

Very Short Term Time-Series Forecasting of Solar Irradiance Without Exogenous Inputs*

Christian A. Hans and Elin Klages**

Control Systems Group, Technische Universität Berlin, Germany

Abstract. This paper compares different forecasting methods and models to predict average values of solar irradiance with a sampling time of 15 min over a prediction horizon of up to 3 h. The methods considered only require historic solar irradiance values, the current time and geographical location, i.e., no exogenous inputs are used. Nearest neighbor regression (NNR) and autoregressive integrated moving average (ARIMA) models are tested using different hyperparameters, e.g., the number of lags, or the size of the training data set, and data from different locations and seasons. The hyperparameters and their effect on the forecast quality are analyzed to identify properties which are likely to lead to good forecasts. Using these properties, a reduced search space is derived to identify good forecasting models much faster.

1 Introduction

In the last years, the capacity of globally installed photovoltaic generators has continuously increased [1]. With this growth, the intermittent nature of their infeed has become a major challenge in the operation of electric grids [2]. One important way to address this challenge is to use accurate forecasts to predict the infeed of photovoltaic power plants with a high temporal resolution [3].

Various studies on short term forecasting of photovoltaic power plant infeed [4–8] and solar irradiance [9–11] have been published. Autoregressive integrated moving average (ARIMA) models have been widely adopted in this context [4, 7, 9]. In the last decades forecasting methods that use techniques from artificial intelligence have become more prominent. Most of them use artificial neural networks [4, 8–11]. Others employ support vector [6, 8] and nearest neighbor regression [4, 5].

Despite additional effort that comes with the use of exogenous inputs, e.g., cloud cover or air temperature, only in [4, 7, 10] forecasts without exogenous inputs are considered. Moreover, in most publications [4–7, 10, 11] the forecasting models are obtained using only one data set. This makes it hard to draw general conclusions from the analysis that can be transferred to other locations or seasons of the year. Furthermore, in [4–6, 9–11] the model selection process is not explicitly discussed. Consequently, even though high forecast accuracies could be achieved for single data sets, due to the missing description of the selection processes, the findings cannot be directly used for different data sets. For example, in [11] the authors state that they used a search and in [9] that they tried different model structures to choose an artificial neural networks. However, they did not provide information on the search space. In [6], a search space for support vector regression forecasts is provided but no information on the selection criterion or the final model structure and parameters is given. The structure of an artificial neural networks forecasting model is examined in [10] using a sensitivity analysis. Unfortunately, the publication does not contain sufficient details to replicate the analysis for different data. To the knowledge of the authors only in [5] a search space is provided and the results are analyzed to gain information about suitable model structures. However, the analysis is only performed on data from one location. This makes it hard to draw conclusions that allow identify a reduced search space and in practice often leads to an exhaustive search for suitable models.

In this paper we derive a significantly reduced search space that exhibits a high probability of finding an accurate forecasting model. As potential forecasting methods, ARIMA and nearest neighbor regression (NNR) were considered due to their wide adoption. Furthermore, their hyperparameters, e.g., the number of autoregressive lags or neighbors, are mostly discrete which allows

* This work was partially supported by the German Federal Ministry for Economic Affairs and Energy (BMWi), Project No. 0324024A.

** Currently at the Institute of Modelling and Computation, Technische Universität Hamburg, Germany.

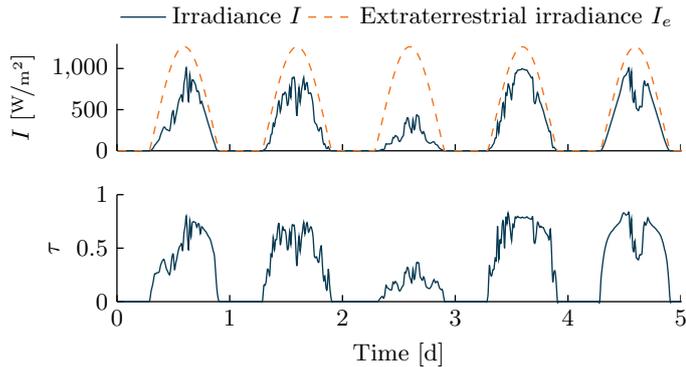


Fig. 1: Global horizontal irradiance, extraterrestrial irradiance and transmissivity over a duration of 5 d. The irradiance was taken from [14], the extraterrestrial irradiance was estimated using [15].

to form a discrete search space where points in a certain range can be explored. For each point, i.e., for each set of hyperparameters in this search space, forecasting models are trained and tested using data from different locations and seasons. Based on the resulting forecast accuracies, conclusions for a reduced search space are drawn. These include the choice of NNR over ARIMA, the number of training data points and autoregressive lags as well as handling of night data and use of transmissivity instead of irradiance. The selection process is explained in much detail allowing others to perform a similar search on different data.

For simplicity and robustness, this work focuses on forecasts of solar irradiance without exogenous inputs using only historic irradiance data, time and location. Motivated by the smallest intraday interval of energy trading in Germany a 15 min prediction step is chosen [12]. To cover a full charge or discharge of a medium size storage unit, a prediction horizon of 3 h is considered.

The remainder of this paper is organized as follows. Solar irradiance and data preprocessing are discussed in Section 2. In Section 3, basics on time-series forecasting are discussed. The forecasting methods considered in this work are presented in Section 4. Then, the model selection procedure is illustrated in Section 5. Finally, in Section 6 the results of the hyperparameters search are analyzed.

1.1 Preliminaries

Throughout this paper, real numbers are denoted by \mathbb{R} , nonnegative real numbers by \mathbb{R}_0^+ , positive real numbers by \mathbb{R}^+ , natural numbers by \mathbb{N} and nonnegative integers by \mathbb{N}_0 . The transpose of a matrix x is x^T and the Euclidean norm of a vector x is $\|x\|_2$. The sum over all elements of a set $\mathbb{K} \subset \mathbb{N}$ where each element $i \in \mathbb{K}$ is taken exactly once is denoted $\sum_{i \in \mathbb{K}} x_i$.

2 Solar Irradiance

The sunlight reaching the outer earth's atmosphere is called extraterrestrial solar radiation. It can be estimated from the energy emitted by and the position of the sun. The sunlight reaching a horizontal plane on earth per unit area at time t , is called global horizontal irradiance $I_t \in \mathbb{R}_0^+$. It includes direct normal irradiance, which originates directly from the sun, and diffuse irradiance, which includes scattered and ground reflected components, [13].

For forecasting, it can be beneficial to normalize the solar irradiance I_t by the extraterrestrial irradiance $I_{e,t} \in \mathbb{R}_0^+$ (see Fig. 1). This yields to transmissivity defined as $\tau_t = I_t/I_{e,t}$ [8]. In this work, the solar position algorithm from [15] is used to estimate the extraterrestrial irradiance $I_e(t)$. It only requires the position of the respective surface to calculate the zenith angle $\zeta_t \in [0, 2\pi)$, i.e., the incidence angle of the sun light on a horizontal plane on the earth's surface [16]. Using ζ_t , the extraterrestrial irradiance can be determined by $I_{e,t} = \epsilon_t I_s \cos(\zeta_t)$, where ϵ_t is a correction factor and $I_s = 1360.8 \text{ W/m}^2$ is the solar constant.

3 Time-Series Forecasting

A time series is a collection of $n \in \mathbb{N}$ chronologically ordered observations y_1, y_2, \dots, y_n . In this collection, every element y_t refers to an observation performed at a time instant $t = 1, \dots, n$. In this work, we focus on univariate forecasts, i.e., forecasts using only historic irradiance to forecast future irradiance. Thus, all elements $y_t \in \mathbb{R}$ for $t = 1, \dots, n$ are real-valued scalars.

Broadly speaking, a forecasting method is a procedure to estimate future values $\hat{y}_{t+j|t}$ of a time series based on past values $Y_t \in \mathbb{R}^n$, i.e., $\hat{y}_{t+j|t} = f(Y_t)$. Here, $\hat{y}_{t+j|t}$ refers to the value at prediction step $j \in \mathbb{N}$ with last known value at time $t \in \mathbb{N}$. In case of univariate forecasting, Y_t only includes the present and past values of the time series being forecast, i.e., $Y_t = [y_t \ y_{t-1} \ \dots \ y_{t-n+1}]^T$.

The vector Y_t can also be based on non-consecutive lags, e.g., to represent seasonal behavior. An example for such a vector is $Y_t = [y_t \ y_{t-s+1} \ y_{t-s}]^T$, where $s \in \mathbb{N}$ is the length of the season, e.g., one day. In this example, the present observation, y_t , the observation from the last day, y_{t-s+1} , at the same time as the predicted value, \hat{y}_{t+1} , and from the current time minus 24 h, y_{t-s} , are included.

3.1 Multi-Step Forecasts

To create multi-step ahead forecasts, a recursive strategy can be applied [17, Section 4.]. The forecasting model is repeatedly used for one step forecasts, where the elements of Y_t are adapted at each step until prediction step $J \in \mathbb{N}$ is reached. For example, if Y_t includes the last three data points, $Y_t = [y_t \ y_{t-1} \ y_{t-2}]^T$ then the procedure for $J = 12$ prediction steps would be

$$\hat{y}_{t+1|t} = f(Y_t) = f([y_t \ y_{t-1} \ y_{t-2}]^T), \quad (1a)$$

$$\hat{y}_{t+2|t} = f(\hat{Y}_{t+1|t}) = f([\hat{y}_{t+1|t} \ y_t \ y_{t-1}]^T), \quad (1b)$$

⋮

$$\hat{y}_{t+12|t} = f(\hat{Y}_{t+11|t}) = f([\hat{y}_{t+11|t} \ \hat{y}_{t+10|t} \ \hat{y}_{t+9|t}]^T). \quad (1c)$$

For seasonal models, the vector $\hat{Y}_{t+j|t} \in \mathbb{R}^n$ is equally adapted at each step in a similar manner.

3.2 Evaluation of Forecast Accuracy

Handling of Night Data. Forecasts of data points during the night, i.e., of zero irradiance values, were not included in the evaluation of the forecast accuracy. Therefore, observations with a zenith angle $\zeta_t \leq 90.83^\circ$ were excluded based on [18]. The angle ζ_t was taken from the data sets if available. Otherwise it was estimated using [15].

Root Mean Square Error (RMSE). For each prediction step $j = 1, \dots, J$ over prediction horizon J , the RMSE is calculated to analyze the error of each step separately. Assuming that $m \in \mathbb{N}$ forecasts were performed, the RMSE of prediction step j is

$$\text{RMSE}_j = \sqrt{\frac{1}{m} \sum_{t=1}^m (\hat{y}_{t+j|t} - y_{t+j})^2}. \quad (2)$$

4 Forecasting Methods

This section presents the forecasting methods employed in this work. First, the persistence model used as reference is discussed. Then, ARIMA models and nearest neighbor regression are illustrated.

4.1 Persistence Models

The basic idea of the persistence model is that future values are assumed to be equal to known past values. In this work, the simple model, $\hat{y}_{t+1|t} = y_t$ is used. It is also possible to consider a seasonal persistence model using, e.g., data from the previous day, week or year. This can be described by $\hat{y}_{t+j|t} = y_{t+j-s}$ where s is the seasonal period.

4.2 Autoregressive Integrated Moving Average (ARIMA) Models

ARIMA models are widely used for time-series forecasting. Future values are estimated using a linear combination of previously observed values and forecasting errors. Also, differencing can be applied to obtain stationary data. An ARIMA model can be described by [19]

$$\Phi(B)\nabla^d y_t = \theta_0 + \Theta(B)e_t, \quad (3)$$

where B is the backwards shift operator with $B^m y_t = y_{t-m}$, and ∇ is the backwards difference with $\nabla^d y_t = (y_t - y_{t-1})^d$, $d \in \mathbb{N}_0$. Furthermore, $\Phi(B) = 1 - \phi_1 B^1 - \dots - \phi_p B^p$ with $\phi_1, \dots, \phi_p \in \mathbb{R}$, $p \in \mathbb{N}_0$ is the autoregressive part, and $\Theta(B) = 1 - \theta_1 B^1 - \dots - \theta_q B^q$ with $\theta_1, \dots, \theta_q \in \mathbb{R}$, $q \in \mathbb{N}_0$ the moving average part. Moreover, e_t is the difference between measured value and forecast.

To fit a model to data that exhibits seasonality, ARIMA models can be extended by seasonality of period s . These, so called seasonal ARIMA (SARIMA) models, can be described by

$$\Phi(B)\Phi(B)\nabla^d \nabla_s^D y_t = \theta_0 + \Theta(B)\Theta(B)e_t, \quad (4)$$

where $\Phi(B) = 1 - \Phi_1 B^s - \dots - \Phi_P B^{P s}$ with $\Phi_1, \dots, \Phi_P \in \mathbb{R}$, $P \in \mathbb{N}_0$ is the seasonal autoregressive part and $\Theta(B) = 1 - \Theta_1 B^s - \dots - \Theta_Q B^{Q s}$ with $\Theta_1, \dots, \Theta_Q \in \mathbb{R}$, $Q \in \mathbb{N}_0$ the seasonal moving average part. Moreover, $\nabla_s^D y_t = (y_t - y_{t-s})^D$, $D \in \mathbb{N}_0$ represents seasonal differencing.

For the implementation of the ARIMA models, the econometrics toolbox in MATLAB 2015b was used. Within the toolbox, maximum likelihood estimation is used to find the model parameters.

4.3 Nearest Neighbor Regression (NNR)

In what follows, based on [20] we introduce NNR. In NNR the pattern of historic data at time instant t , $Y_t \in \mathbb{R}^n$, is compared to previous patterns $Y_i \in \mathbb{R}^n$, $i = 0, \dots, N-1$ in the reference sample $\mathbb{D} = \{(Y_0, y_1), (Y_1, y_2), \dots, (Y_{N-1}, y_N)\}$, $N \in \mathbb{N}$ to forecast \hat{y}_t .

Although NNR allows to use arbitrary elements in Y_i , e.g., $Y_i = [y_{i-13} \ y_{i-111} \ y_{i-2678}]^T$, we follow the notion of autoregressive and seasonal autoregressive lags of (seasonal) ARIMA models to enable a comparison of both methods. For NNR this means that according to the number autoregressive lags p the vectors Y_i in \mathbb{D} are formed. Thus, analog to ARIMA models the entries of the NNR reference sample with $p = 3$ autoregressive lags has the form $Y_i = [y_i \ y_{i-1} \ y_{i-2}]^T$.

The same holds for the seasonal autoregressive lags. In a similar fashion as in (4), the seasonal autoregressive part is multiplied with the autoregressive part. Consequently, a reference sample with $p = 3$ autoregressive lags, $P = 2$ seasonal autoregressive lags and a season of $s = 10$ results in

$$Y_i = [y_i \ y_{i-1} \ y_{i-2} \mid y_{i-9} \ y_{i-10} \ y_{i-11} \ y_{i-12} \mid y_{i-19} \ y_{i-20} \ y_{i-21} \ y_{i-22}]^T.$$

Thus, as in (4), the elements of the previous seasons that are associated with $\hat{y}_{i+1|i}$ (here y_{i-9} and y_{i-19}) are included in Y_i . This is due to the zero order term, i.e., the 1, in $\Phi(B)$ and $\Theta(B)$.

A simple forecasting model can be obtained by combining the $k \in \mathbb{N}$ elements in \mathbb{D} with the smallest distance to Y_t . Using the set of k nearest neighbors of Y_t , $\mathbb{K}(Y_t) \subseteq \mathbb{D}$, this can be written as

$$\hat{y}_{t+1|t} = f(Y_t, \mathbb{D}) = 1/k \sum_{Y_i \in \mathbb{K}(Y_t)} y_{i+1}. \quad (5)$$

Note that the number of neighbors k can be fixed or determined by a maximum distance $\varepsilon \in \mathbb{R}^+$. With distance $d(Y_i, Y_t)$, the set of neighbors closer than ε to Y_t is $\mathbb{K}(Y_t) = \{Y_i \in \mathbb{D} \mid d(Y_i, Y_t) \leq \varepsilon\}$. Note that in this work, the Euclidean norm is used as distance, i.e., $d(Y_i, Y_t) = \|Y_i - Y_t\|_2$.

The model (5), can be modified using a weighted average. In this work, weights inverse to the distance between x_t and the neighbor x_i , have been considered, i.e.,

$$\hat{y}_{t+1|t} = f(Y_t, \mathbb{D}) = \begin{cases} 1/k \sum_{Y_i \in \mathbb{K}(Y_t)} \frac{1}{d(Y_i, Y_t)} y_{i+1}, & \text{if } d(Y_i, Y_t) \neq 0 \ \forall Y_i \in \mathbb{K}(Y_t), \\ y_{t+1}, & \text{if } \exists Y_l \in \mathbb{K}(Y_t) \text{ with } d(Y_t, Y_l) = 0. \end{cases} \quad (6)$$

Although, (5) requires no training, the model and the selected reference sample will be referred to as trained model for simplicity. For implementation, the Statistical Learning Toolbox for MATLAB [21] was used. Note that in this work, approximated nearest neighbors which is faster than the exact nearest neighbors algorithm is used as no significant difference in the forecast accuracy between the two could be observed.

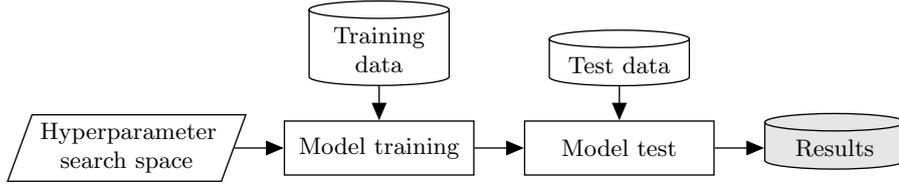


Fig. 2: Scheme used to obtain forecast accuracy data for different models.

5 Hyperparameter Search

We aim to compare the forecast accuracy for different forecasting model structures and different sets of historic data. Therefore, an exhaustive search was performed to evaluate what models approximate the test data best. In our search, different model structures are trained using different sets of historic data. To identify model structure and historic data, hyperparameters are used. Moving along Fig. 2, in this section first training and test data are discussed. Then, the hyperparameter search space is posed. Finally, training and test of the models is illustrated.

5.1 Training and Test Data

For training and test data sets [14, 22, 23] were used. From each data set three different seasons were selected to include different climatic situations in the hyperparameter search. For every location and every season, i.e., for 9 data sets, the observations were divided into training data (up to 2 months) and independent test data (1 week) that was not used for training. The test data was chosen to always follow directly the data used for training. As shown in Fig. 3, the nine weeks of test data include different climatic situations, e.g., a sunny week in December from the National Renewable Energy Laboratory (NREL) Clark Station, a cloudy week in March from the Atmospheric Radiation Measurement (ARM) Facility, and various weeks with sunny, foggy and cloudy days.

Note that all data sets [14, 22, 23] have a temporal resolution of 1 min. The 15 min average values y_1, \dots, y_N were derived from the original time series y'_1, \dots, y'_{15N} as $y_t = 1/15 \sum_{q=1}^{15} y'_{q+15(t-1)}$.

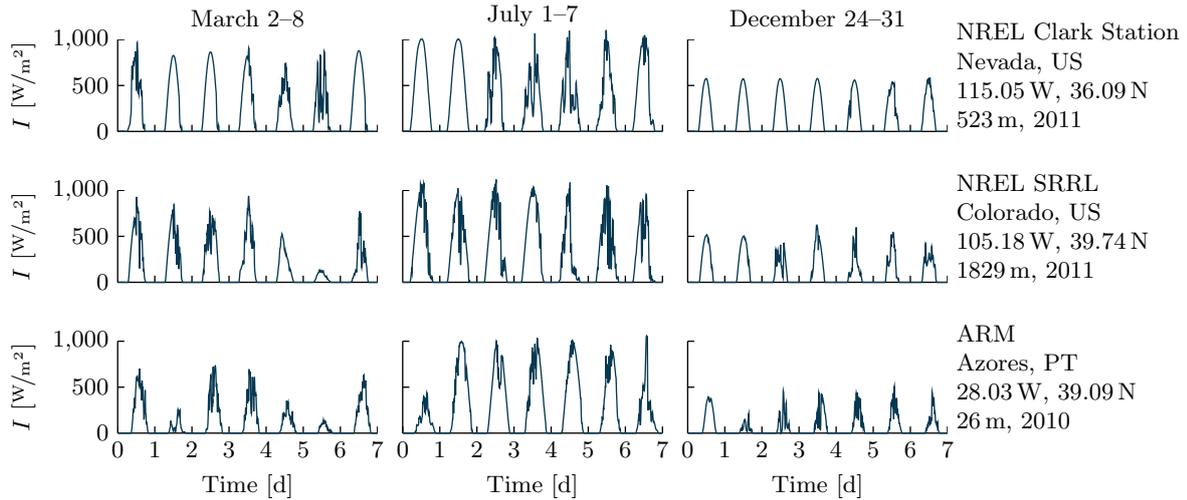


Fig. 3: Measured solar irradiance data used to test the trained models.

Table 1: Sets of hyperparameters for different forecasting methods considered in search.

(a) ARIMA models.		(b) NNR.	
Autoregressive lags p	0, 1, ..., 10	Autoregressive lags p	1, 2, ..., 20
Moving average lags q	0, 1, ..., 10	Weight	uniform, inverse to distance
Differencing d	0, 1, 2	Number of neighbors k	1, 2, ..., 20
		Threshold ³ ε	0.01, 0.05, 0.1, 0.5, 1
(c) SARIMA models.		(d) SNNR.	
Autoregressive lags p	0, 1, 3	Autoregressive lags p	1, 2, ..., 11
Moving average lags q	0, 1, 3	Seasonal autoregressive lags P	1, 2, ..., 7
Seasonal autoregressive lags P	0, 1, 2, 3	Weight	uniform, inverse to distance
Seasonal moving average lags Q	0, 1, 2, 3	Number of neighbors k	1, 2, ..., 20
Differencing d	0, 1	Threshold ³ ε	0.01, 0.05, 0.1, 0.5, 1
Seasonal differencing D	0, 1		

5.2 Hyperparameters

Data Hyperparameters. The hyperparameters that concern the training data are independent of the forecasting method and the model structure. We considered modifications of training data by 1. pre- and postprocessing, 2. handling of night data and 3. number of data points used.

1. Regarding the pre- and postprocessing, two approaches were considered. (a) The time series of irradiance, I_t , is directly used to train the models without further preprocessing. (b) Transmissivity $\tau_t = I_t/I_{e,t}$ is derived from I_t and extraterrestrial irradiance $I_{e,t}$. Then, the models are trained to forecast $\hat{\tau}_{t+j|t}$. Using $\hat{\tau}_{t+j|t}$ the irradiance forecast is then calculated via $\hat{I}_{t+j|t} = I_{e,t} \hat{\tau}_{t+j|t}$.

2. Three cases related to the handling of data points at night were considered: (a) include all data points from day and night, (b) only include data points between 5 am and 8 pm local time, and (c) only include data points between sunrise and sunset, i.e., observations with $\zeta_t > 90.83^\circ$ (see Section 3.2). As it can be cumbersome to implement models with seasonality for a varying number of data points per day, case (3) was not considered for seasonal models.

3. Regarding the number of training data points, 1, 3, 7, 14, 30, and 60 days of data were considered. For seasonal models, the data was chosen to include at least 14 days.

Model Hyperparameters. For every forecasting method, the model structure can be uniquely described using model specific hyperparameters. Note the difference to model parameters of a trained model. For example, the number of autoregressive lags p is a hyperparameter. The coefficients of the autoregressive part, ϕ_1, \dots, ϕ_p , of a trained ARIMA model are model parameters. The subsets of hyperparameters considered in the search for are shown in Table 1. Note that in Tables 1(b) and 1(d) either the number of neighbors k or the maximal distance to the neighbors ε is used. Further note that the season s was chosen to be 1 d for SARIMA and seasonal NNR (SNNR).

5.3 Model Training

For each combination of hyperparameters described in Section 5.2, we trained nine models, i.e., one for each location and each season (see Section 5.1). This led to more than 250,000 trained forecasting models with different hyperparameters. Note that some combinations of hyperparameters did not result in suitable models. Especially for ARIMA, no stable model could be derived for certain hyperparameters. Furthermore, for SARIMA 3.3% of the hyperparameters were excluded, mostly due to very long training procedures that rendered them unsuitable for practical use.

³ Multiplied by 1360.8 W/m^2 for irradiance.

5.4 Model Test

In the model test, the performance of the trained models was evaluated. Therefore, every trained model was tested by forecasting irradiance of the unknown test data (see Fig. 3) that followed the training data. Due to zero irradiance, data points during the night were excluded from the evaluation based on the zenith angle ζ_t (see Section 3.2). For each of the data point in the test data, a 12 step ahead, i.e., 3 h, prediction was obtained. Using the forecast values and the test data, the RMSE_j (see Section 3.2) was then calculated for $j = 1, \dots, 12$ and stored. In the next section, we will analyze the resulting RMSE_j for every combination of hyperparameters, location and season to identify models that are likely to provide good forecasts.

6 Analysis

In this section, we analyze the results of the hyperparameter search in Section 5. Our goal is to draw conclusions that help to reduce the space of future hyperparameter searches. More precisely, we aim to find forecasting methods and ranges of hyperparameters that increase the probability of finding a suitable model for unseen data. Before starting the analysis, some remarks are posed.

Remark 1. As reference, we use the persistence model on transmissivity as it outperformed the persistence model on irradiance for all test data sets. In the reference model historic transmissivity is derived using $\tau_t = I_t/I_{e,t}$. Then a persistence forecast is performed, i.e., we set $\hat{\tau}_{t+j|t} = \tau_t$ for $j = 1, \dots, J$. Finally, the irradiance forecast is determined as $\hat{I}_{t+j|t} = \hat{\tau}_{t+j|t} I_{e,t+j}$ for $j = 1, \dots, J$.

Remark 2. The style of the box plots used throughout this work is illustrated in in Fig. 4. Here, m marks the median. The box around the median contains all data from the 25th (q_1) to the 75th (q_3) percentile. The left whisker marks the lowest occurring value within $q_1 - 1.5(q_3 - q_1)$ and the right whisker marks the highest occurring value within $q_3 + 1.5(q_3 - q_1)$. Due to numerous outliers, only every 100th outlier is shown. To assure that the forecast with the highest accuracy is considered the lowest outlier is always included.

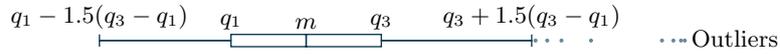


Fig. 4: Box plot with median m , quartiles q_1 and q_3 , whiskers and outliers.

Remark 3. Most distributions of the forecast accuracy discussed in this section are similar for the different test data sets shown in Fig. 3. Therefore, most distributions were combined in one plot that includes the values of all locations.

Remark 4. Throughout this paper, the forecast at prediction steps $j = 1, 4, 12$ are analyzed to approximate the evolution of the forecast accuracy over prediction horizon $J = 12$. For the 15 min sampling time considered, they correspond to a forecast 15 min, 1 h and 3 h ahead.

6.1 Forecasting Method

We first compare the forecasting methods ARIMA and NNR along with the persistence model on transmissivity that acts as a reference (see Remark 1). In Fig. 5, box plots illustrating the distributions of forecast accuracies of all models included in the search are shown.

The box plots in Fig. 5 show that for predictions 15 min ahead, almost no model outperforms the persistence model. For a prediction step of 1 h some models outperform the persistence forecaster. Furthermore, the RMSE of forecasts performed with NNR varies less than the RMSE of ARIMA forecasts. This indicates a higher probability of finding a good forecasting model using NNR. For

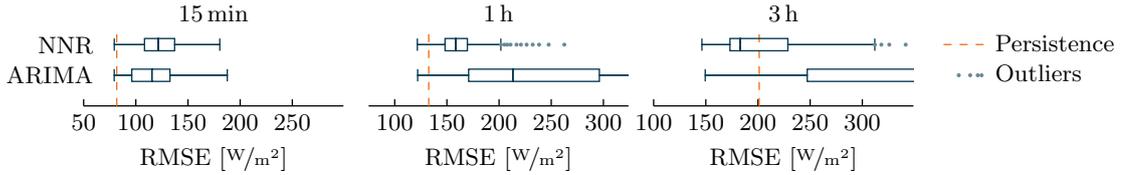


Fig. 5: RMSE for 15 min, 1 h and 3 h ahead predictions with ARIMA and NNR models.

predictions of 3 h ahead, the potential improvement over the persistence model is significant for all methods. As already observed for 1 h ahead predictions, the RMSE of NNR varies less. Additionally, the 25th percentile is remarkably lower for NNR. The highest forecast accuracy achieved is very alike for NNR and ARIMA for all prediction steps. Still, NNR achieves much lower medians and lower 25th and 75th percentiles. Consequently, the probability of finding a suitable NNR model is much higher. Therefore, ARIMA and SARIMA models will be disregarded in the next steps of the analysis. Consequently, in what follows we only investigate how the search spaces for NNR and SNNR can be reduced to increase the probability of finding an accurate forecasting model.

6.2 Preprocessing of NNR and SNNR

As stated in Section 5.2, models using irradiance or the normalized transmissivity values were considered in the search. In Fig. 6 the forecast accuracy of both approaches is shown. It can be observed that the box plots for 15 min forecasts are very alike. Also, little difference can be seen for 1 h ahead predictions. However, for 3 h ahead predictions the transmissivity based forecasts often outperform the irradiance based forecasts. The only exception from this observation could be found for data from the NREL Clark Station during March 2–8 (see Fig. 7), where the best models use irradiance data. Still the lower whiskers and the medians in this case are very alike.

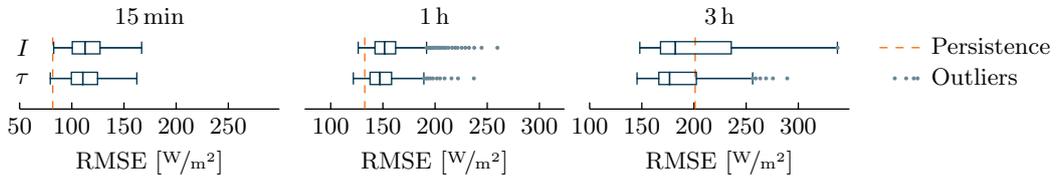


Fig. 6: RMSE of NNR and SNNR using different preprocessing for all locations and seasons.

The box plots indicate that using transmissivity yields a higher probability of finding an accurate model. Although for data set [22] a better irradiance model for the 3 h prediction could be found, the overall distribution in this case shows no significant additional difference to Fig. 6. As the single best model is hard to find, we are confident that using transmissivity yields a higher probability of finding a good model. Therefore, irradiance forecasts are excluded from the reduced search space.

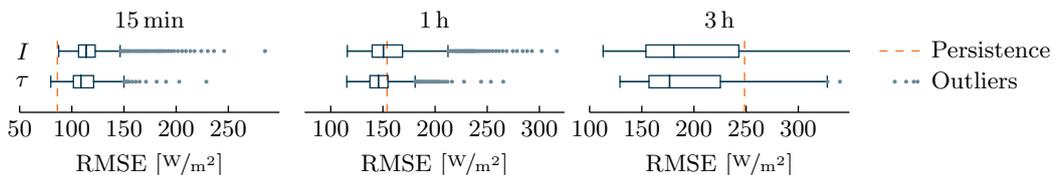


Fig. 7: RMSE of NNR and SNNR using different preprocessing at NREL Clark Station, March 2–8.

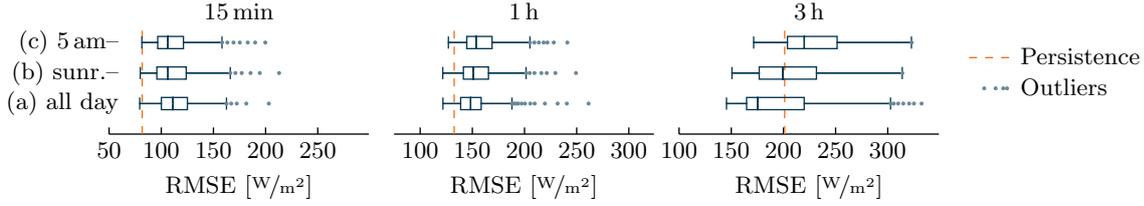


Fig. 8: RMSE of NNR and SNNR for different handling of night data: (a) entire day, (b) sunrise to sunset and (c) 5 am to 8 pm.

6.3 Handling of Night Data of NNR and SNNR

As stated in Section 5.4, data points during the night were not included in the evaluation of the forecast accuracy. However, in some cases they were used in the data of the reference sample to allow for smooth transitions between two days. As stated in Section 5.1, we considered training data that (a) includes the data of the entire day (and night), (b) only includes data points from sunrise to sunset and (c) only includes data points from 5 am to 8 pm.

In Fig. 8, the forecast accuracy for different handling of night data is shown. It can be observed that the forecast accuracy is similar for the 15 min forecast. For forecasts 1 h ahead, the difference between the approaches remains small. However, for forecasts 3 h ahead, using all data points for training, i.e., including night values, outperforms the other approaches. Therefore, all data points including night data are considered in the reduced search space.

6.4 Size of the Reference Sample of NNR and SNNR

The reference sample is chosen such that it includes that last values of historic data up to the most recent data point. A large reference sample, i.e., a set \mathbb{D} with a high number of elements, therefore includes data that reaches further into that past than a small reference sample.

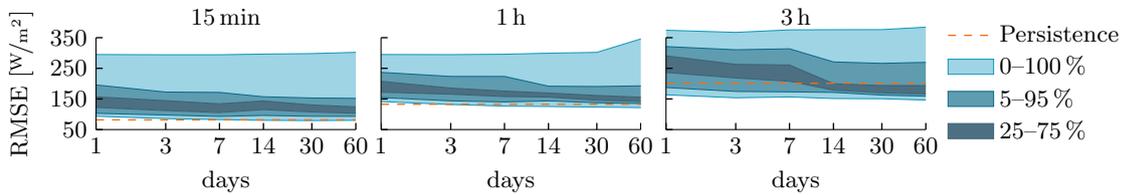


Fig. 9: RMSE of NNR and SNNR using different sizes of the reference sample.

In Fig. 9, the forecast accuracy using a reference sample of 1 d, 3 d, and 7 d as well as two weeks (14 d), three weeks (21 d), one month (30 d) and two months (60 d) is shown. The plots for all prediction steps shows a small decrease of the 25th and the 75th percentile as well as the lowest RMSE for an increasing number of elements in the reference sample. Considering the largest RMSE of each prediction step, the inverse effect can be observed. In general little difference can be observed in the forecast accuracy, when varying the size of the reference sample. However, as the forecast accuracy increases using more data, a reference sample of 60 d is used in the reduced search space.

6.5 Autoregressive (AR) Lags of NNR and SNNR

This section focuses on the AR lags, i.e., the data points of the historic data used in the elements of the reference sample. We differentiate between AR and seasonal AR lags. The former refer to 15 min steps prior to the first forecast time instant. The latter refer to 24 h steps, used to include

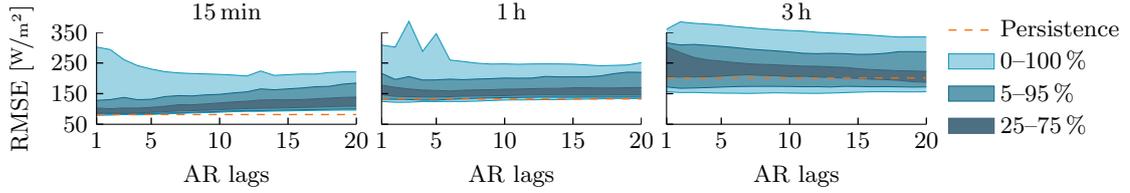


Fig. 10: RMSE of NNR using different AR lags.

the days prior to the forecast. It can be observed in Fig. 10 that for all prediction steps the models with the highest forecast accuracy has a small number of AR lags. However, no clear tendency can be identified that supports choosing a particular number of AR lags.

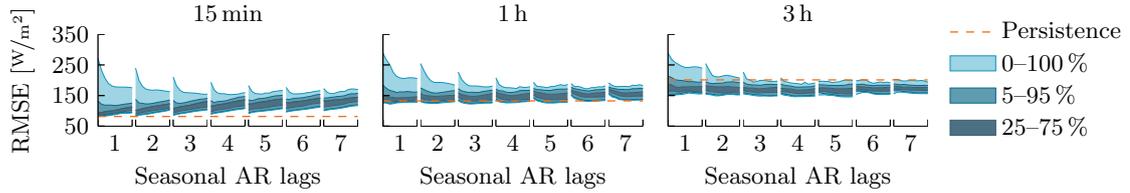


Fig. 11: RMSE of SNNR using different seasonal AR lags. For every seasonal AR lag, the non-seasonal AR lags increase from 1 on the left side of the plot to 11 on the right side of the small subplot.

Fig. 11 shows the RMSE for different numbers of seasonal AR lags. Here, for every seasonal AR lag, the number of AR lags was increased from 1 (left) to 11 (right). It can be observed that for 15 min forecasts the RMSE slightly increases with the number of seasonal AR lags. Furthermore, for every seasonal AR lag, the RMSE also increases with increasing number of non-seasonal AR lags. In contrast, for the 3 h forecasts, the RMSE slightly decreases with the number of seasonal lags. Comparing Figs. 10 and 11, shows that the variance decreases for the 1 h and 3 h forecasts with increasing number of seasonal lags. This results in a higher probability of finding a good model for a higher number of seasonal AR lags. Unfortunately, the 15 min forecasts do not show the same effect. Still, as the 15 min forecast barely outperform the persistence model it seems reasonable focus on the improvement for larger horizons and consider only seasonal models in the reduced search space.

6.6 Weights of NNR and SNNR

As stated in Section 5.2 uniform weights and weights inverse to the distance to the neighbors were considered. In the following, we analyze the effect of the distance to neighbors on the distribution of the RMSE. As can be seen in Fig. 12, the box plots hardly differ for all prediction steps. This indicates that there is no strong correlation between forecast accuracy and weights. However, because of an easier implementation, uniform weights were chosen for the reduced search space.

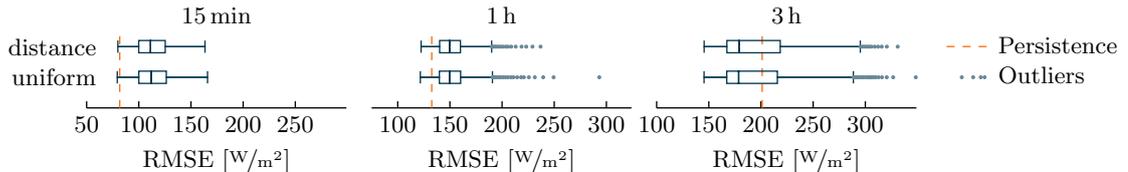


Fig. 12: RMSE using uniform weights and weights inverse to the distance of the neighbor.

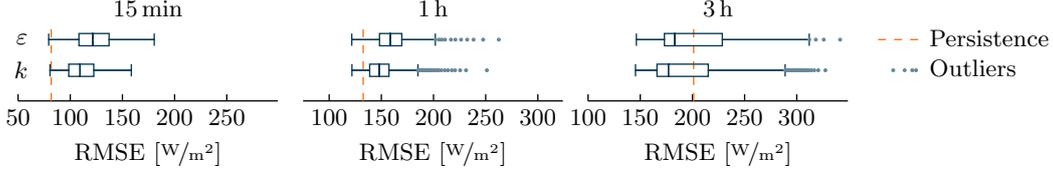


Fig. 13: RMSE using a threshold ε or a fixed number of neighbors k to define the neighborhood \mathbb{K} .

6.7 Definition of the Neighborhood of NNR and SNNR

As stated in Section 4.3, the neighborhood can be chosen to be a fixed number of neighbors, k , or a maximum distance, ε . The comparison in Fig. 13 shows that the highest accuracy achieved with both hyperparameters is very alike for all prediction steps. However, the 25th and 75th percentile, the upper whisker and the median are lower for models using a fixed number of neighbors k . Therefore, in the reduced search space, only models using a fixed number of neighbors k are considered.

6.8 Number of Neighbors k of NNR and SNNR

In Fig. 14, the forecast accuracy of all k nearest neighbor models included in the search are shown. It can be observed, that the forecast accuracy mostly changes using 1 to 5 neighbors. The distribution of the RMSE becomes narrower for most cases with an increasing k . Furthermore, the lowest RMSE decreases slightly for $k \geq 10$. Therefore, models using between 10 and 20 neighbors are considered the reduced search space.

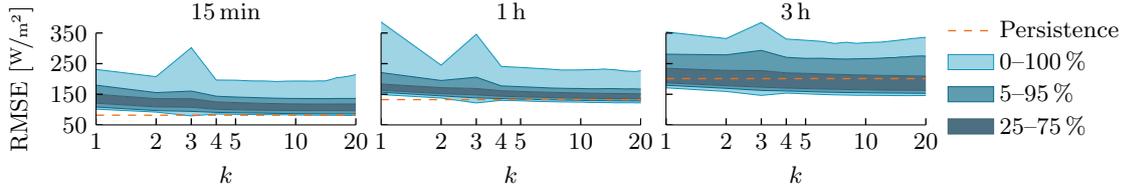


Fig. 14: RMSE of NNR and SNNR using a different number of neighbors k .

6.9 Summary

In summary, the following conclusions can be drawn from the analysis.

1. The probability of finding a sufficiently accurate model to forecast irradiance is much higher for NNR and SNNR than for ARIMA and SARIMA models (see Section 6.1).
2. Data pre- and postprocessing of NNR and SNNR models:
 - (a) Use transmissivity instead of irradiance (Section 6.2).
 - (b) Include night data in the reference sample (Section 6.3).
 - (c) Use data from the past 60 d in the reference sample (Section 6.4).
3. Hyperparameters of NNR and SNNR models:
 - (a) Seasonal models are beneficial (Section 6.5).
 - (b) Favor uniform weights over weights inverse to distance (Section 6.6).
 - (c) Choose the neighbors using a fixed number k instead of a maximum distance (Section 6.7).
 - (d) Search for models with 10 to 20 neighbors (Section 6.8).

Based on these findings the hyperparameter search space, i.e., the number of potential models could be reduced from more than 250,000 to less than 1,000. The search space for suitable k nearest neighbor (k NN) forecasters now only includes the number of autoregressive lags with range $1, 2, \dots, 11$, number of seasonal autoregressive lags with range $1, 2, \dots, 7$ and number of neighbors with range $10, 11, \dots, 20$.

7 Conclusion

A comparison of ARIMA and NNR for short term solar forecasts showed that it is more likely to find good forecasting models using NNR. Based on this finding, we derived a reduced search space of models that are likely to provide a good NNR forecasting model on unseen data.

Future work concerns an extension of the current approach to include exogenous data. Additionally, other artificial intelligence based methods, e.g., neural networks and support vector regression, are planned to be investigated.

Bibliography

- [1] REN21 Secretariat: Renewables 2017 global status report. Technical report, REN21 (2017)
- [2] Denholm, P., Hand, M.: Grid flexibility and storage required to achieve very high penetration of variable renewable electricity. *Energy Policy* **39**(3) (2011) 1817–1830
- [3] West, S.R., Rowe, D., Sayeef, S., Berry, A.: Short-term irradiance forecasting using skycams: Motivation and development. *Sol. Energy* **110** (2014) 188–207
- [4] Pedro, H.T., Coimbra, C.F.: Assessment of forecasting techniques for solar power production with no exogenous inputs. *Sol. Energy* **86**(7) (2012) 2017–2028
- [5] Pedro, H.T., Coimbra, C.F.: Nearest-neighbor methodology for prediction of intra-hour global horizontal and direct normal irradiances. *Renewable Energy* **80** (2015) 770–782
- [6] Shi, J., Lee, W.J., Liu, Y., Yang, Y., Wang, P.: Forecasting power output of photovoltaic systems based on weather classification and support vector machines. *IEEE Trans. Ind. Appl.* **48**(3) (2012) 1064–1069
- [7] Bacher, P., Madsen, H., Nielsen, H.A.: Online short-term solar power forecasting. *Sol. Energy* **83**(10) (2009) 1772–1783
- [8] Zeng, J., Qiao, W.: Short-term solar power prediction using a support vector machine. *Renewable Energy* **52** (2013) 118–127
- [9] Reikard, G.: Predicting solar radiation at high resolutions: A comparison of time series forecasts. *Sol. Energy* **83**(3) (2009) 342–349
- [10] Sfetsos, A., Coonick, A.: Univariate and multivariate forecasting of hourly solar radiation with artificial intelligence techniques. *Sol. Energy* **68**(2) (2000) 169–178
- [11] 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) (2010) 807–821
- [12] EPEX SPOT SE: 15-minute intraday call auction (2016) URL: <https://www.epexspot.com/document/29113/> (accessed Apr. 24, 2018).
- [13] ISO 9488:1999: Solar Energy – Vocabulary. International Organization for Standardization, Geneva, CH (1999)
- [14] ARM: Climate research facility, surface meteorology system (MET). Eastern North Atlantic Facility ARM Data Archive: Oak Ridge, US (2010) (accessed Jul. 14, 2011).
- [15] Reda, I., Andreas, A.: Solar position algorithm for solar radiation applications. *Sol. Energy* **76**(5) (2004) 577–589
- [16] Quaschnig, V.: *Regenerative Energiesysteme*. München, Hanser (2009)
- [17] Bontempi, G., Taieb, S.B., Le Borgne, Y.A.: Machine learning strategies for time series forecasting. In: *European Business Intelligence Summer School*, Springer (2012) 62–77
- [18] Meeus, J.H.: *Astronomical algorithms*. Willmann-Bell, Inc. (1991)
- [19] Box, G.E.P., Jenkins, G.M., Reinsel, G.C.: *Time Series Analysis: Forecasting and Control*. Prentice-Hall, Inc. (1994)
- [20] Altman, N.S.: An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician* **46**(3) (1992) 175–185
- [21] Lin, D.: *Statistical learning toolbox* (2006)
- [22] Andreas, A., Stoffel, T.: Nevada power: Clark station; Las Vegas, Nevada, US (data). Technical Report DA-5500-56508, NREL, Golden, US (2006)
- [23] Andreas, A., Stoffel, T.: NREL solar radiation research laboratory (SRRL): Baseline measurement system (BMS); Golden, Colorado, US (data). Technical Report DA-5500-56488, NREL, Golden, US (1981)