

COMBINATION OF PHYSICS BASED SIMULATION AND MACHINE LEARNING FOR PV POWER FORECASTING OF LARGE POWER PLANTS

Nicolas Holland¹, Xiulan Pang², Wiebke Herzberg¹, Jefferson Bor¹ and Elke Lorenz¹
Fraunhofer Institute for Solar Energy Systems ISE, Heidenhofstr. 2, 79100 Freiburg, Germany¹
SPIC Huanghe Hydropower Development Co., LTD. HHDC²

ABSTRACT: Forecasting of PV power essentially relies on forecasting of the solar irradiance. There are different sources of irradiance forecasts like numerical weather predictions as well as satellite and measurement-based methods. Those different sources are typically combined to derive optimized forecasts for different forecast horizons. To derive PV power forecasts on this basis, new machine learning methods are a valuable addition to physics-based simulations. In our approach we combine not only different sources of irradiance forecasts but also the physics-based simulations with machine learning methods in order to produce optimized PV forecasts for large scale power plants and apply the method on a power plant (>800MW) in the Qinghai province of China. This PV power plant is frequently subject to curtailment by the grid operator for large power values, which presents a challenge for model training that we address by applying a mix of active power control thresholding and a machine learning model trained to detect curtailment in PV power measurements.

Keywords: PV Forecast, Satellite, Machine learning

1 INTRODUCTION

Forecasting the power output of large PV power plants is often attached with significant challenges ranging from the presence of different technologies within the power plant to external power regulations that effect the power production and can have little to do with weather situations. We developed a PV power forecasting method applied and tested for a power plant of over 800MW installed capacity in the Qinghai province of China. Due to grid regulations, power plant operators are obligated to provide power production forecasts to the grid operator and are only allowed to feed in power according to their predictions [1]. In cases where the predicted power is below what could actually be produced, grid operators may apply curtailment, which leads to missed opportunities for selling power and, in cases in where the produced power is significantly below the predicted power, there can be penalty fines that have to be paid to the grid operator. This situation creates a demand for reliable forecasting. Due to the nature of the regulation the 2-hour ahead forecast plays an important role in the evaluation of the forecasts.

Our first approach was a model, that relied on numerical weather predictions (NWP) and measurement based global horizontal irradiance (GHI) forecasts (persistence) which were combined using a regression model to create an optimized GHI forecast. This GHI forecast was then used as input for a PV power simulation alongside ambient temperature data resulting in a power forecast. The method is largely based on previous work by Kühnert et. al [2] and yielded promising results, we first described the approach at the IEEE PVSEC-46 [3].

Since then we improved the forecast performance by introducing satellite based GHI forecasts and by replacing the regression model with a neural network. Instead of combining GHI forecasts, we apply PV power simulations on the GHI forecasts first and use the neural network to combine the power predictions into an optimized PV power prediction.

One of the major challenges to model training is the fact that many of the power measurements are affected by curtailment. When training a model with respect to power measurements, curtailed power creates a bias to underestimate the true power output potential, because the power plant output is artificially lower than what weather

and GHI inputs would suggest. We address this challenge by both applying curtailment detection methods on the training data and therefore removing parts of the dataset that would hinder model training, and by adjusting the model training loss function such that curtailment artefacts that remain undetected within the training data have a less severe impact on the resulting model. Both the introduction of satellite data, the neural network and the curtailment detection improve the forecasting system to the point where it can achieve the level of accuracy required by the grid regulation.

2 DATA

2.1 On-site measurements

The Qinghai based power plant has an over 800MW installed capacity and contains different PV technologies including mono- and bifacial modules as well as tracked and fixed orientations.

GHI measurements were taken with a horizontal standard pyranometer and power measurements were obtained for the overall PV power output of the entire plant. Both GHI and power measurements were in 1-minute resolution, which were averaged to 15-minute time series. The grid regulation requires forecasts to be in 15-minute resolution. Next to the measurements, the on-site available data includes Active Power Control (AGC) data, which is a signal sent by the grid operator to regulate and curtail the power plant. Our curtailment detection methods rely both either directly or implicitly on this AGC data.

The data which our analysis is based on ranges from March 2020 to April 2021 and has a few weeks of missing data in August 2020. AGC data was available from March to Mai of 2020.

2.2 Numerical weather predictions

In addition to on-site measurements the forecasting method relies also on numerical weather predictions. We use forecasts by the European Centre for Medium-Range Weather Forecasts (ECMWF), which are updated four times per day and have a forecast horizon of up to several days. To resample the hourly ECMWF values to a 15-minute resolution, the values are divided by modelled clearsky irradiance to compute clearsky indices, which then are interpolated to produce 15-minute values. We use

the clearsky model by Dumortier [4] to model clearsky irradiance. Temperature forecasts by ECMWF were used as input to the PV simulation method.

3 METHOD

The entire PV power forecast procedure consists of several methods that are combined and adjusted to form one forecasting system. First, we use methods for predicting GHI irradiance like measurement-based persistence forecasts and satellite-based forecasts. Next, these forecasts together with external NWP data are input to a physics-based PV power simulation, which are then the bases for a neural network combi model. The neural network was trained on a dataset that was subject to our curtailment detection, which itself consists of an AGC-based thresholding method or a machine learning model for cases in which AGC data is not available.

3.1 Persistence forecasts

Persistence GHI forecasts (G_{pfc}) are based on both on-site GHI measurements and a clearsky irradiance model. Given a GHI measurement at time t and assuming the current weather situation does not change, a persistence forecast is computed by multiplying the current clearsky index with a clearsky model,

$$G_{pfc,t+1} = \frac{G_{meas,t}}{G_{cs,t}} G_{cs,t+1}$$

with $G_{meas,t}$ and $G_{cs,t}$ being the GHI measurement and clearsky model prediction at time t . Both the irradiance measurements and modelled clearsky irradiance are averaged to 15-minute resolution time series.

3.2 Satellite based forecasts

Satellite based GHI forecasts (G_{sat}) are derived from satellite images using the heliosat method [5] and cloud motion vectors [6]. We use satellite data in the visible spectrum from the geostationary Himawari satellite operated by the Japan Meteorological Agency. The satellite makes new measurements every 10 minutes, which canonically leads to GHI forecasts with a 10-minute resolution. The core aspect of the forecasting algorithm is the calculation of cloud index information from satellite images which in combination with clearsky models produces GHI estimations. Using two consecutive cloudindex images allows to compute a vector field describing the cloud motion between the two images. This vector field is referred to as either optical flow or cloud motion vectors. Assuming that the cloud movement is persistent the vector field is then used to extrapolate cloud movement into the future thereby creating cloudindex forecasts and ultimately GHI forecasts.

The cloud motion vectors are computed using the deepflow method [7] implemented in the OpenCV library [8]. The entire satellite based GHI forecast process can be summarized as this,

$$\begin{aligned} \text{heliosat}(CI_t) &\rightarrow G_{sat,t} \\ \text{deepflow}(CI_{t-1}, CI_t) &\rightarrow \begin{pmatrix} u \\ v \end{pmatrix}_t \\ \text{map}\left(CI_t, \begin{pmatrix} u \\ v \end{pmatrix}_t\right) &\rightarrow CI_{t+1} \\ \text{heliosat}(CI_{t+1}) &\rightarrow G_{sat,t+1} \end{aligned}$$

With heliosat and deepflow being the in [6] and [7] described methods, CI_t the derived cloudindex at time t , $\begin{pmatrix} u \\ v \end{pmatrix}_t$ the cloud motion vector and CI_{t+1} or $G_{sat,t+1}$ the forecasted cloudindex or irradiance. The forward extrapolation in time is done via a simple inverse mapping,

$$CI_{t+1}(x, y) = CI_t(-u_t(x, y) + x, -v_t(x, y) + y)$$

The resulting 10-minute resolution GHI forecasts are resampled to match the 15-minute resolution of the other irradiance and power time series. Figure 2 shows an example of a satellite image as it is used in the forecasting process, the Qinghai region is in the upper left area of the image.

Another approach that we observed with great interest is the satellite-based forecasting method by Harty et. al. [9]. While our cloud motion vectors used for cloudindex forecasting are solely derived from optical flow computed from satellite images, they combine optical flow with wind vector fields derived from numerical weather predictions and compute cloud motion vectors that appear to be outperforming the optical flow only forecasts and could potentially improve our system as well in future works.

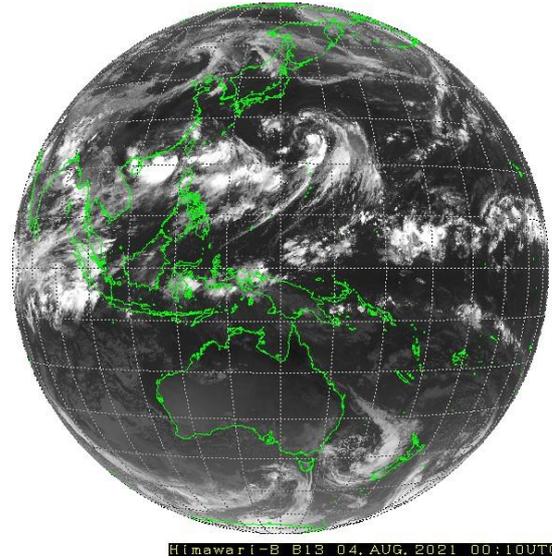


Figure 2: Himawari satellite image in visible spectrum, image taken on 2021-08-04 00:10 UTC. Satellite images like this one are the bases for cloud motion vector and satellite irradiance forecasts. Image by Japan Meteorological Agency [10].

3.3 PV simulation

The first step in creating PV power forecasts is to simulate the power plant with different GHI forecasts as input. The splitting of GHI into its direct and diffuse components is done using the DIRINT model [11] and the conversion from horizontal irradiance to plane of array irradiance is done using the model by Perez et al [12]. We use the implementations of those models from the PVLIB library [13]. The models used to derive PV power from plane of array irradiance are implemented in-house and are described in [14].

Since the power plant consists of different technologies and configurations the simulation is helped by organizing the power plant into suitable sub-plants which are simulated independently. The results of those simulations

are then summed up to describe the entire power plant output. We apply this simulation approach on every GHI forecast individually so that, for all $s \in \{nwp, pfc, sat\}$

$$\text{simulate}(G_s, T, \theta) \rightarrow P_s$$

with G_s being the irradiance forecast generated through method s and T, θ being the ambient temperature and sun angles ie. zenith and elevation. Figure 3 shows the result of the PV simulation applied on satellite GHI forecasts with a 2-hour forecast horizon.

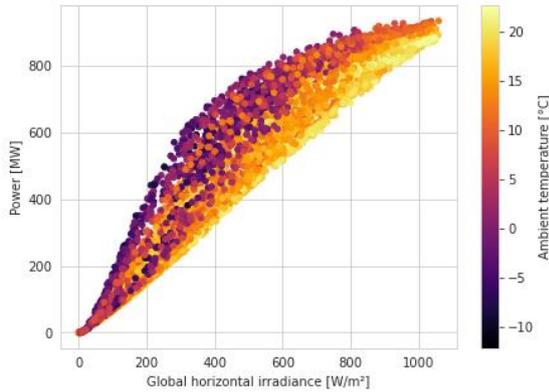


Figure 3: Scatter plot of PV simulation and GHI satellite data with ambient temperature from ECMWF forecasts. We see a spread of power values for similar GHI values, this can be explained by them occurring at times with different solar angles, which can have a large effect especially on parts of the power plant with tracked modules. We also see a temperature dependency, as illustrated here with different colors.

3.4 Curtailment detection

Figure 4 shows an example of a day where power measurements were likely curtailed, and the curtailment was detected through one of our models. Our methods rely on either the availability of AGC data to make a direct prediction over whether a power measurement was affected by curtailment or a prediction is made by a machine learning model trained on curtailment information, which in turn can have been produced using AGC. The actual truth of when the power plants were curtailed is held by the grid operator and was in this case not known to us. Our curtailment detection method was first described in a talk given at PVSEC-30 2020 [15].

3.4.1 Curtailment detection using AGC

The AGC approach is basically a thresholding approach where the difference between AGC and power measurement is compared with a fixed percentage of the overall installed capacity,

$$C_{AGC}(P_t) = \begin{cases} 1, & |AGC_t - P_t| < 1\% \text{ inst. kWp} \\ 0, & \text{otherwise} \end{cases}$$

with $C_{AGC}(P_t)$ being a truth value over whether the power measurement at time t was curtailed. The percentage threshold was adjusted through manual inspection.

3.4.2 Curtailment detection using RDF

In the absence of AGC information the curtailment detection relies on a random decision forest (RDF) which

was trained on labels that were produced by the AGC method. Besides both power and irradiance measurements P_{meas}, G_{meas} the input features used to make RDF predictions include simulated power P_{sim} and clearsky index $\frac{G_{meas}}{G_{cs}}$ and a set of time series features that are derived from those four. We use the Scalable Hypothesis method [16] to derive those features, which is a two-step process. First, we compute a large set of time series features that could potentially be relevant, before rating them individually as to how well every feature can predict the ground truth curtailment labels. Individual features are rated using the Benjamini-Yekutieli procedure [17] and because this process produces relevant time series features without manual feature engineering, it is often referred to as automatic feature extraction. For the feature computation we use the implementation from the tsfresh library [18] and for the RDF we use the implementation in the scikit-learn library [19]. The training set for the RDF model was the time range from March 2020 to Mai 2020 in which AGC data was available.

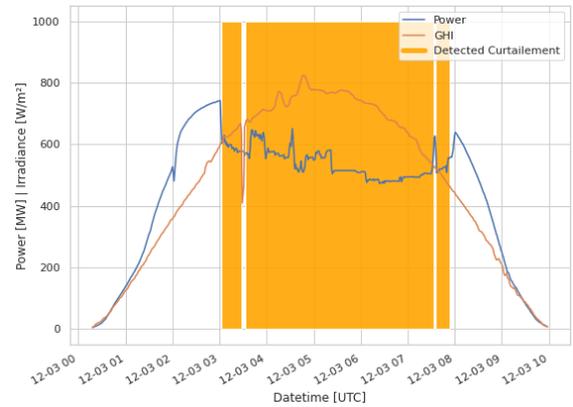


Figure 4: Power and GHI measurements for single day in December. There are visible curtailment effects in the power measurement which were detected by the RDF model. Note that the model was trained on and is therefore meant to be used on minute resolution data.

3.5 Model training

To combine the different power forecasts into a single best power prediction we use a densely connected neural network consisting of two hidden layers which take as input both the three different power forecasts as well as the original GHI forecasts and solar zenith angles. The neural network is trained to produce complete forecasts, ie all the forecast steps are produced by the same network at once. The differentiation between forecast steps is achieved by adding the forecast horizon as an additional input feature. This enables the model to learn weighting satellite, persistence, and NWP inputs dependent on the forecast horizon. The training is done using Adam optimization [20] and we use the network implementation from the keras library [21]. Model training was done on the curtailment filtered dataset.

3.5.1 Train/test split

The entire dataset we used for the analysis ranges from March 2020 to April 2021, the data used for forecast model training ranges from March 2020 to September 2020 with everything starting from October 2020 being solely used for testing.

3.5.2 Loss function

We started our forecast model training minimizing the root-mean-square-error (RMSE) as a loss function. In the regular RMSE, given by

$$RMSE = \sqrt{\frac{1}{N} \left(\sum_t (P_{p,t} - P_{m,t})^2 \right)}$$

with $P_{p,t}, P_{m,t}$ being the predicted and measured power at time t , over- and underestimations are treated equally.

However, even with curtailment detection applied on training data, there are still issues which can have effects on model training. Both undetected curtailment as well as simple lack of data in situations with high GHI values, and a correlation clearly smaller than one introduces incentives for the model training to produce models which tend to underestimate the true output potential of the power plant for larger GHI values. To mitigate the effect of this tainted training data we make an adjustment to the loss function of the model training. We introduce a skewed version of the RMSE where over- and underestimations are weighted differently by preset parameters α, β

$$\sqrt{\frac{1}{N} \left(\alpha \sum_{t \in \Omega_\alpha} (P_{p,t} - P_{m,t})^2 + \beta \sum_{t \in \Omega_\beta} (P_{p,t} - P_{m,t})^2 \right)}$$

with

$$\Omega_\alpha = \{t: P_{p,t} < P_{m,t}\}$$

$$\Omega_\beta = \{t: P_{p,t} > P_{m,t}\}$$

and $\alpha > \beta$, the exact values for α, β are defined through manual inspection. This way the underestimations are weighted higher, leading to more impact on the model training, which is desired assuming that some of the overestimations are caused by curtailment and can therefore be accepted.

3.6 Metrics

In addition to the RMSE as a performance metric we also consider the accuracy metric issued by Chinese grid operators, which is used to assess and regulate power plants in the Qinghai province as well as other provinces [1],

$$1 - 2 \sum_{t=1}^N \left(\left| \frac{P_{p,t}}{P_{p,t} + P_{m,t}} - \frac{1}{2} \right| \frac{|P_{p,t} - P_{m,t}|}{\sum_{i=1}^N |P_{p,i} - P_{m,i}|} \right)$$

with N the number of values per day and $P_{p,t}, P_{m,t}$ being the predicted and measured power at time t . The metric is used to assess the forecast accuracy of entire days and is typically applied on the 2 hours ahead forecast. Only value pairs with either the predicted or measured power above 3% of the installed capacity are considered in the assessment.

4 RESULTS AND DISCUSSION

4.1 Curtailment

The detection of curtailment is an important part of both model training and evaluation, as curtailed power measurements can introduce strong biases into models trained on them. While in the presence of AGC

information, we prefer using the thresholding method to detect curtailed events, our results show that the machine learning based RDF method can be a feasible alternative for situations without AGC information. Figure 5 shows the curtailment detection results for the two methods.

The dataset is split into two parts, with the first part being used as training set for the RDF model and the second as test set. Without curtailment detection there is a bias in the training set of 9.2% with respect to the installed capacity, and 6.2% in the test set. With curtailed datapoints filtered out, there remain biases of 1.9% on each of the two parts of the dataset, which is a big improvement but also suggests that there might be undetected curtailment.

When comparing the AGC method with the RDF method on the part of the dataset where both methods are applicable, we find that the AGC method flags 810 of the 5855 datapoints while the RDF method flags 882 datapoints, with 633 being flagged by both methods in agreement. In total, the curtailment detection process removes 15.2% of datapoints in the entire dataset. The figure also shows that there is an underrepresentation of situations with large power outputs in the dataset.

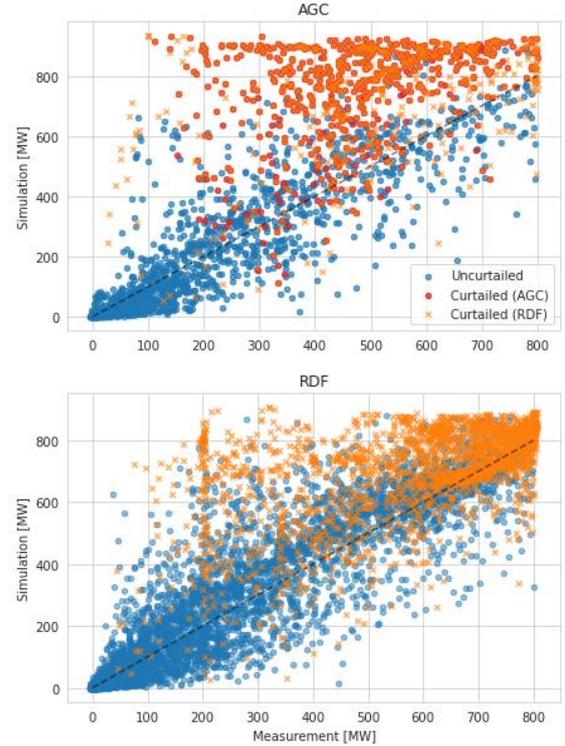


Figure 5: Scatter plots of satellite-based over measured power with curtailment detected using the AGC and RDF methods. Top: Training set (March to Mai 2020) with AGC available. Bottom: Test set (June 2020 to April 2021) without AGC available. There appear to be common events where power measurements are curtailed to 200MW, many of which were identified by the curtailment detection.

4.2 Forecast accuracy

The forecasts are evaluated using both the RMSE metric, which is widely used, especially in PV power forecasting and the daily accuracy metric, which is used in the Qinghai province where the power plant is located. Figure 6 compares power forecasts made from satellite or NWP data as well as combi models trained with regular RMSE

and our skewed RMSE as loss functions for the curtailment filtered part of the dataset. Both the skewed and regular combi models have learned to reduce the scattering while simultaneously introducing a bias towards underestimating the power for situations with above 700MW power output. This can be caused by undetected curtailment still being present in the training data and because in cases with a correlation < 1.0 the RMSE is being reduced by reducing the amount of extreme values. The skewed model has learned to be less biased towards underestimation for high power values in comparison, but also has more overestimations in situations with small values.

Since the daily accuracy metric is of importance for the Qinghai region, figure 7 shows the distribution of daily accuracies from the skewed combi model. We find that the combi model achieves to meet the 75% accuracy threshold requirement in 50.8% of days. It is important to note that datapoints that were affected by the grid operator through curtailment are to be excluded from this evaluation metric. Since it is only the grid operator who knows the exact ground truth over curtailment status, their forecast evaluation may differ, and our evaluation is meant to generate an expectation of what their results could be. Detailed evaluations for the different forecast models and datasets with respect to both RMSE and daily accuracy are shown in table 1 and 2 respectively. Figure 8 shows an example of a single forecast for a day with mostly clearsky conditions which illustrates the effect of either undetected curtailment or underrepresented high irradiance situations in the training set on both combi models and how the skewed model has learned to mitigate the effect.

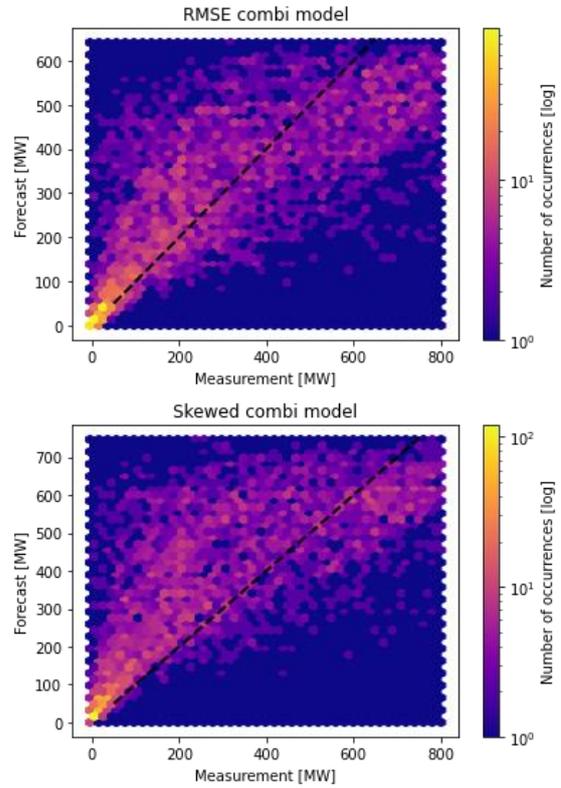


Figure 6: Scatter plots of forecasted power over measured power for the 2-hour ahead horizon comparing satellite and NWP based forecasts with the combi models. Note that the Forecast axis do not have equal range on all the plots.

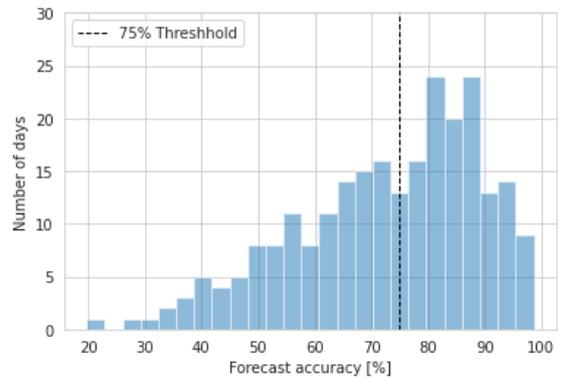
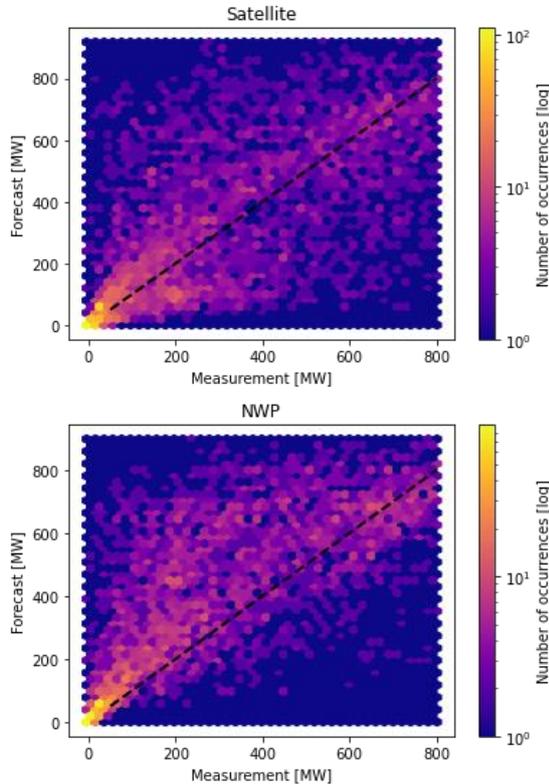


Figure 7: Distribution of daily accuracy values for the skewed combi model evaluated on the entire dataset with curtailment filtered out. With 75% accuracy being the threshold set by the energy regulation which must be met as a minimum requirement in order to avoid being penalized for providing insufficient forecasts.

When comparing the model performances on the different datasets in table 1 and 2 we find that

- All the models, as expected, have higher RMSE values on the dataset that is not curtailment filtered.
- The satellite-based forecast outperforms the NWP forecast on the 2-hour ahead horizon.
- Since the combi models have reduced scattering, their RMSE values are lower than the satellite and NWP forecasts.

- The RMSE of the model that was trained on RMSE is evidently lowest.
- Some of the performance differences found with respect to RMSE were also found with respect to the daily accuracy metric, e.g. both the combi models were outperforming the satellite and NWP forecasts and that the combi models appear to perform better on the evaluation set, which indicates fewer curtailment situations being present there.

| Dataset | Satellite | NWP | CombiS | CombiR |
|---------|-----------|-------|--------|--------|
| T_A | 22.7% | 23.4% | 19.1% | 17.0% |
| E_A | 18.5% | 18.9% | 14.6% | 14.7% |
| F_A | 21.9% | 22.5% | 18.3% | 16.5% |
| T_C | 19.3% | 20.5% | 17.9% | 15.0% |
| E_C | 17.7% | 19.3% | 15.4% | 13.8% |
| F_C | 19.0% | 20.3% | 17.5% | 14.7% |

Table 1: Results for RMSE metric computed on the training dataset (T), the test dataset (E) and the full dataset (F) for both the curtailment filtered part of the dataset (C) and all the data (A). The RMSE values are normed with respect to the installed capacity. CombiS is the combi model trained with the skewed RMSE and CombiR the model with the regular RMSE.

| Dataset | Satellite | NWP | CombiS | CombiR |
|---------|-----------|-------|--------|--------|
| T_A | 65.7% | 68.9% | 71.0% | 73.6% |
| E_A | 76.7% | 77.2% | 79.3% | 80.5% |
| F_A | 71.1% | 73.0% | 75.1% | 77.0% |
| T_C | 64.2% | 68.2% | 69.0% | 72.2% |
| E_C | 74.3% | 74.7% | 76.5% | 78.1% |
| F_C | 68.9% | 71.3% | 72.5% | 75.0% |

Table 2: Mean of daily accuracy metric evaluated on the same data as in table 1. Both the RMSE and skewed combi models outperform the forecasts based solely on satellite or NWP data having similar performance improvements when compared to the RMSE metric.

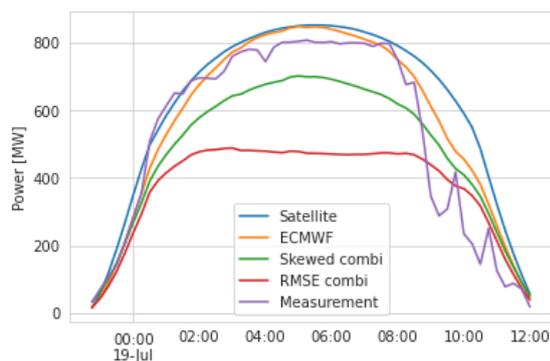


Figure 8: PV power forecasts computed on a day in July 2020 compared with the power measurement of that day. At the time at which the forecasts were computed both the satellite and NWP based methods predicted mostly clearsky conditions which for the majority of the day turned out to be the case. Despite the optimistic input predictions both the skewed and RMSE combi models appear to have bias towards underestimating the power production with the RMSE combi model being significantly more pessimistic. It is precisely this pessimistic behavior for otherwise optimistic situations, which lead us to introduce the skewing of the loss function

into the model training.

5 SUMMARY AND OUTLOOK

With machine learning in general relying heavily on clean and rich datasets, the absence of reliable clean ground truth data makes model training under this condition a challenge. Our results show that model training for power forecast enhancement is feasible even with curtailment issues tainting power measurements and introducing biases. Our curtailment detection method was able to find many of the curtailed measurements, but the training of the model remains to be challenging, especially because the dataset gets reduced and not all the curtailment events are detected. Through the introduction of a skewed loss function that treats over- and underestimations differently, the aim of avoiding large underestimations for situations with high power values is achieved. However, in this first approach it is also related to reduced forecast quality for low power values. Further investigations on metrics that can reflect model performances under assumed curtailment and model training specialized for curtailment situations will follow, as will further curtailment detection methods, for example, through the usage of uncurtailed reference arrays.

6 REFERENCES

- [1] Northwest China Energy Regulatory Bureau of National Energy Administration of the People's Republic of China, <http://xbj.nea.gov.cn/website/Aastatic/news-196276.html>
- [2] Wolff, B., J. Kühnert, E. Lorenz, O. Kramer, D. Heinemann: Comparing support vector regression for PV power forecasting to a physical modeling approach using measurement, numerical weather prediction, and cloud motion data", *Solar Energy*, Volume 135, October 2016, Pages 197-208, ISSN 0038-092X, doi:10.1016/j.solener.2016.05.051.
- [3] Holland, Nicolas, et al." Solar and PV forecasting for large PV power plants using numerical weather models, satellite data and ground measurements." 2019 IEEE 46th Photovoltaic Specialists Conference.
- [4] Dumortier, D., 1995. Modelling global and diffuse horizontal irradiances under cloudless skies with different turbidities. Daylight II, jou2-ct92- 0144, final report vol. 2. Tech. rep., CNRS-ENTPE.
- [5] Hammer A, Heinemann D, Hoyer C, Kuhlemann R, Lorenz E, Muller R, Beyer HG. Solar energy assessment using remote sensing technologies. *Remote Sensing of Environment* 2003; 86(3): 423–32.
- [6] J. Kühnert, E. Lorenz, D. Heinemann: 'Satellite-Based Irradiance and Power Forecasting for the German Energy Market' in *Solar Energy Forecasting and Resource Assessment*, Editor: Jan Kleissl (Elsevier 2013).
- [7] Weinzaepfel, P., J. Revaud, Z. Harchaoui, C. Schmid, 2013: Deepflow: Large displacement optical flow with deep matching. – ICCV – IEEE International Conference on Computer Vision 1385–1392.
- [8] Bradski, G., 2000: The opencv library. – Dr. Dobb's Journal of Software Tools.
- [9] Harty, Travis M., William F. Holmgren, Antonio T. Lorenzo, and Matthias Morzfeld. "Intra-hour cloud index forecasting with data assimilation." *Solar Energy* 185 (2019): 270-282.
- [10] Himawari satellite image, Full Disk, Meteorological

Satellite Center of Japanese Meteorological Agency,
<http://www.data.jma.go.jp/mscweb/data/himawari/>,
visited 2021-08-04.

- [11] Perez, R., P. Ineichen, E. Maxwell, R. Seals and A. Zelenka, (1992). "Dynamic Global-to-Direct Irradiance Conversion Models". ASHRAE Transactions-Research Series, pp. 354-369
- [12] Perez, R., Ineichen, P., Seals, R., Michalsky, J., Stewart, R., 1990. Modeling daylight availability and irradiance components from direct and global irradiance. *Solar Energy* 44 (5), 271-289.
- [13] William F. Holmgren, Clifford W. Hansen, and Mark A. Mikofski. "pvlb python: a python package for modeling solar energy systems." *Journal of Open Source Software*, 3(29), 884, (2018). <https://doi.org/10.21105/joss.00884>
- [14] Muller, B., L. Hardt, et al. (2015). "Yield predictions for photovoltaic power plants: empirical validation, recent advances and remaining uncertainties." *Progress in Photovoltaics: Research and Applications*: published online
- [15] Holland, Nicolas, et al. "Trainable curtailment detection for the enhancement of PV power forecasting." 2020 PVSEC-30 & GPVC 2020, November 8 - 13, 2020 ICC JEJU, Jeju, Republic of Korea
- [16] Christ, M., Kempa-Liehr, A.W. and Feindt, M. (2016). Distributed and parallel time series feature extraction for industrial big data applications. *ArXiv e-prints*: 1610.07717 URL: <http://adsabs.harvard.edu/abs/2016arXiv161007717C>
- [17] Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, 1165–1188
- [18] tsfresh, (2020). Retrieved from <https://tsfresh.readthedocs.io/en/latest/index.html>
- [19] Pedregosa et al., Scikit-learn: Machine Learning in Python, *JMLR* 12, pp. 2825-2830, 2011.
- [20] Kingma, Diederik P., and Jimmy Ba. "Adam: A method for stochastic optimization." *arXiv preprint arXiv:1412.6980* (2014).
- [21] François Chollet et. al., Keras, (2020), GitHub repository, <https://github.com/keras-team/keras>