

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

Georgios Mitrentsis and Hendrik Lens

PV power forecasting models are predominantly based on machine learning algorithms which do not provide any insight into or explanation about their predictions (black boxes). To this end, we propose a two stage probabilistic forecasting framework able to generate highly accurate forecasts yet offering full transparency on its predictions. It is revealed that the model was able to develop complex nonlinear relationships which follow known physical properties as well as human logic and intuition.

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]. To this end, 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. As expected, probabilistic forecasts are more valuable compared to deterministic ones since they provide uncertainty information about the future power outcomes [3]. Yet probabilistic RES forecasting 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]. Furthermore, most 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. In this context, system operators may be reluctant in using those black box models in applications where safety critical decisions are to be made, e.g., power system

operation and control. 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.

To address those issues, a two stage interpretable probabilistic framework is proposed, which can generate highly accurate probabilistic forecasts while providing full transparency on its predictions. In the first stage, we propose the application of the natural gradient boosting (NGBoost) algorithm [4] for yielding probabilistic PV power forecasts. In the second stage, we calculate the Shapley additive explanation (SHAP) values [5] in order to interpret the derived model. Note 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.

Field and Weather Data

The power data employed in this study were acquired from two PV parks located in Southern Germany, with nominal capacity 3.2MVA and 1.8MVA, respectively. 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 $t-15$, $t-30$, and $t-45$, 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) and are assumed as the corresponding weather forecasts at those two locations (since the weather stations are located a few kilometers away from the respective PV parks). 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 data from one whole year while the test set comprises data from the following month.

Results

Figure 1 shows four indicative results as generated by the NGBoost algorithm using data from PVP1. As depicted in Figure 1, each plot illustrates a completely different generation pattern due to seasonal variations. Yet the NGBoost algorithm was able to yield highly accurate and sharp probabilistic forecasts, where the point forecasts lie

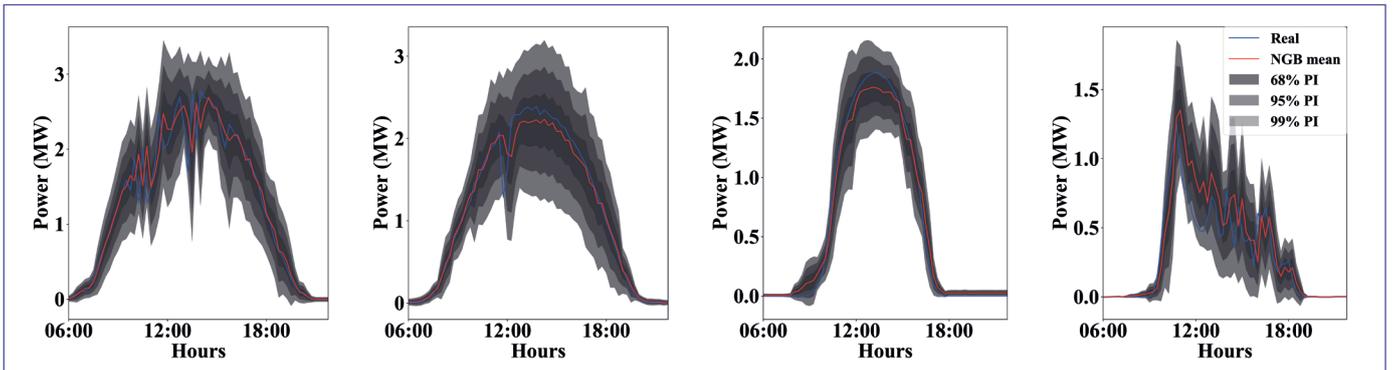


Figure 1 Indicative model output for four different test scenarios

always within the boundaries of the predicted distributions.

In order to interpret and understand the output of the probabilistic forecasting model, we deploy the SHAP summary plot for both the point forecast and the uncertainty estimation (prediction intervals of the generated distribution). The SHAP summary plots summarize the local explanations in a compact yet information-rich way for both model outputs (point forecast, prediction intervals), as shown in Figure 2. 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. 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 (left plot in Figure 2), 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 accurate 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 prediction intervals associated with the predictions (right plot in Figure 2) 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

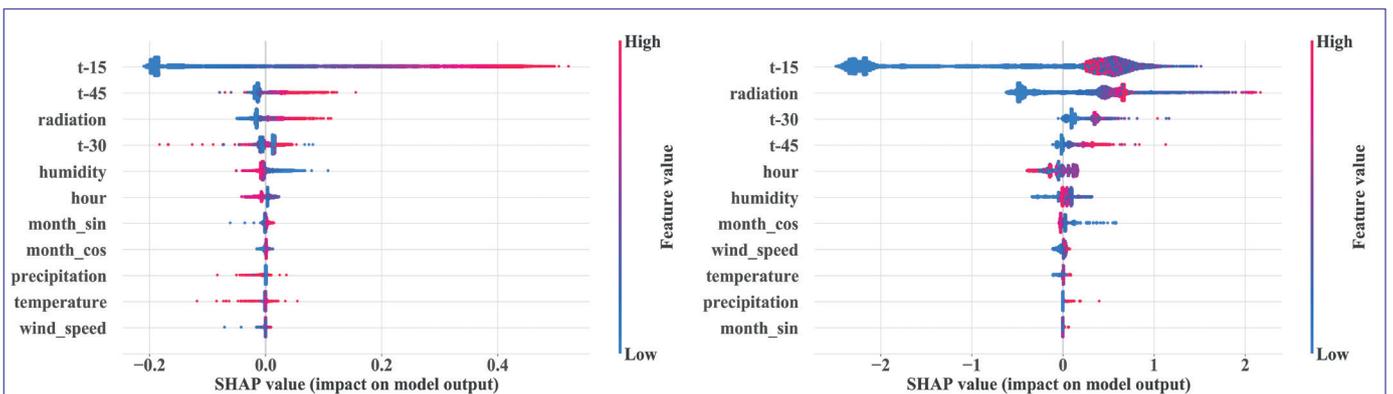


Figure 2 SHAP summary plots of PVP1 for the point forecast (left plot) and for the prediction intervals (right plot)

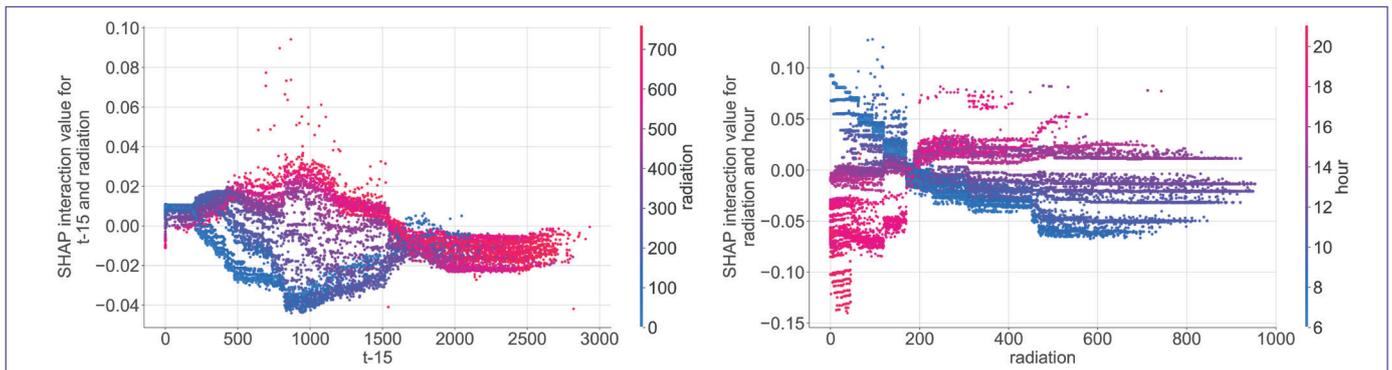


Figure 3 Two indicative SHAP interaction plots of PVP1 for the point forecast (left plot) and for the prediction intervals (right plot)

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 Figure 3. 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 the left plot of Figure 3 for the case of point forecast. As derived by this figure, $t-15$ values within the range of 0.5MW-1.5MW can have an utterly different influence on the model predictions depending on the radiation values. As expected, high radiation leads to higher and low radi-

ation to lower power predictions. For higher lagged power values greater than 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 indicative interaction plot between the time of the day (hour) and radiation is shown in the right plot of Figure 3 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 200W/m², the interaction between those features decreases. Yet early morning hours combined with high radiation tend to decrease the predictive uncertainty.

Conclusion

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, 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. Hatziaargyriou, J. Milanovic, C. Rahmann, V. Ajarapu, C. Canizares, 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] 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.
- [5] 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.

G. Mitrentsis, PhD candidate, H. Lens, Professor for power plant and grid systems, IFK - University of Stuttgart
georgios.mitrentsis@ifk.uni-stuttgart.de