

A chained neural network model for photovoltaic power forecast

Carola Gajek, Alexander Schiendorfer, Wolfgang Reif

Angaben zur Veröffentlichung / Publication details:

Gajek, Carola, Alexander Schiendorfer, and Wolfgang Reif. 2020. "A chained neural network model for photovoltaic power forecast." *Lecture Notes in Computer Science* 11943: 566–78. https://doi.org/10.1007/978-3-030-37599-7_47.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



A Chained Neural Network Model for Photovoltaic Power Forecast

Carola Gajek^(✉), Alexander Schiendorfer, and Wolfgang Reif

Institute for Software & Systems Engineering,
University of Augsburg, Augsburg, Germany
`{gajek,schiendorfer,reif}@isse.de`

Abstract. Photovoltaic (PV) power forecasting is an important task preceding the scheduling of dispatchable power plants for the day-ahead market. Commercially available methods rely on conventional meteorological data and parameters to produce reliable predictions. These costs increase linearly with a rising number of plants. Recently, publicly available sources of free meteorological data have become available which allows for forecasting models based on machine learning, albeit offering heterogeneous data quality. We investigate a chained neural network model for PV power forecasting that takes into account varying data quality and follows the business requirement of frequently introducing new plants. This two-step model allows for easier integration of new plants in terms of manual efforts and achieves high-quality forecasts comparable to those of raw forecasting models from meteorological data.

Keywords: Machine learning · Neural networks · Photovoltaic power forecast

1 Motivation

In the wake of the energy revolution, more and more volatile power plants based on renewable energy sources such as wind turbines or photovoltaic (PV) plants enter the market. In Germany, for instance, solar energy accounted for 8.4% of the total electricity generated in 2018 – five years earlier it was only 5.7% [3]. This increasing ratio affects the stability of the power grid due to the intermittent generation of PV plants. Cloud movements influence the solar irradiation and consequently cause fluctuations in the generated power.

In order to guarantee balance in the energy grid, supply and demand must be approximately equal at all times. While the output generated by dispatchable plants such as gas turbines can be increased in the event of an energy decay, this is not possible for PV plants. So, if volatile plants produce too little electricity, dispatchable power plants have to compensate. In order to be able to estimate

The research presented in this paper was carried out in cooperation with and partially funded by Stadtwerke München. We thank Barnabas Kittlaus and Florian Nafz for their contributions and providing PV data.

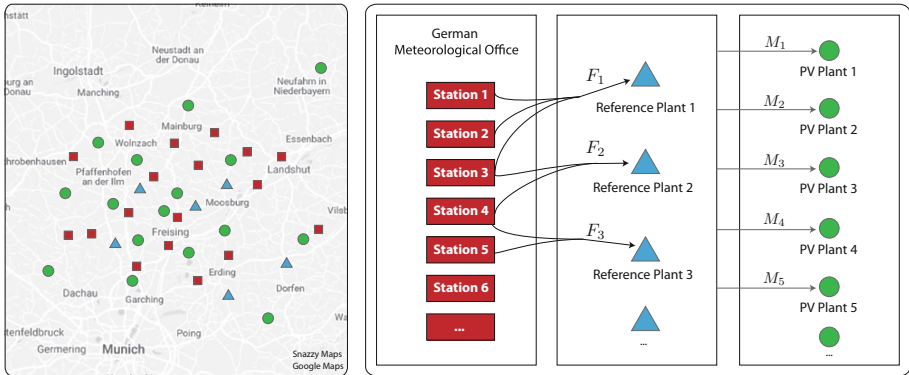


Fig. 1. Two-stage process by means of chained machine learning models: power outputs for reference PV plants (shown as blue triangles) are forecasted from raw meteorological data (at several locations, depicted as red rectangles) provided by the German Meteorological Office. The outputs of non-reference PV plants are then predicted from the output predictions of reference plants. Locations are picked for illustration purposes only and do not coincide with real plant or station locations. (Color figure online)

in advance when and how much energy will probably have to be compensated, forecasts are needed especially for volatile plants. The more accurate the power forecasts are, the less energy will have to be compensated in the short term, leading to more efficient schedules for the power plants involved [11]. Large energy operators responsible for a large number of PV systems (such as Stadtwerke München) have to create such forecasts for each individual plant, typically based on meteorological data. They can either be obtained commercially or must be created by the operators themselves based on physical equations or machine learning techniques. In order to be able to produce acceptable predictions in the latter case, usually a great amount of historical data is needed for each plant, which leads to high administrative expenses. If the underlying data consist of weather information from a local weather service, the costs of high-resolution weather information for the location of the power plant are usually quite high. Thus, costs and effort of data acquisition and management can quickly exceed the profit generated by accurate forecasting.

Consequently, in this paper, we take on the task of producing accurate PV forecasts under several business constraints for a specific application scenario:

Costs of Forecasts The driving motivation behind our study is to examine if accurate PV power predictions can be obtained from freely available meteorological data, as opposed to existing baseline predictions.

Heterogeneous Data Quality Apart from being free, the meteorological data available at relevant locations vary heavily in terms of quality (see Sect. 2).

Data Preparation and Maintenance Efforts New plants should be easily integrated into the system – optimally at little manual data cleansing costs.

Instead of trying to improve the state of the art in PV power forecasting in general, our goal is to apply machine learning to achieve acceptable power forecasts in this setting, even if the available meteorological data is limited. For doing so, we developed three prediction models: the first model introduced in Sect. 3.1 takes meteorological data from the region around a PV plant to forecast its power. Although this direct prediction model yields the best results in terms of error reduction, the data preparation effort is significant for every single plant due to the heterogeneous availability of weather data.

Therefore, we propose a chained model (Sect. 3.3) that includes precise predictions from weather data using the forecast model for a selected set of reference plants and learns mappings from the resulting predictions (Sect. 3.2) to every other remaining plant, as Fig. 1 visualizes. Our approach takes into account the organizational change processes involved with preparing and training new prediction models and offers a transformation strategy by starting with predictors for few reference plants to serve as a baseline for learned mappings and gradually converting them to have their own prediction models.

2 Heterogeneity of Data Quality

The communal energy service provider Stadtwerke München operates a *virtual power plant* which is a network of several small decentralized energy producers and consumers [12] to regulate short-term fluctuations in the energy grid. This coalition can jointly achieve production capacities matching those of ordinary power plants. In order to maintain the balance of the energy grid under the influence of many players, every provider has to forecast how much energy will be generated or consumed. For the so-called *day-ahead trading* at the energy market in Germany, these predictions are required daily in 15-minute resolution for the following day. Since unforeseen changes due to outages of some plants or changing weather conditions can quickly occur, *intraday trading* allows to compensate the deviating energy on the same day [5]. Due to their commercial relevance, we focus on day-ahead forecasts in this work.

We considered a subset of 120 PV plants of the virtual power plant throughout Germany, each of them with a maximum power between 100 and 20kW. Historical actual power output data was available in a 15-minute resolution from December 2017 through June 2018. As the virtual power plant continues to grow, several new PV plants can accrue per month, which also need to be predicted. Due to the direct dependence of photovoltaic power on weather condition, we used numerical weather records and forecasts from the German Meteorological Office DWD¹ as a basis for the power forecasts. The hourly forecasts originate from the weather model *cosmo-d2*, which is generated for an equidistant grid of 2.2km over Germany. By contrast, the historical data is recorded at various weather stations throughout Germany in different time resolutions. Most importantly, all of this data is distributed openly and for free which can in turn help to reduce prediction costs.

¹ Deutscher Wetterdienst <https://www.dwd.de>

The most natural way to obtain historical weather data for a PV plant would be to consider the records from the nearest weather station. Given the historical data from DWD, we face two problems: first, each station is only recording a small subset of all possible weather attributes, and second, multiple time gaps and series of error values occur in the recorded data – for some stations more than for others. This variation in data availability of the individual stations together with their uneven distribution over Germany lead to very heterogeneous data quality for all PV plants. For example, for some plants there could be a high-quality and close weather station, whereas others may be surrounded by many distant weather stations that simultaneously suffer from poor data quality or have recently been installed and therefore lack historical data.

Related Work

Due to their relevance for trading in the energy market, forecasting models for PV plants have been investigated thoroughly [1, 7, 11]. These models target several forecast horizons, ranging from a few seconds to days or weeks ahead. Similarly, they address a variety of spatial horizons, from a single site to regional forecasts. The authors of [1] provide a thorough survey of the most common approaches to PV forecasting, including analytical equations to model the irradiance of a PV system as the most important predictor for the power output (see, e.g., [4]) and statistical or machine learning models such as neural networks [8] or random forests [9]. In [8], a multilayer perceptron model is trained to forecast solar irradiance which is then converted into power outputs using a physical model. By contrast, our forecasting model (see Sect. 3.1) learns the mapping from irradiance values of multiple weather stations to a power output. This is due to the fact that we can obtain irradiance forecasts from DWD.

However, given our available (free) data, none of these forecasting models is directly applicable which is why we opted for a custom machine learning model that is trained in our specific setting and does not require, e.g., precise tilt angles and orientations of PV plants as inputs.

3 Forecasting Models

In our proposed approach (see Fig. 1), we build prediction models P_1, \dots, P_m for all m PV plants based on meteorological data from DWD. Since the generated forecasts are intended for day-ahead trading, they must be created for the following day in 15-minute resolution using current meteorological *predictions*. Consequently, estimating PV power is a regression problem where the input features \vec{x} correspond to a vector of numerical meteorological parameters for a certain point in time (given as forecasts themselves) and the target quantity y is the continuous output of a PV plant at this point in time. Although a first impression might suggest that a joint model could be trained and applied to forecast the outputs of all PV plants simultaneously, there are several reasons why we decided to build a prediction model for every single plant instead:

- (a) In a joint model, poor data quality of some PV plants could harm prediction accuracy for all other plants.

- (b) The joint model has to be completely re-built and re-trained when new plants are added to the virtual power plant. For separate models, by contrast, only the added plants require building new forecasting models and the already existing ones can simply remain. This aspect is very important in our scenario as new plants are added frequently to the virtual power plant.
- (c) Handling single models is scalable because they can be trained and evaluated in parallel.

3.1 A Power Forecasting Model from Weather Data

As mentioned before, we use publicly accessible weather data from DWD as basis for the power forecasting models F_i of a single PV plant i . Historical records of pairs of meteorological data and PV power outputs are used to train the models such that weather *forecasts* can afterwards be mapped to power *forecasts*. In order to use a weather attribute as a feature for the data set, it needs to occur both in historical and forecast data. Unfortunately, this is only the case for four attributes where we discarded the attribute `cloud coverage` due to its unavailability in more than 90% of the stations, leaving us with

- solar diffuse irradiation,
- solar global irradiation and
- air temperature two meters above ground

as our remaining available features. Fortunately, irradiation and temperature tend to be the most important features for PV power prediction [10, 14].

The historical weather values used to train our models are measured at fixed weather stations distributed over Germany in varying time resolutions (e.g. ranging from every minute to yearly), whereas the weather forecasts are collected from the weather model `cosmo-d2` that offers points arranged in a dense, equidistant grid over Germany (called `cosmo-d2` grid in the following). However, these forecast values are only available in an hourly resolution. Since the input data need to have the same resolution, the historical values are aggregated to hourly resolution to match the available forecast data. As mentioned before, we still have to provide power forecasts in *15-minute* resolution even though our weather features only have *hourly* resolution. Therefore, we map the hourly weather input onto four consecutive quarter-hourly power forecasts for the same hour.

Mapping coarse hourly weather condition onto finer quarter-hourly power forecasts presents us with a problem: the model cannot decide how the four output values should develop. For similar weather conditions, we would expect a positive development of them in the morning while the sun is still rising, but a negative one in the afternoon due to the sunset. Therefore, we add the hour value of the data point in a one-hot-encoded way to the features.

Because only few weather stations measure the three relevant weather attributes and the records contain many gaps or error values, we include the attributes of the stations in an area around a plant into its feature set. Based on the measured average speed of clouds from [6], we choose a radius of 90 km

Table 1. Statistical values describing the minimal, maximal, mean and median number of weather attributes from DWD data that were used as features for the forecasting models of the individual PV plants.

	Min	Max	Median	Mean
solar diffuse irradiation	0	7	2	2.23
solar global irradiation	0	7	2	2.77
air temperature two meters above ground	8	28	20	19.03
All three features	11	41	21	24.03

around the plant that corresponds to twice the average distance traveled by clouds per hour. Since an instance can only be used productively if there exist non-error values for all features, we consider a weather attribute from a station only if it provides enough data points so that the overall amount of instances is not reduced too much. With respect to the DWD data, we first sort the stations around a plant by their distance in ascending order. Afterwards, while iterating over the ordered stations, we exclude attributes from those with a ratio of missing or error values exceeding 1% of the current amount of instances.

According to this strategy, the weather-based data sets for all 120 PV systems have been generated. Table 1 shows that due to the heterogeneity of weather data, the data sets contain much more temperature attributes than for irradiation. Moreover, some PV plants even suffer from not having any weather station in their environment measuring irradiation values. Since these features are much more important for accurate power forecasts than temperature values, the corresponding forecasting models are expected to perform poorly.

For PV power forecasts, we have to resort to *cosmo-d2* grid points as opposed to specific weather stations provided in the historical data that were used for training. Therefore, we replaced the input features of the weather stations with those of the closest *cosmo-d2* grid points during forecasting. Since the grid resolution is about 2.2 km, the distance between the forecast points and weather stations is reasonable.

In an ideal world with lots of resources available for data processing and storage, the forecasting model is a good choice to get power forecasts for all PV plants, as our evaluation in Sect. 4 shows. With limited resources, however, a few problems arise: the historical weather records for each PV system must be both continuously updated and always processed into a consistent amount of data. In addition, the available weather stations of the DWD can change, for example, by adding new stations or removing previously used stations. In the latter case in particular, the model of the affected PV plant must then be completely re-built and re-trained due to the omission of the associated feature for this station. Thus, the entire knowledge of the old model is lost.

Besides our strategy of selecting feasible weather stations for a PV plant, there is a number of other strategies to create a joint data set from the given

historical records of the weather stations. Due to the heterogeneity of the available data, it may be possible that for different PV systems different merging strategies lead to the best data sets. Comparing multiple strategies for all plants to find the best one would in turn require a lot of computational or manual data preprocessing effort. The heterogeneity of the data poses another problem: As can be seen in Table 1, some considered PV plants do not have any nearby weather stations that measure irradiation attributes in sufficient quality. Unfortunately, these attributes are crucial for PV prediction. An intuitive solution would be to gradually increase the radius around the stations until an acceptable number of weather stations measuring irradiation values were found which in turn would cause increasing computational effort.

3.2 A Mapping Model Learning from Reference Power Forecasts

Instead of computing power forecasts using the weather-based forecasting model F_i for all m PV plants (including the aforementioned data preparation and cleaning efforts), we will build models F_{r_1}, \dots, F_{r_k} only for a small selection of plants, the so-called *reference plants* $\{r_1, \dots, r_k\} \subset \{1, \dots, m\}$, and map the resulting forecasts to the remaining $m - k$ plants via *mapping models* M_1, \dots, M_{m-k} . Precisely, a PV prediction model P_i of plant i is obtained as follows:

$$P_i = \begin{cases} F_i & \text{if } i \text{ is reference plant, i.e., } i \in \{r_1, \dots, r_k\} \\ M_i(F_{r_1}, \dots, F_{r_k}) & \text{if } i \text{ is non-reference plant} \end{cases}$$

This approach is based on the principle of regionality, as nearby plants are expected to be usually exposed to similar weather conditions and consequently show similar power curves. The power forecasts of the reference plants are mapped to those for the remaining non-reference plants. Various factors, such as the orientation of the system with respect to both the cardinal direction and the angle of the solar panels, their location along the latitude or the topology of the environment, for example tall shading buildings, affect the output power of a PV plant. Since this information is individual, in some cases not available, and in general difficult to express analytically, we also use machine learning techniques for the mapping model. To avoid having to manually pick a subset of reference plants for each PV plant, the forecasts of all of them are used as input for the models. A resulting advantage is that the location of the target plant does not have to be known in order to obtain forecasts.

The mapping model takes the forecast values of all reference plants for a certain quarter-hourly point in time to produce the power forecast of a regarded plant for the same point. For training, historical actual power outputs of the plant as well as the corresponding forecasts of the reference plants are required in 15-minute resolution. The reference plants should be selected from all over Germany and all together provide as gapless forecast values as possible for the considered period. This approach offers the advantage that all models have the same input, only the output has to be exchanged for each plant individually. Even if a new plant is added to the virtual power plant, there is no effort for generating new input for it.

3.3 A Chained Forecasting Model

To obtain reliable power predictions for a large and increasing number of PV plants based on the available data, we combine the two methods presented so far. This means that for every non-reference PV plant, the two models are composed: for the reference plants, the forecasting models F_{r_1}, \dots, F_{r_k} predict power values based on weather data; the mapping model M_i then takes all these power forecasts and estimates the output of a non-reference plant i , as shown in Fig. 1.

The *curse of dimensionality* is the most important reason why we decided to compose the models sequentially instead of combining them into one: In the latter case, the feature set would consist of all weather attributes of all weather stations associated with reference plants. Then, too few training instances could be available to train this model due to the time gaps and error values in the historical records. By splitting the model into two successive models, in contrast, significantly fewer features are used in both steps, which we expect to improve training for both and, consequently, to lead to better predictions.

Using this stepwise approach, a small loss of accuracy is tolerable in favor of a significantly lower manual data preparation effort. The models depend on much less historical weather data from DWD, since we only need them for the small portion of reference plants. Our strategy is to pick those few reference plants carefully by evaluating available raw weather data in terms of proximity and data quality. This expenditure is well spent since once we have high-quality forecasting models for the reference plants, we can easily train additional mapping models for the remaining plants using the same inputs (forecasts of reference PV plants) but only different power targets.

4 Evaluation

For our evaluation, we investigate whether our proposed models based on freely available weather data reach the quality of the commercial baseline predictions and whether the chained model achieves acceptable results compared to the pure forecasting models – at lower development costs.

We noted that the provided actual power values of some PV plants contained days with no power output at all. Since PV plants generate a small amount of energy even on winter days, this is obviously incorrect data and was consequently excluded from the data set. These zero power days could be caused by maintenance, defects or snow on the solar panels and would significantly skew the results. After this preprocessing step, the remaining data for each plant was split into a training set with about 70% of the instances, a validation and a test set with 15%, respectively. For a better visualization on the validation and test sets, all instances of a whole day were assigned to precisely one of the three sets.

Instead of optimizing hyperparameters individually for each of the 120 PV plants considered in this paper, we opted for a uniform hyperparameter setting across all plants. By doing so, we assume that the underlying functions to be learned are similar for all PV systems and therefore it is enough to tune the

capacity of a model only once for all of them. Therefore, a small subset of eleven plants was randomly selected such that they are distributed all over Germany with different maximal output, in the following referred to as *PV1* to *PV11*. We used the root mean squared error (RMSE) as error metric, adhering to conventions about PV forecasting evaluations in the literature [1].

As a preliminary experiment in terms of model adequacy, we evaluated the three machine learning techniques *random forest*, *Gaussian process* and *feed-forward neural network* for the mapping model using given baseline forecasts for 13 reference plants. Random forests have already been successfully applied to PV forecasting in literature [14] and can represent a large number of functions without having to severely limit the inductive bias as they make no prior assumptions regarding smoothness or other properties of the function. This is an important aspect for our scenario, as the form and the complexity of the mapping between the reference predictions and the performance values of the remaining plants is not known in advance.

Gaussian processes were chosen because of their probabilistic nature, which allows us to obtain predictions and simultaneously quantify their uncertainty. Moreover, once the covariance (or kernel) function has been selected, many hyperparameters of the model are set automatically during training so that they do not need to be chosen manually [13]. Like random forests, they can describe many different functions, which can nevertheless be restricted if necessary by the kernel function. Finally, neural networks have proven to be a universal tool in practice since they are very adaptive and can fit a wide variety of functions. They have also been successfully used for PV forecasting multiple times [1].

In our implementation, we used the python library *GPy*² for Gaussian processes, *scikit-learn*³ for random forests and *Keras*⁴ with the *TensorFlow*⁵ backend for neural networks. With random forests, the tuned hyperparameters included the **number of estimators**, the **minimum samples per leaf** and the **maximal number of features** per split. After model selection, i.e. grid searching for appropriate values of these hyperparameters, the number of estimators was set to the default value 10, the maximal number of features to 4 and the minimum samples per leaf was set to 1.

For Gaussian processes, we applied the procedure for automatic kernel selection described in [2]: At first, the best standard kernel of the framework is chosen. Afterwards, it is combined stepwise through addition or multiplication with other standard kernels until no further improvement of the resulting Gaussian process can be achieved. The best results were achieved using an additive combination of a linear and an exponential kernel for all input dimensions together after 200 optimization steps. For the neural networks, different architectures with one or two hidden layers were compared. The activation functions for these layers were taken from the standard functions of the framework, whereas the identity

² GPy <https://sheffielddml.github.io/GPy/>

³ scikit-learn <https://scikit-learn.org/>

⁴ Keras <https://keras.io/>

⁵ TensorFlow <https://www.tensorflow.org/>




Fig. 2. Comparison of validation errors (RMSE) of all tested methods for the mapping model and the baseline model.

function was selected for the output layer since we consider a multivariate regression problem. Models with one hidden layer containing ten neurons and the activation function `softplus` achieved the best results.

Figure 2 shows the error values achieved by the different models. All three techniques perform at approximately the same level for each plant, with the neural network performing best for most PV plants.

Compared to the mapping model, Gaussian processes are not equally suitable for our forecasting model due to the higher dimensionality of the weather data. The computational complexity of Gaussian processes depends cubically on the dimension [13] which makes this technique hard to apply in high-dimensional settings. Moreover, we also refrained from using random forests for our forecasting model since we expected the neural network to capture the feature interactions between the continuous weather attributes more easily. Hence, the experiments for the forecasting model were only performed with neural networks.

We used the same reference plants as in the preliminary experiments for the mapping model. Thus, no optimized selection of the reference plants has yet taken place. For each model, the number of features is different due to the heterogeneity of weather stations discussed in Sect. 3.1. For architecture tuning, we chose the number of hidden neurons as a multiple of the number of inputs. This makes the different models much more similar to each other than using a fixed number of neurons because the ratio of neurons in input and hidden layer stays the same. When tuning the forecasting models of the eleven selected power plants, three layers with the same number of neurons in the input as in the hidden layer and activation function `softplus` in the hidden layer turned out to be the most successful architecture.

Finally, for the chained model, we reused the techniques of the two basic models we identified as performing best. Whereas the weather-based forecasting models of the first part are completely adopted, the subsequent mapping models are re-tuned with the resulting reference forecasts since in the preliminary

Table 2. Evaluation of the direct forecasting model and the two-stage chained model on a shared test set in comparison with the baseline model of Stadtwerke München. Plants marked with an asterisk (*) are used as reference plants and therefore only predicted by the forecasting model. Absolute values are given as RMSE in kilowatts, relative values are error improvements compared to the baseline model.

PV plants	Chained Improvement		Forecasting Improvement		Baseline model
model			model		
<i>PV1*</i>			144.6	16.2%	172.6
<i>PV2</i>	37.0	41.2%	22.5	64.3%	62.9
<i>PV3*</i>			82.0	9.2%	90.3
<i>PV4</i>	181.8	8.0%	102.6	48.1%	197.6
<i>PV5</i>	108.7	5.2%	98.2	14.5%	114.8
<i>PV6</i>	54.5	10.6%	54.9	9.9%	61.0
<i>PV7</i>	189.3	−3.6%	146.4	19.9%	182.8
<i>PV8</i>	83.8	8.2%	71.1	22.2%	91.3
<i>PV9</i>	96.8	−0.2%	83.6	13.6%	96.7
<i>PV10</i>	199.9	44.9%	222.8	38.6%	362.8
<i>PV11</i>	110.3	14.5%	106.7	17.3%	129.0
Average	117.2	10.8%	104.9	20.2%	131.4
Sum	1288.7	10.8%	1153.4	20.2%	1445.7
Reference plants*					
Average			144.0	20.2%	180.5
Sum			1872.6	20.2%	2346.0
All 120 plants					
Average	112.0	20.0%	108.3	22.6%	139.9
Sum	13098.4	20.0%	12674.5	22.6%	16373.5

experiment they were trained on external forecast data. The resulting mapping model tuned on the selected eleven plants contains seven neurons in the hidden layer, once again applying the activation function `softplus`.

As Table 2 visualizes, both the direct forecasting and the chained model were able to achieve better results than the baseline predictions obtained commercially by Stadtwerke München. Using the forecasting model for predicting outputs of reference plants, the RMSE could be improved by 20.2%. Mapping these forecasts to the outputs of the remaining PV plants, this improvement could almost be retained. Regarding all 120 PV plants, the direct weather-based forecasting model performs better than the chained model – however, this comes at the price of a much higher manual workload, as we discussed in Sect. 3.

5 Conclusion and Future Work

We propose a chained model for PV power forecasting that takes into account the organizational change process of preparing and training new prediction models and follows requirements of business processes like easy adaption for accruing new plants. Since our models are based on meteorological data that is suffering from spatially varying data quality, we show how to deal with this heterogeneity. The evaluation shows that the chained model achieves better results than the baseline and even performs comparable to the direct forecasting model which requires significantly higher data preparation effort for every single plant.

In future work, one could use so-called *clear sky models* to interpolate hourly weather forecasts to quarter-hourly ones (like in [14]) for the forecasting model instead of mapping hourly weather forecasts to four consecutive quarter-hourly power forecasts. Moreover, transfer learning could help leverage knowledge extracted from plants with high data quality to improve models based on lower data quality. Finally, weather forecasts with short temporal horizon may serve as acceptable historical records since weather forecasts tend to be at good quality for this horizon.

References

1. Antonanzas, J., Osorio, N., Escobar, R., Urraca, R., Ascacibar, F.J., Antonanzas, F.: Review of photovoltaic power forecasting. *Sol. Energy* **136**, 78–111 (2016)
2. Duvenaud, D.: Automatic model construction with Gaussian processes. Ph.D. thesis, University of Cambridge (2014)
3. Fraunhofer ISE: Jährlicher Anteil der Solarenergie an der Stromerzeugung in Deutschland. https://www.energy-charts.de/ren_share_de.htm?year=all&source=solar-share&period=annual. Accessed 25 Mar 2019. (in German)
4. Inman, R.H., Pedro, H.T., Coimbra, C.F.: Solar forecasting methods for renewable energy integration. *Prog. Energy Combust. Sci.* **39**(6), 535–576 (2013)
5. Kraus, K.M.: Vortrag: “der regelenergiemarkt. In: Wirkungsweise und Wechselbeziehungen mit Spot-und Intra-Day-Märkten”, Regelenergie-Symposium, Leipzig September (2004)
6. Lappalainen, K., Valkealahti, S.: Analysis of shading periods caused by moving clouds. *Sol. Energy* **135**, 188–196 (2016)
7. Liu, J., Fang, W., Zhang, X., Yang, C.: An improved photovoltaic power forecasting model with the assistance of aerosol index data. *IEEE Trans. Sustain. Energy* **6**(2), 434–442 (2015)
8. Mellit, A., Pavan, A.M.: A 24-h forecast of solar irradiance using artificial neural network: application for performance prediction of a grid-connected pv plant at trieste, italy. *Sol. Energy* **84**(5), 807–821 (2010)
9. Mohammed, A.A., Yaqub, W., Aung, Z.: Probabilistic forecasting of solar power: an ensemble learning approach. In: Neves-Silva, R., Jain, L.C., Howlett, R.J. (eds.) *Intelligent Decision Technologies. SIST*, vol. 39, pp. 449–458. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-19857-6_38
10. Monteiro, C., Fernandez-Jimenez, L., Ramirez-Rosado, I.J., Muñoz Jiménez, A., Lara-Santillan, P.: Short-term forecasting models for photovoltaic plants: analytical versus soft-computing techniques. *Math. Problems Eng.* **2013**, 1–9 (2013)

11. Pierro, M., et al.: Multi-model ensemble for day ahead prediction of photovoltaic power generation. *Sol. Energy* **134**, 132–146 (2016)
12. Ruiz, N., Cobelo, I., Oyarzabal, J.: A direct load control model for virtual power plant management. *IEEE Trans. Power Syst.* **24**(2), 959–966 (2009)
13. Williams, C.K., Rasmussen, C.E.: *Gaussian Processes for Machine Learning*, vol. 2. MIT Press, Cambridge (2006)
14. Zamo, M., Mestre, O., Arbogast, P., Pannekoucke, O.: A benchmark of statistical regression methods for short-term forecasting of photovoltaic electricity production, part I: deterministic forecast of hourly production. *Sol. Energy* **105**, 792–803 (2014)