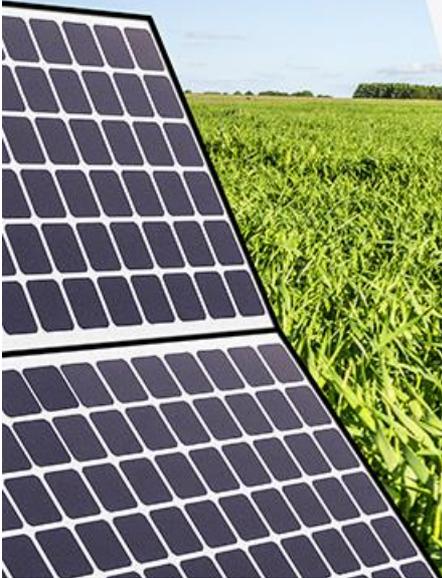
Importantly, it was demonstrated that feature importance is particularly advantageous with respect to understanding the output of a model. By using the theoretically optimal SHAP values as feature attributions, we are able to interpret the model, identify the learned feature relationships, explain individual predictions, reduce the computational requirements for the training without jeopardizing the model accuracy, detect possible bugs, and gain trust in the model. Finally, a detailed analysis of the derived SHAP values revealed that the proposed forecasting models came up with some nonlinear feature relationships that follow physical properties as well as human logic and intuition. This outcome can have a significant impact on tackling the missing trust in machine learning models. Consequently, their full potential in challenging power system problems can be optimally exploited.

References

- [1] R. Ahmed, V. Sreeram, Y. Mishra, and M. Arif, “A review and evaluation of the state-of-the-art in PV solar power forecasting: Techniques and optimization,” *Renewable and Sustainable Energy Reviews*, vol. 124, p. 109792, 2020.
- [2] N. Hatzigargyriou, J. Milanović, C. Rahmann, V. Ajjarapu, C. Cañizares, I. Erlich, D. Hill, I. Hiskens, I. Kamwa, B. Pal *et al.*, “Stability definitions and characterization of dynamic behavior in systems with high penetration of power electronic interfaced technologies,” *IEEE PES Technical Report PES-TR77*, 2020.
- [3] H. Wang, Z. Lei, X. Zhang, B. Zhou, and J. Peng, “A review of deep learning for renewable energy forecasting,” *Energy Conversion and Management*, vol. 198, p. 111799, 2019.
- [4] D. Van der Meer, J. Munkhammar, and J. Widén, “Probabilistic forecasting of solar power, electricity consumption and net load: Investigating the effect of seasons, aggregation and penetration on prediction intervals,” *Solar Energy*, vol. 171, pp. 397–413, 2018.
- [5] R. Doherty and M. O’malley, “A new approach to quantify reserve demand in systems with significant installed wind capacity,” *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 587–595, 2005.
- [6] A. Papavasiliou, S. S. Oren, and R. P. O’Neill, “Reserve requirements for wind power integration: A scenario-based stochastic programming framework,” *IEEE Transactions on Power Systems*, vol. 26, no. 4, pp. 2197–2206, 2011.
- [7] M. Sun, C. Feng, and J. Zhang, “Probabilistic solar power forecasting based on weather scenario generation,” *Applied Energy*, vol. 266, p. 114823, 2020.
- [8] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee, “From local explanations to global understanding with explainable AI for trees,” *Nature Machine Intelligence*, vol. 2, no. 1, pp. 56–67, 2020.
- [9] T. Duan, A. Anand, D. Y. Ding, K. K. Thai, S. Basu, A. Ng, and A. Schuler, “NGBoost: Natural gradient boosting for probabilistic prediction,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 2690–2700.
- [10] S. M. Lundberg, B. Nair, M. S. Vavilala, M. Horibe, M. J. Eisses, T. Adams, D. E. Liston, D. K.-W. Low, S.-F. Newman, J. Kim *et al.*, “Explainable machine-learning predictions for the prevention of hypoxaemia during surgery,” *Nature Biomedical Engineering*, vol. 2, no. 10, pp. 749–760, 2018.

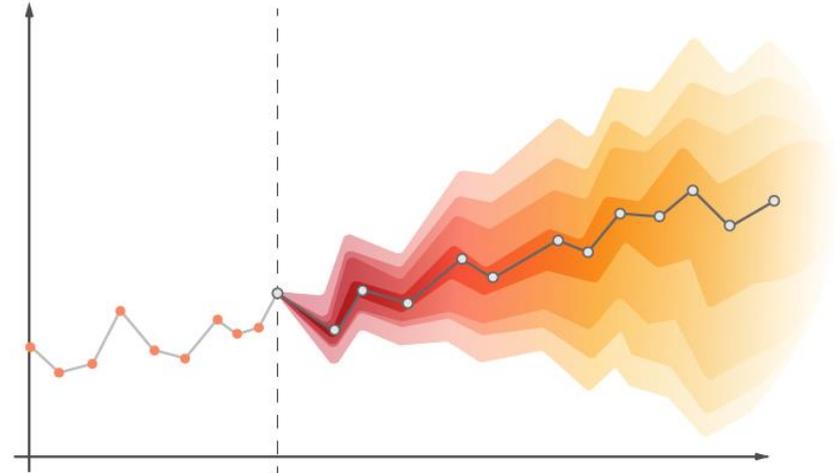


University of Stuttgart

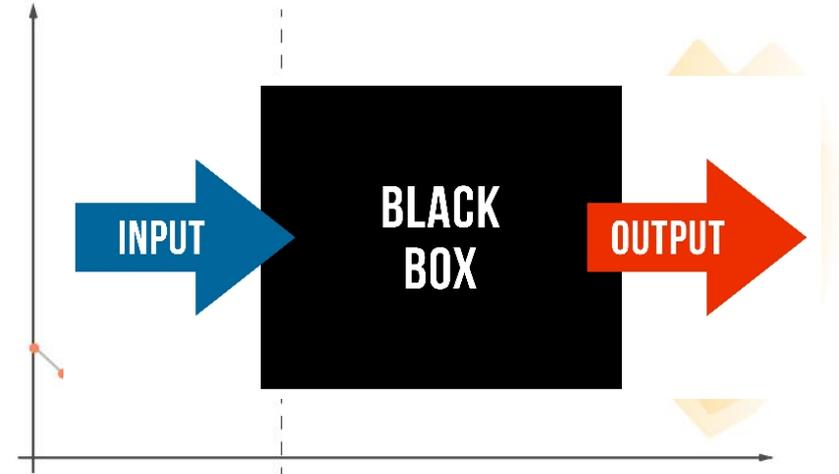


**Enabling Energy
Transition:
An Interpretable AI
Model for
Probabilistic Solar
Power Forecasting**

Motivation

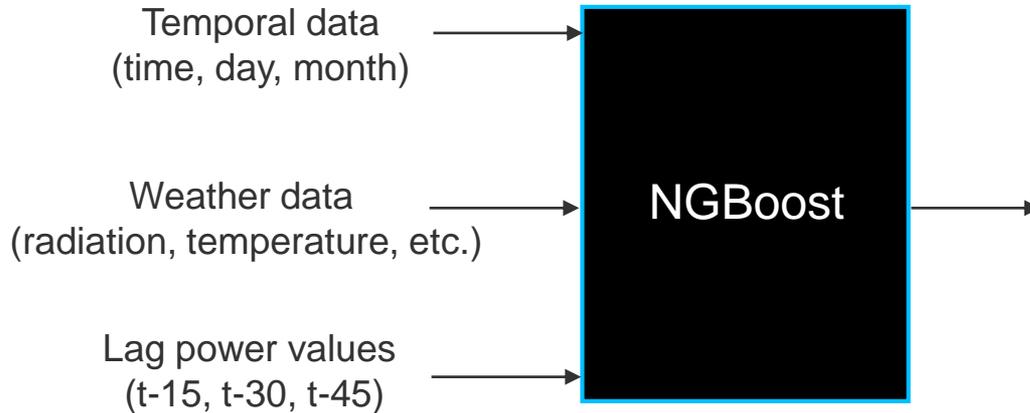


Motivation



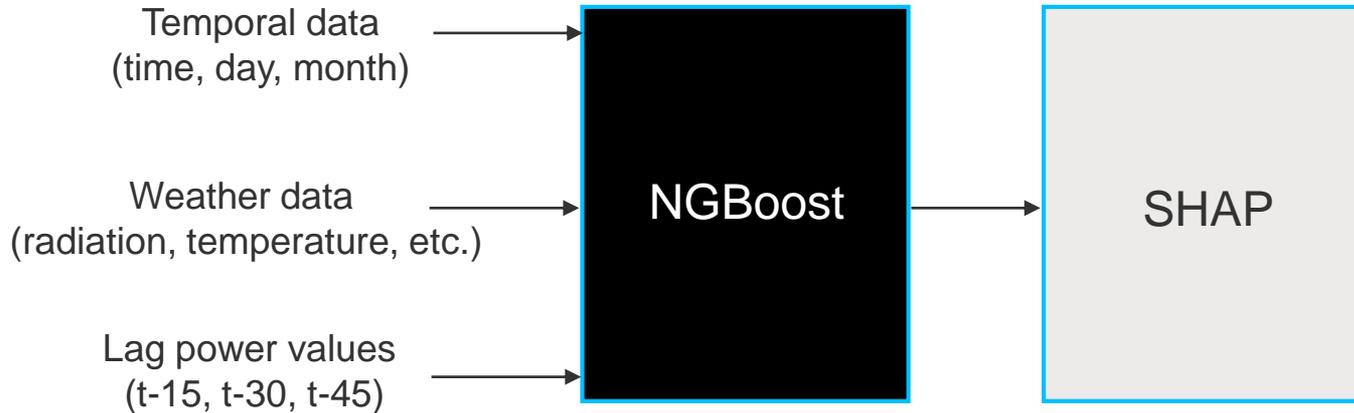
Our approach

- 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 is able to provide theoretical optimal **explanations** about the predictions of a model.



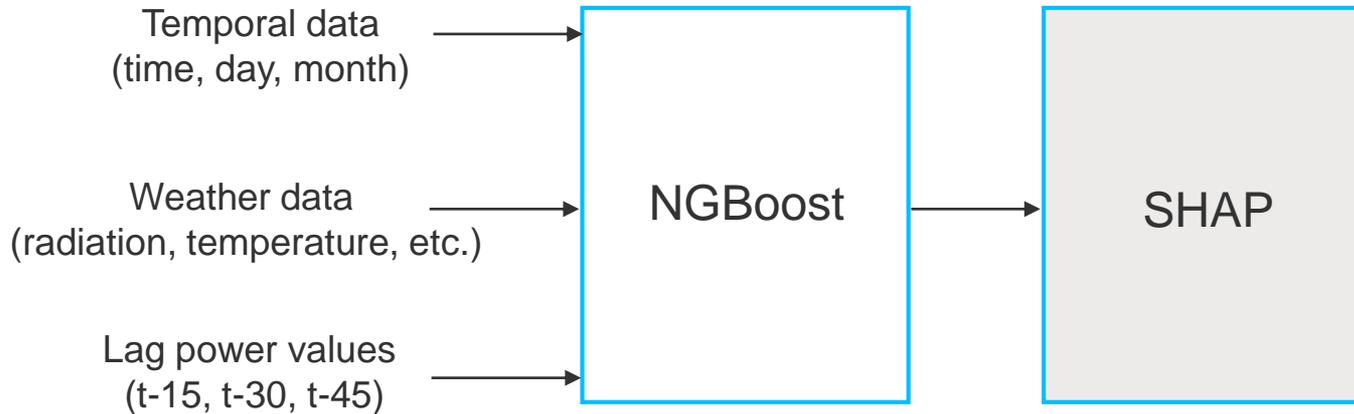
Our approach

- 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 is able to provide theoretical optimal **explanations** about the predictions of a model.



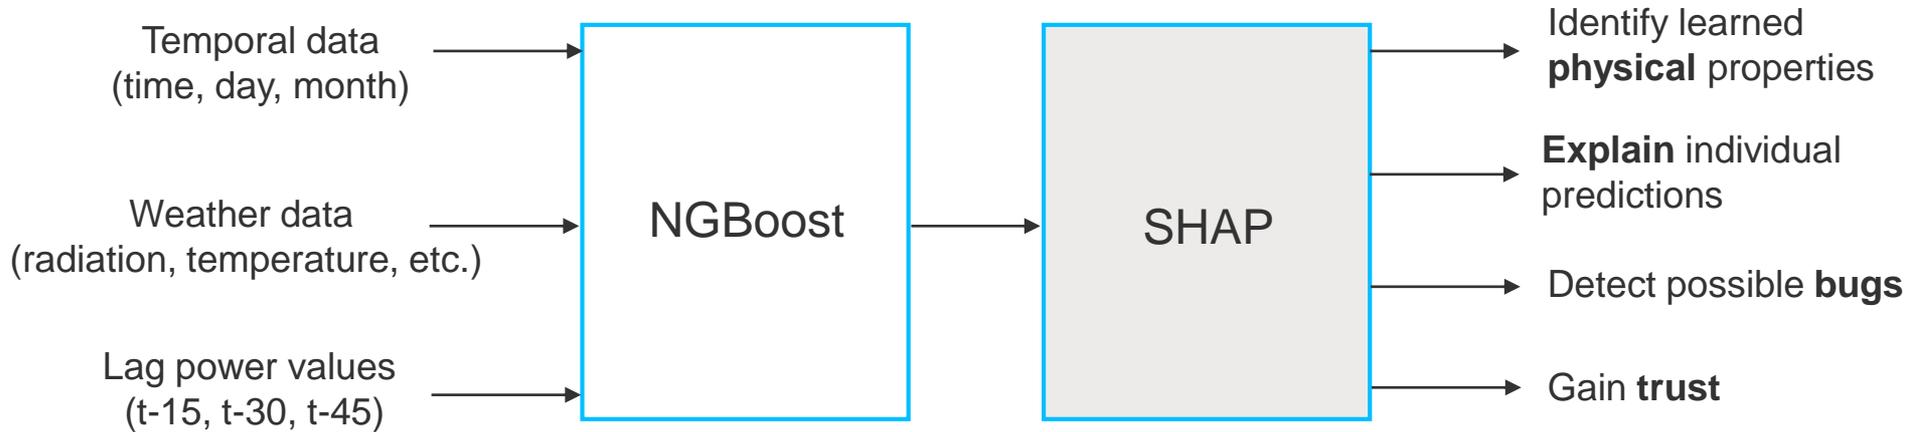
Our approach

- 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 is able to provide theoretical optimal **explanations** about the predictions of a model.

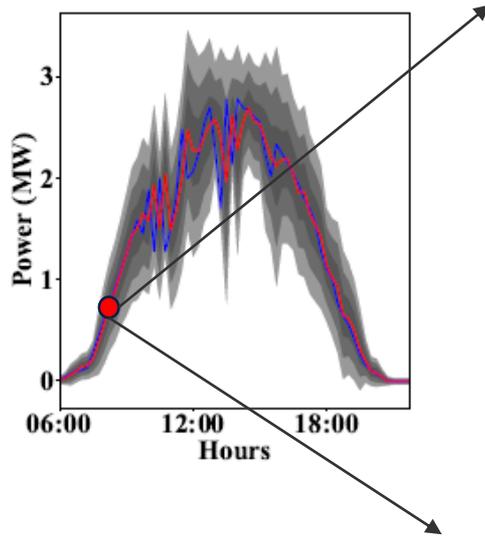


Our approach

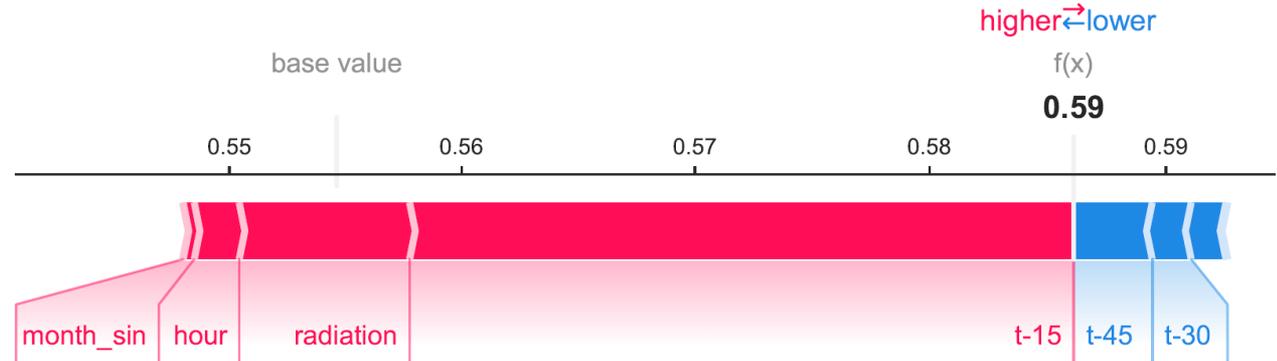
- We propose the application of the natural gradient boosting (**NGBoost**) algorithm for yielding **probabilistic** PV power **forecasts**.
- We combine it with Shapley additive explanation (**SHAP**) values, which is the unique consistent attribution method that is able to provide theoretical optimal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



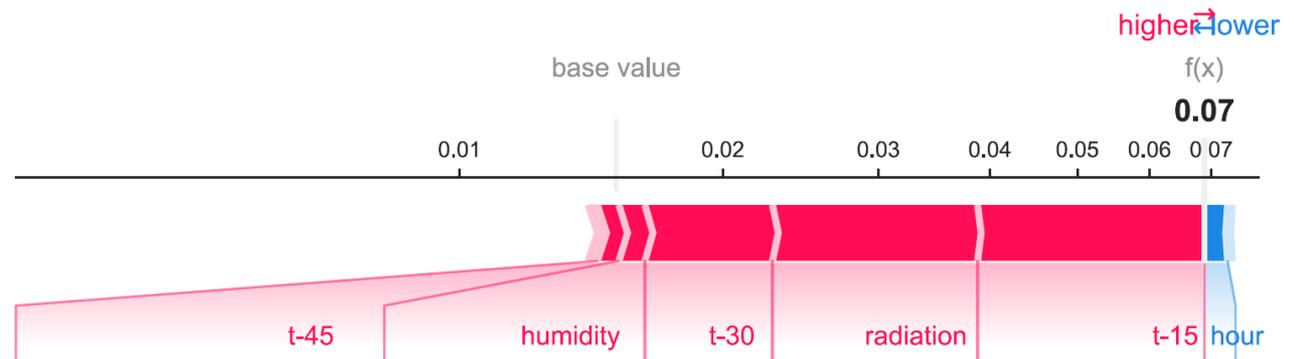
Our approach



Point forecast

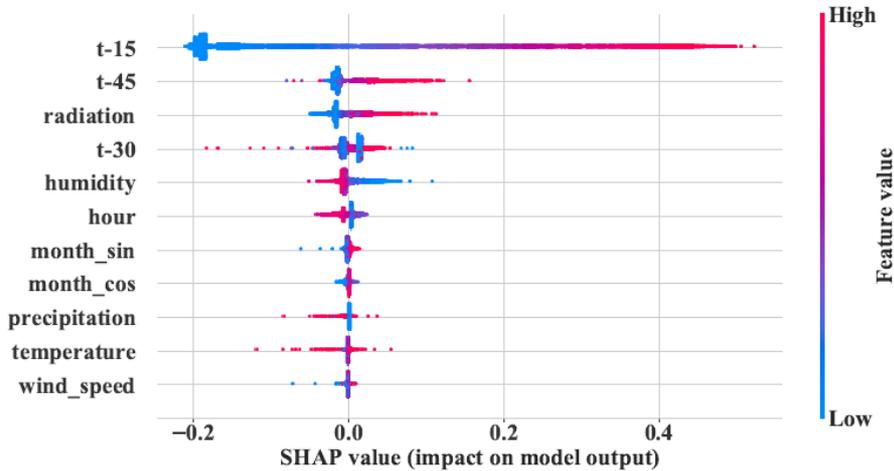


Uncertainty estimation (PIs)

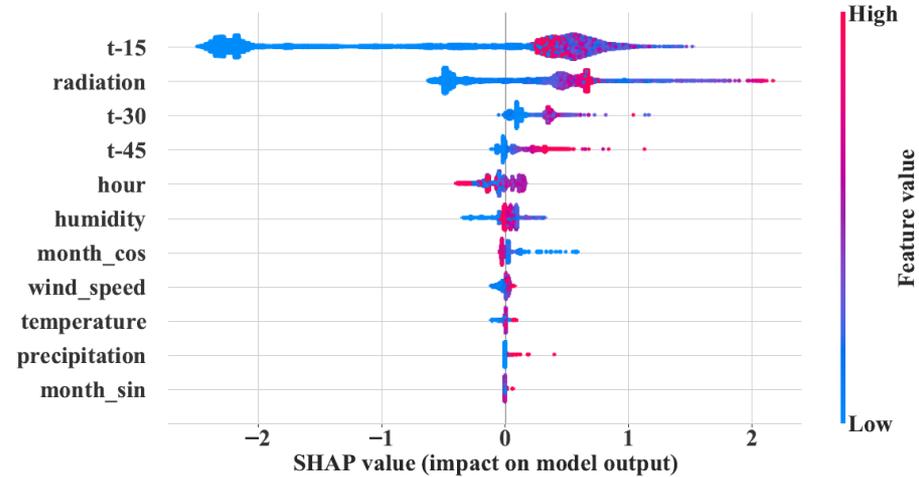


Model interpretation

Point forecast



Uncertainty estimation (PIs)

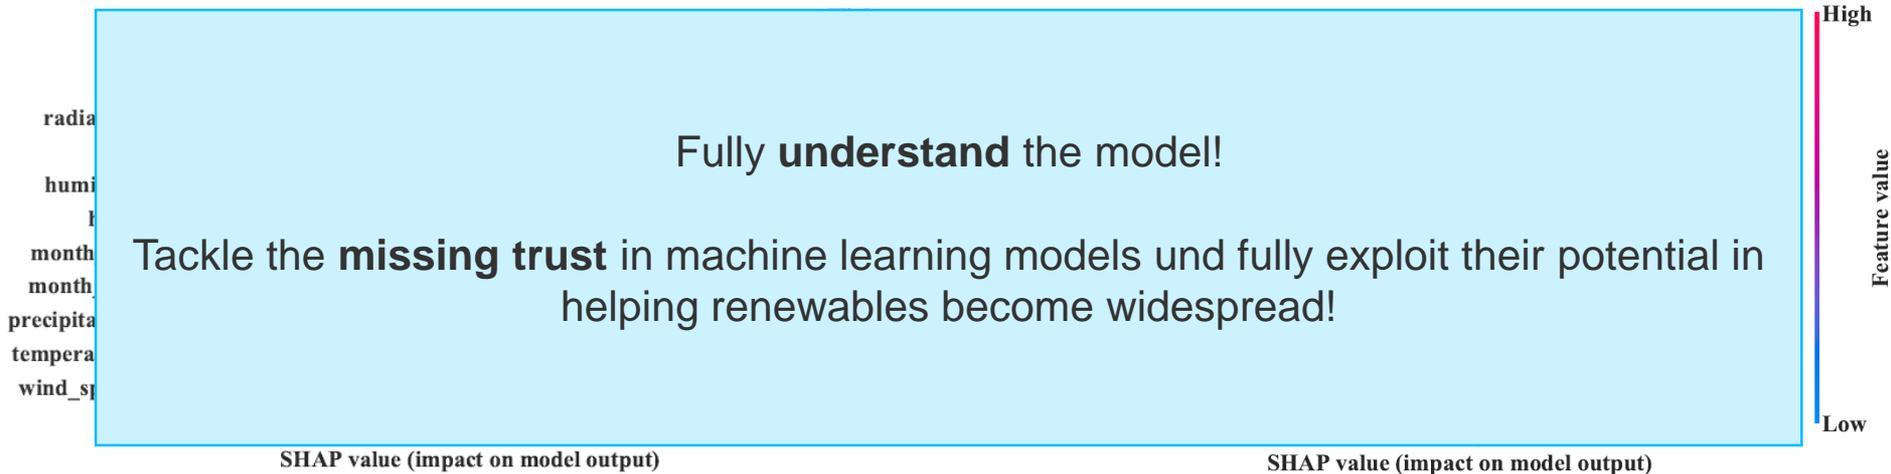


- Nonlinear **physics based interactions** between features **were learned**.
- The behavior of the model is aligned with **human logic and intuition**.

Model interpretation

Point forecast

Uncertainty estimation (PIs)



- Nonlinear **physics based interactions** between features **were learned**.
- The behavior of the model is aligned with **human logic and intuition**.



Universität Stuttgart

Vielen Dank!



Georgios Mitrentsis

georgios.mitrentsis@ifk.uni-stuttgart.de

+49 (0) 711 685 68916

Universität Stuttgart

Institut für Feuerungs- und Kraftwerkstechnik

Pfaffenwaldring 23 • 70569 Stuttgart