Irradiance Prediction Intervals for PV Stochastic Generation in Microgrid Applications

Enrica Scolari[a,*], Fabrizio Sossan[a], Mario Paolone[a]

[a] *École Polytechnique Fédérale de Lausanne, Switzerland*

Abstract

The increasing interest in integrating volatile resources into microgrids implies the necessity of quantifying the uncertainty of photovoltaic (PV) production using dedicated probabilistic forecast techniques. The work presents a novel method to construct ultra-short-term and short-term prediction intervals (PIs) for solar global horizontal irradiance (GHI). The model applies the k-means algorithm to cluster observations of the clear-sky index according to the value of selected data features. At each timestep, the features are compared with the actual conditions to identify the representative cluster. The lower and upper bounds of the PI are calculated as the quantiles of the irradiance instances belonging to the selected cluster at a target confidence level. The validation is performed in 3 datasets of GHI measurements, each one of 85 days. The model is able to deliver high performance PIs for forecast horizons ranging from sub-second to intra-hour ahead without the need of additional sensing systems such as all-sky cameras.

Keywords: Solar forecasting, Prediction interval, Ultra-short term, k-means algorithm.

1. Introduction

The thrust toward increasing the penetration of non-dispatchable renewable generation in the electrical grid requires to redefine conventional practices to

*Corresponding author
Email address: enrica.scolari@epfl.ch (Enrica Scolari)

assure reliable power system operation, see for example [1, 2]. A paradigm increasingly advocated in the recent technical literature to cope with the variability of stochastic generation is the development of robust and predictive controls. They take advantage of short-term forecasts of renewable generation in order to plan adequate counteractions to prevent, or mitigate, operational issues related to renewables power fluctuations. Examples include the dispatchability of renewables, achieving self-consumption of locally generated electricity, and the short-term redispatching of conventional generation units, see for example [3, 4, 5, 6]. The period of the redispatch control action normally depends on the availability of the reserve in a given power grid and on the performance of the forecasting tools accounting for the uncertainties. For the case of microgrids, their limited physical extension and the low granularity of the resources involve the necessity of coupling the reserves dispatch with their real-time control. In this respect, a new protocol for real-time control of microgrids has been presented in the recent literature; in this framework, the control decision is updated with a sub-second resolution, [7, 8]. Since photovoltaic (PV) systems represent one of the major resources in modern microgrids, the availability of irradiance forecasting is beneficial to address the aforementioned challenges at forecast horizons from sub-second up to intra-hour, [2]. A further concern associated with the dense penetration of PV installations in distribution systems and microgrids is the lack of the spatial smoothing effect, resulting in large variations of the solar irradiance. As an example, Figures 1a and 1b respectively show daytime global horizontal irradiance measurements (GHI, recorded at the EPFL campus by using a pyranometer) and the power consumption of a group of EPFL buildings equipped with a 95 kWp PV-roof system. As visible, GHI variations (which varies up to 85% in magnitude in less than two minutes) cause very steep fluctuations of the power production and consumption. The availability of high-quality ultra-short-term and short-term GHI forecast enables the possibility of taking preemptive control actions and mitigating the effect of its fluctuations.

In general, the choice of the forecast method is strictly related to the target

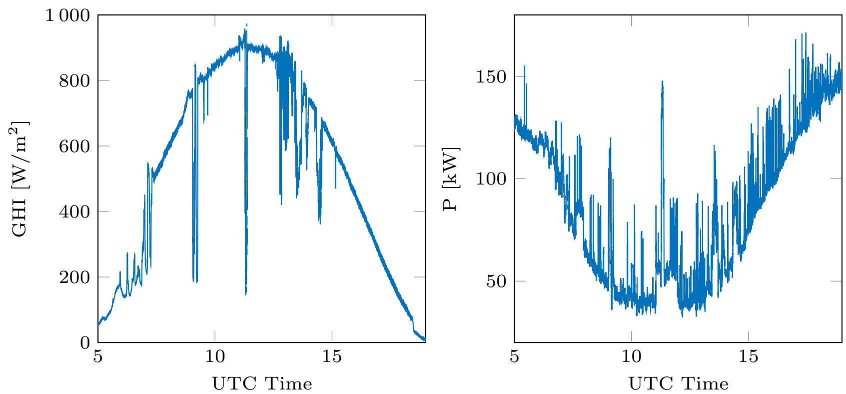

(a) GHI Profile (b) Aggregated Power Profile

Figure 1: GHI and aggregated power profile (load and PV production) as a function of the UTC Time, registred at the EPFL campus on the 15^{th} of May 2016.

forecast horizon and geographical scale. As explained in [9], day-ahead regional irradiance forecasting relies on satellite observations and numerical weather predictions (NWPs). However, we here focus on local and shorter term forecasts (lower than one hour) where Artificial Intelligence (AI) methods are generally applied, [10]. The use of all-sky cameras is a promising solution for intra-hour forecast horizons, as introduced in [11, 12]. To the best of our knowledge, the only method addressing the problem of solar irradiance forecast at sub-second time scale is the one proposed in [13].

Two main kinds of forecast are conceived: a first kind (deterministic) considers only the point forecast while the second one (probabilistic) includes information accounting for the intrinsic uncertainty of the prediction and it is more appropriate when dealing with control and decision making in modern power systems, [14]. Especially in the case of fast irradiance ramps, generally difficult to predict, PIs are necessary to define the worst-case scenario that should be considered in the control decision process.

Regarding GHI point predictions, the simplest forecast model is the persistent one, which is commonly used as a benchmark for performance evaluation. It assumes that the GHI remains constant with the forecast horizon. In general, most of the point forecast techniques are based on AI methods. A more deterministic approach consists in detecting the position of the clouds, deducing clouds motion and calculating the time when a cloud covers the sun, e.g. [11, 15]. Apart from cloud detection and motion, sky images contain more information impacting the GHI prediction: examples are the cloud cover and the type of clouds. This kind of information can be combined with machine learning methods to compute the forecast, e.g. [16, 17].

Several works address the problem of probabilistic forecast and propose PIs computation models. Probabilistic solar power forecast is proposed in [18, 19] where a set of likely predictions (i.e. an ensemble) is provided using a historical set of variables and deterministic meteorological models. Authors of [18] use a distance criterion to retrieve similar past forecasts, under the assumptions that their errors are likely to be similar to the errors of the current forecast. These methods refer to 0-72 hours forecast horizons, considering hourly power data. In [20], a hybrid model is proposed, integrating Support Vector Machine (SVM), ANN and sky imaging techniques to deliver real-time PIs for direct normal irradiance (DNI) for 5, 10, 15, 20 minutes ahead. At each time step, the computational time is less than 5 seconds. Another stochastic approach in [21] proposes the design of a k-nearest neighbors (KNN) algorithm. The KNN algorithm is used to predict the GHI and DNI and their uncertainty intervals, for time horizons from 5 to 30 minutes. More recently, Authors of [22] proposed a data-driven method to construct GHI probability densities for one hour-ahead predictions, using nonparametric bootstrap and a map of solar position. The developed method has low computational complexity, requiring 0.56 seconds on a personal computer. In [23], point forecasts are generated using AutoRegressive Moving Average (ARIMA) and the associated PI is calculated using a Generalized AutoRegressive Conditional Heteroskedasticity model (GARCH), considering a prediction horizon from 10 minutes to 6 hours. The use of recursive formulas, to update the model parameters in real-time, allows to reduce the computational complexity of the method.

Having stated this, we note that the available literature lacks of a unique forecasting tool for prediction horizons ranging from sub-second up to intra-hour and capable of operating at low levels of aggregation, where the level of volatility is higher due to the reduced spatial smoothing effect. While many methods have been proposed for intra-hour GHI forecasting and might be applied to deliver ultra-short-term predictions, there is at least the compelling need of re-assessing their performance in the light of the requirements of real-time control of local power systems. Moreover, computational complexity becomes a key concern when considering the high reporting rate of ultra-short predictions, implying that available forecasting methods with complex on-line training procedures (like ANNs, heuristic optimization-based and sky imaging) might not be suitable.

We propose a novel nonparametric method for ultra-short term forecasting of the global horizontal irradiance (GHI) to deliver predictions with a forecast horizon in the range from 500 ms to 5 minutes, thus suitable both in the context of real-time control of microgrids and energy management strategies. PIs computation is based on a well-known pattern recognition technique called k-means clustering, [24]. A training dataset is first clustered considering two empirically selected influential variables. Then, PIs are calculated by extracting the quantiles of the cluster which resembles at most the actual conditions. A clear-sky model is also introduced for the de-trending of the GHI time series. The method does not require any information from sky-imaging since it only relies on measurements of the GHI, it is computationally efficient and needs a limited training dataset. As later qualified in the paper, the real-time generation of the PIs takes, for one time instance, less than 0.5 ms. Thus, the method is applicable even when the control decision has to be taken at sub-second time scale.

The paper is organized as follows: Section 2 defines the problem and introduces the nomenclature, Section 3 explains in details the methodology adopted to deliver the PIs and discusses its computational complexity. Section 4 describes the available datasets and their characterization. Section 5 presents the results, comparing them with available benchmark methods. Section 6 shows the main conclusions.

2. Preliminaries on the Adopted Nomenclature

PIs give a range of possible values in which the future realization is expected to lie with a given confidence level α, [25]. At the timestep t, we denote the one-step-ahead PI at confidence level α as composed by the upper and lower bounds:

$$\left(I_{t+1|t}^{\uparrow\alpha}, I_{t+1|t}^{\downarrow\alpha} \right). \tag{1}$$

Note that we do not have any assumption on the distribution of the time series since we adopt a nonparametric approach.

GHI measurements are pre-processed in order to remove the daily and seasonal components due to changes of the sun position. This is achieved by introducing the clear-sky index K, which is defined as the ratio between the measured GHI and the clear-sky irradiance, respectively denoted by I and I_{cs}:

$$K = \frac{I}{I_{cs}}. \tag{2}$$

The clear-sky irradiance is the irradiance that would reach the ground in clear-sky conditions, i.e. absence of clouds. It is obtained by applying the clear-sky model implemented in the geographical information system GRASS, which also account for topological shading [26, 27].

3. Methods

The principle behind the proposed forecasting approach is that, if a realization of the solar irradiance happened in the past under certain measurable conditions, then it is likely to happen again under the same confluence of circumstances.

The proposed PI estimation method consists in clustering historical data according to the value of certain influential variables, introduced in the following. The clusters are therefore used as empirical conditional probabilities of future realizations and used to compute the PI by calculating the quantiles according to a given confidence level. In particular, these influential variables should be representative of the irradiance fluctuations since it is the main cause of the uncertainty associated with solar forecasts. These variables, inputs of the clustering process, are selected according to the literature that considers the average and the variability of the clear sky-index as the most influential ones, [21, 28]. We consider a training dataset of historical clear-sky index observations $K_1, \ldots, K_N$, from which we extract the following influential variables:

- the average clear-sky index value on a mobile window of length n considering the most recent data points:

$$M_i = \frac{1}{n} \sum_{j=i-n}^{i} K_j, \quad i = n+1, \ldots, N \tag{3}$$

  of which we consider the normalized version M_i^*. Namely, we normalize the sequence $M_0, \ldots, M_1$ to a length of 1; [1]

- the clear-sky index variability:

$$V_i = \sqrt{\frac{1}{n} \sum_{j=i-n}^{i} (K_j - K_{j-1})^2}, \quad i = n+1, \ldots, N \tag{4}$$

  which is a measurement of GHI fluctuations. As for the previous case, we consider the normalized version V_i^*.

Normalization of the influential variables is required to enable a fair comparison between parameters with different scale. For each observation, the vector $\mathbf{p_i}$ of influential variables is:

$$\mathbf{p_i} = (M_i^*, V_i^*), \qquad i = n+1, \ldots, N. \tag{5}$$

[1] The normalized version of a vector $\mathbf{X}$ is a vector $\mathbf{X}^* = \mathbf{X}/|\mathbf{X}|$, where $|\mathbf{X}|$ is the norm of $\mathbf{X}$.

The process to compute PIs is performed in two ways:

- Method A: we cluster the original clear-sky index time series;
- Method B: we cluster the differentiated clear-sky index time series:

$$\Delta K_i = K_i - K_{i-1}, \;\; i = 2, \ldots, N, \tag{6}$$

  to verify if differencing leads to better prediction performance.

3.1. Clustering of the training set

The k-means iterative algorithm is firstly used to classify historical observations of clear-sky index according to predefined influential variables. K-means clustering is a partitioning algorithm that allocates each observation into exactly one of the k clusters, each one defined by a representative centroid. In particular, k centroids are at first randomly selected (the first centroids are simply uniformly random observations). Then, each vector of the training dataset is assigned to the closest centroid, and the centroid is iteratively recalculated as the mean of the vectors of each class until convergence is reached (i.e., centroids do not change anymore between iterations).

Method A

We apply the k-means algorithm to cluster the vectors $\mathbf{p_i}$ belonging to the training set, being k the number of clusters. The algorithm assigns to each vector $\mathbf{p_i}$ a cluster index l between 1 and k and determines the centroids locations $\mathbf{c_l} = (M_l^*, V_l^*)$ for $l = 1, \ldots, k$.

We denote the generic cluster G_l as composed by all the clear-sky indexes K_{i+1} for which $\mathbf{p_i}$ has index l.

Method B

We apply the same clustering procedure described above for Method A.

However, we denote the generic cluster ΔG_l as composed by all the differentiated clear-sky realizations ΔK_{i+1} for which $\mathbf{p_i}$ has index l.

An example of the k-mean partitioning of the influential variables is shown in Fig. 2 where the normalized clear-sky index average and variability are clustered.

It is worth noting that the k-means clustering of the training dataset can be performed off-line on historical data. This is a key aspect if considering the high reporting rate of predictions for microgrid applications since it allows to reduce the computational complexity.

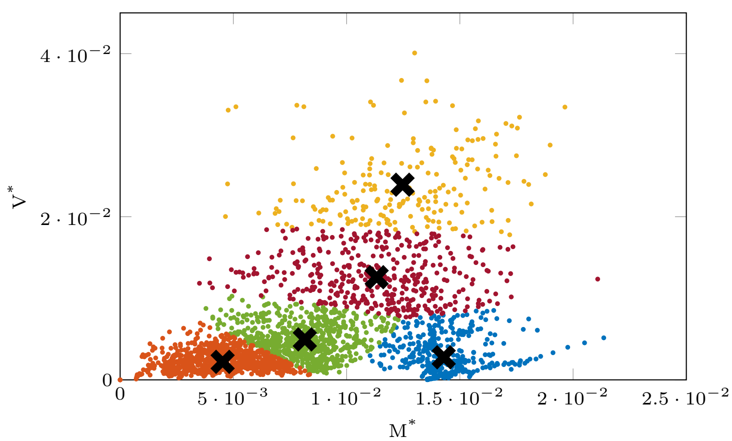

Figure 2: Example of k-means clustering obtained for $k = 5$. The x and y axis represent the normalized clear-sky index average and variability, respectively. The black marker signs the centroids.

3.2. Prediction Intervals

In this section we describe how PIs are computed, distinguishing between the two proposed methods.

Method A

Starting from the clusters $G_1, \ldots, G_k$ defined in the previous section, the PIs at the target confidence level α can be computed as:

$$q_l^{\uparrow\alpha} = (1+\alpha)/2 \text{ quantile of } G_l, \ \ l = 1, ..., k \tag{7}$$

$$q_l^{\downarrow\alpha} = (1-\alpha)/2 \text{ quantile of } G_l, \ \ l = 1, ..., k. \tag{8}$$

For increased computational efficiency, we note that also this operation can be performed off-line, and the PIs for each class can be stored.

Say being at time t, the objective is to perform the on-line computation of the PI for the next time interval $t+1$. The vector of influential variables at t is denoted by $\mathbf{p_t} = (M_t^*, V_t^*)$. It is calculated normalizing the raw influential variables M_t, V_t with respect to those available in the training data set. The next step is the calculation of the Euclidean distances between $\mathbf{p_t}$ and the centroids $\mathbf{c_l}$:

$$d_l = \|\mathbf{c_l} - \mathbf{p_t}\|^2, \quad l = 1, ..., k \tag{9}$$

which is used as a similarity criterion to select the cluster representative of the future clear-sky outcome. We indicate with $\hat{l}$ the index corresponding to the cluster with minimum distance. It is used to select the quantiles used in the PI computation as:

$$K_{t+1|t}^{\uparrow\alpha} = q_{\hat{l}}^{\uparrow\alpha}. \tag{10}$$

$$K_{t+1|t}^{\downarrow\alpha} = q_{\hat{l}}^{\downarrow\alpha} \tag{11}$$

Method B

Starting from the clusters $\Delta G_1, \ldots, \Delta G_k$ obtained from the differentiated GHI time series, the PIs at the target confidence level α can be computed as:

$$q_l^{\downarrow\alpha} = (1-\alpha)/2 \ \text{ quantile of } \ \Delta G_l, \ \ l = 1, ..., k \tag{12}$$

$$q_l^{\uparrow\alpha} = (1+\alpha)/2 \ \text{ quantile of } \ \Delta G_l, \ \ l = 1, ..., k. \tag{13}$$

Also in this case, the quantiles extraction is computed off-line. The on-line computation of the PIs consists in finding the index $\hat{l}$ of the cluster with centroid at the minimum distance from $\mathbf{p_t}$. It is used to select the quantiles used in the PI computation as:

$$K_{t+1|t}^{\uparrow\alpha} = K_t + q_{\hat{l}}^{\uparrow\alpha}, \tag{14}$$

$$K_{t+1|t}^{\downarrow\alpha} = K_t + q_{\hat{l}}^{\downarrow\alpha}. \tag{15}$$

i.e., the current measurement is summed to the upper and lower quantiles of the differentiated time series.

It is important to note that, so far, PIs are computed to forecast the clear-sky index. The last step consists in computing the PI for the GHI:

$$\hat{I}^{\uparrow\alpha}_{t+1|t} = K^{\uparrow\alpha}_{t+1|t} I_{cs,t+1}, \tag{16}$$

$$\hat{I}^{\downarrow\alpha}_{t+1|t} = K^{\downarrow\alpha}_{t+1|t} I_{cs,t+1}. \tag{17}$$

In the results section, performance of methods A and B are evaluated by comparing the estimated PIs as defined in (16)-(17) with the GHI measurements.

3.3. Selection of the Parameters

The parameters we need to specify when applying the k-means clustering procedure are:

- The number of samples n used in (3) and (4);
- The number of cluster k used for the partition of the training dataset;
- The length of the training dataset N.

The selection of the parameters values is a sensitivity process that is exhaustively evaluated in Section 5. The assessment is performed in a searching dataset, and then the selected values are applied in a testing dataset for performance evaluation.

While it was seen from the results that variations of n in the range from 2 to 5 do not, in general, alter modeling performance, the values of k and N are interdependent and the selection of these two parameters needs to be carried out simultaneously (i.e. we need to find the combination of k and N with the best performance). For each specific case, a diagnostic analysis should be performed. We consider here two main approaches to fix k and N as described in the next subsections.

3.3.1. A-posteriori Selection of k with Exhaustive Search (ES)

Several (k, N) combinations are attempted in a searching dataset: the candidate combination is the one with best a-posteriori prediction performance. These values are then used to evaluate the performance of the testing dataset, supposing that it exhibits similar characteristics of the searching set. This approach can become computationally expensive, therefore motivating the development of methods for the a-priori selection of the free parameters, as described in the next subsection.

3.3.2. A-priori Selection of k with Silhouette Analysis (SA)

The objective is to use Silhouette Analysis, [29], to improve the partitioning of the training dataset, allowing for an a-priori selection of parameter k and in order to avoid the exhaustive approach. The Silhouette Analysis consists in some main steps:

- a small value of k is chosen (e.g. $k = 5$) and the clustering algorithm is run;
- the silhouette value for a generic point i is calculated as:

$$s(i) = \frac{a(i) - b(i)}{\max(a(i), b(i))} \tag{18}$$

  where $a(i)$ is the average distance from point i to the other points in the same cluster, while $b(i)$ is the minimum average distance from instance i to points in a different cluster, minimized over clusters. Parameter s is a measure of how close the instance is to the other instances in its cluster and how far it is to those in the other clusters. In general, a silhouette value close to 1 is desired because it means that the point is well clustered while a value close to -1 means misclassification;
- the mean of the silhouette values is computed. If most points have a high silhouette value, then the clustering is appropriate;
- the value of k is augmented and it is evaluated if having more clusters allows for a better partitioning (higher mean of the silhouette values);

- the number of clusters is selected equal to the value k above which we do not see any improvement in terms of increasing of the average silhouette value.

As shown in subsection 5.3.2, the value of N required to converge at constant k is not sensitive to the characteristics of the dataset. In general, for each forecast horizon and given k, it is possible to identify a value of N above which performance is close to convergence. Above this value, small oscillations are explained by the intrinsic stochasticity of the data. This feature is important for the modeler since it allows for the reduction of the parameters to be found.

3.4. Algorithms time complexity

In this section, we evaluate the time complexity of the real-time computation of the proposed algorithms. This is an important aspect because they are designed with the stated objective of delivering PIs to real-time control processes for electrical power systems.

The algorithms consist of two parts, the training phase and on-line computation of PIs. The former does not have any real-time requirement and can be performed off-line. The latter phase is instead time critical, and it is to perform with a hard-real deadline. First, it consists in calculating the normalized influential variables, (3)-(4), an operation with constant time complexity, $O(1)$, which involves algebraic operations. Then, we have the computation of k norms, $O(k)$, and a minimum search, which can be performed efficiently with a merge search, $O(k\log(k))$. Considering that k is fixed by design, the time complexity of the real-time computation is constant time, $O(1)$, therefore denoting that complexity does not grow with the size of the problem (scalability). Statistics on the execution time of the algorithms are given in Section 5.8.

4. Data

4.1. Data acquisition

Global horizontal irradiance measurements are collected at 50 ms resolution at the Écolecole Polytechnique Fédérale de Lausanne (EPFL) by using an

Apogee SP-230 all-seasons pyranometer which is located at the following GPS coordinates: 46.518397-N, 6.565229-E. We consider three datasets of 85 days each, corresponding to different periods of the year. The first contains irradiance measurements from July to September 2015 (Summer), the second from October to December 2015 (Autumn), and the third from January to March 2016 (Winter). Each dataset is divided into a searching subset of 55 days and a testing one composed of the remaining 30 days.

The original time series is down-sampled to three different resolutions: 500 ms, 1 minute and 5 minutes. These series are used to compute one-step-ahead PIs for the corresponding forecast horizon. Down-sampling is computed by averaging the intermediate samples.

It is worth noting that applying the clear-sky normalization causes very high values of K close to sunrise and sunset. Therefore, we consider only daylight values covering the period of the day for which the clear-sky index does not diverge.

4.2. Data classification

Characterizing the dataset is important for performance comparison and evaluation. Indeed, the robustness of the method should be tested during periods of different irradiance volatility. In our case, we are interested in characterizing the three available datasets: Summer, Autumn, and Winter. First, we give an information regarding the weather of the selected period and location. In particular, we retrieve cloud cover data [2] from MeteoSwiss Idaweb services, [30], from two weather stations in the vicinity of our installation. The average cloud okta values for the three seasons are: 3.86 okta in Summer, 4.67 okta in Autumn and 5.96 okta in Winter.

To be more specific, we introduce a criterion consisting in counting the per-

[2] Cloud cover corresponds to the fraction of the sky obscured by clouds when observed from a given location. The unit of measurement is the okta, ranging from 0 (completely clear-sky) to 8 (completely overcast).

centage of periods with a volatility lower than a given threshold. For each timestep t, we calculate the per-unit difference as $\Delta I_t = (I_t - I_{t-1})/I_{max}$, where $I_{max} = 1000$ W/m^2. For each prediction horizon, we establish a threshold for ΔI_t, above which the observation at time t is considered with high volatility. The threshold is empirically computed as the 99% quantile of the ΔI time series obtained by manually selecting a period of 3 clear-sky days. The values are shown in Table 1 for different forecast horizons.

Table 1: Thresholds.

Time Horizon	**Thresholds**
500 ms	0.0004
1 min	0.011
5 min	0.025

Table 2: Percentage of periods with high irradiance volatility.

	Forecast Horizon		
Season	500 ms	1 min	5 min
Summer	16	17	22
Autumn	5	9	12
Winter	13	15	20

The percentage of periods exceeding the threshold of GHI high volatility is shown in Table 2, for each dataset and for different forecast horizons. As it can be observed, the Summer period is characterized by the highest GHI volatility, followed by Winter and Autumn.

5. Results and Discussion

First, the metrics used for performance evaluation are introduced in Subsection 5.1. Then, Subsection 5.2 shows the advantage given by the introduction of a clear-sky model at different forecast horizons. In Subsection 5.3 the sensitivity of the performance with respect to the selection of the model parameters is discussed. In Subsections 5.4-5.6 the performance of the proposed methods is benchmarked against existing techniques. First, we compare the proposed methodology with the symmetric quantile extraction, which is the simplest way to construct our intervals. We use the empirical quantiles extracted from the distribution of the time series to build the PIs as in (10)-(11) and (14)-(15), respectively. It is important to highlight that the quantiles at time t are extracted from the whole time series, from $t = 0$ to $t - 1$.

As a second benchmark, we compare our method with a model commonly used in forecasting. We first generate a point forecast using AutoRegressive Moving Average model (ARIMA), [31], with Double Exponential Smoothing. Then PIs are constructed assuming a Gaussian distribution of the point forecast error as:

$$K^{\uparrow\alpha}_{t+1|t} = \hat{K}_{t+1} + \eta_\alpha \sqrt{\sigma_t}, \tag{19}$$

$$K^{\downarrow\alpha}_{t+1|t} = \hat{K}_{t+1} - \eta_\alpha \sqrt{\sigma_t}. \tag{20}$$

where $\hat{K}_{t+1}$ is the point forecast obtained by the ARIMA model, η_α is the quantile of the normal distribution corresponding to the target confidence level α and σ_t is the variance of the forecast error.

Unless otherwise indicated, the target confidence level used for the following analysis is fixed equal to 95%.

Subsection 5.7 presents and discusses the reliability diagrams. Finally, statistics of the method execution time are provided in Subsection 5.8.

5.1. Metrics

We use three standardized metrics from the existing literature to evaluate the performance of the proposed methods [20, 32]. The first metric is the PI coverage probability (PICP) which counts the number of times that the realization falls inside the PI for a given confidence level α:

$$\mathrm{PICP} = \frac{1}{L}\sum_{t=1}^{L} c_t \tag{21}$$

where L is the total number of forecast instances of the testing dataset and

$$c_t = \begin{cases} 1, & \hat{I}^{\downarrow\alpha}_{t+1|t} \leq I_{t+1} \leq \hat{I}^{\uparrow\alpha}_{t+1|t} \\ 0, & \text{otherwise.} \end{cases} \tag{22}$$

Then, to account for the fact that the wider the PI, the easier it is to have a realization falling inside it, we measure the PI normalized averaged width (PINAW):

$$\mathrm{PINAW} = \frac{1}{LI_{max}}\sum_{t=1}^{L}(\hat{I}^{\uparrow\alpha}_{t+1|t} - \hat{I}^{\downarrow\alpha}_{t+1|t}), \tag{23}$$

where $I_{max} = 1000$ W/m^2. The value has been selected accordingly to the work in [21], that is later used as benchmark for performance comparison. The third metric quantifies the trade-off between having a large coverage probability and small interval width. It is called coverage width-based criterion (CWC):

$$\mathrm{CWC} = \mathrm{PINAW}(1 + \gamma(\mathrm{PICP})e^{-\mu((\mathrm{PICP})-\mu)}) \tag{24}$$

where

$$\gamma = \begin{cases} 0, & \mathrm{PICP} \geq \alpha \\ 1, & \mathrm{PICP} < \alpha. \end{cases} \tag{25}$$

The parameter μ can be tuned based on how much bad PIs are to penalize, see [32]. We select here $\mu = 10$. PIs should have high PICP (higher or equal to the target confidence level) coupled with a low value of PINAW.

5.2. Clear-sky index and GHI time-series comparison

A first analysis aims at assessing the difference between using the measured irradiance (GHI) or the clear-sky index (K) time series as inputs for the PIs

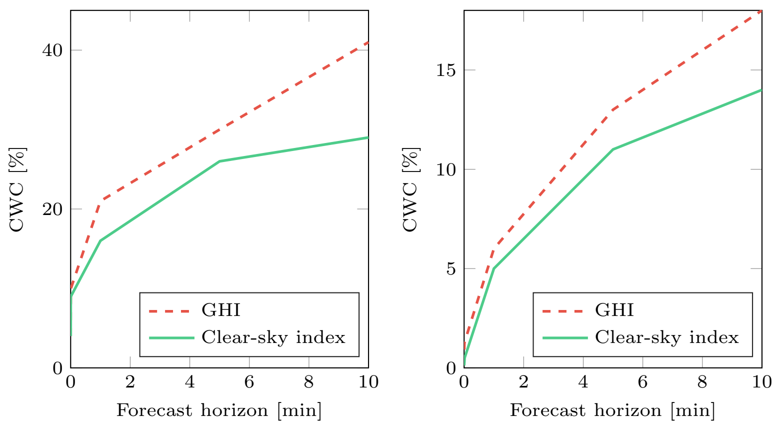

(a) Method A (b) Method B

Figure 3: CWC as a function of the time horizon. Comparison between the use of the GHI time series (original) and the clear-sky index one (normalized).

computation method. In particular, we apply methods A and B and we increase the forecast horizon to evaluate when the inclusion of a clear-sky model becomes advantageous.

Fig. 3 shows that the CWC is, in general, lower (better performance) when using the clear-sky index. In particular, the advantage of using a clear-sky model becomes evident for time horizons longer than 1 minute. As expected, the normalization of the time series becomes more important at higher forecast horizons, when the effect of the sun position becomes more dominant. When referring to ultra-short term horizons, fluctuations of solar irradiance are mainly related to cloud motion and the importance of a clear-sky model becomes marginal. Since the inclusion of a clear-sky model only leads to similar or better prediction performance, the proposed methods are applied to the clear-sky index time series.

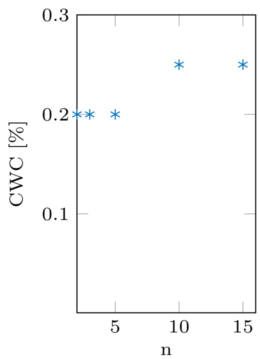

(a) 500 ms

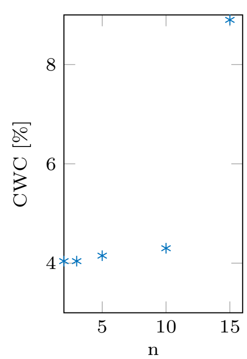

(b) 1 min

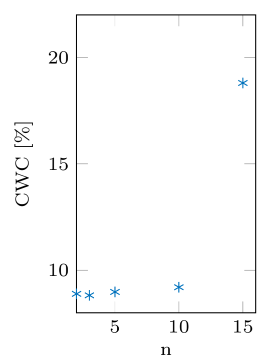

(c) 5 min

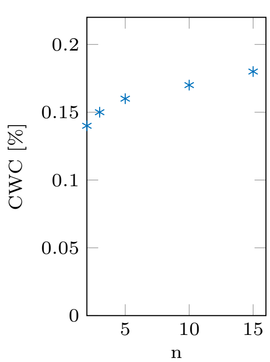

(d) 500 ms

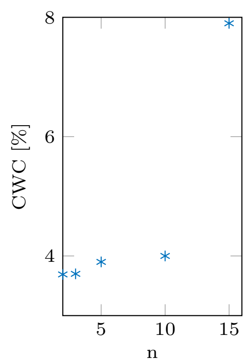

(e) 1 min

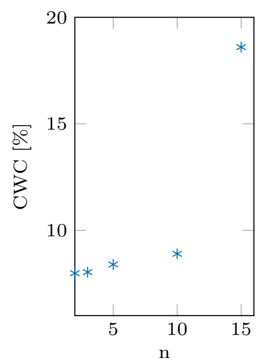

(f) 5 min

Figure 4: CWC as a function of n for different forecast horizons. N is equal to 30 days. k is equal to 5 for cases (a)-(c) and equal to 30 for case (d)-(f).

5.3. Parameters Selection and Sensitivity analysis

5.3.1. Selection of parameter n

Fig. 4 shows the CWC metric as a function of n (3-4) obtained from a-posteriori analysis of the performance of the whole Autumn dataset. Method B is applied. The analysis considers different forecast horizons for both k=5, cases (a)-(c), and k=30, cases (d)-(f). It is possible to see that performance is not very sensitive to variations of n in the range from 2 to 5. From an a-posteriori analysis of our datasets at different forecast horizons, we can conclude that n can be fixed to a value between 2 and 5, for all the considered cases. Indeed, analogous conclusions can be inferred for Method A and different datasets (not shown here because of the similar behavior). These values of n are a good trade-off between having enough significant measurements to compute the influential variables and avoiding to consider realizations that are too far from the actual conditions. The results presented in what follows are obtained with $n = 3$.

5.3.2. Selection of parameter k and N

As explained in Subsection 3.3 two main procedures are proposed to assign parameters k and N. In the case of the exhaustive search, we isolate 55 days of each dataset to perform the exhaustive searching. The number of clusters k can vary between one (single cluster) and the total number of training samples (each data is assigned to its own cluster). Since the computational effort of the k-means algorithm is linearly dependent to the number of clusters and to the number of data, we here limit k to 1000 and N to 30 days. The dashed lines in Figures 5, 6 and 7 show the value of k that returns the best performance for a fixed N. This k is obtained a-posteriori by applying methods A and B to the searching dataset of 55 days ($k_{ES,A}$ and $k_{ES,B}$, respectively). The three figures refer to different datasets. Each figure includes three sub-figures referring to the three forecast horizons, respectively. It is possible to note that, due to the heuristic nature of the methods, the optimal value of k cannot be known a-priori and it varies among the different considered cases (namely, we do not have a global optimum). Therefore, we select k and N as the combination returning the best prediction performance (minimum CWC), found a-posteriori. These values found for the searching set of 55 days are then applied for performance evaluation in the remaining 30 days.

In the case of the silhouette analysis, k is calculated for different N as the one maximizing the average silhouette of the training set and it is shown in Figures 5, 6 and 7 with the solid line. The value of k that maximizes the average silhouette does not correspond to the one returning the best forecasting performance. However, we notice that its value does not vary with N and can be selected independently. In order to determine N in the case of the silhouette analysis, we fix k equal to the one returned by the analysis (k=5) and we evaluate the prediction performance for a different number of training days. Figures 8, 9 and 10 show the CWC (in logarithmic scale) as a function of N. They refer to 500 ms, 1 and 5 minutes, respectively. Each figure consists of two plots, showing the performance for Method A and B, respectively. We can make the following observations:

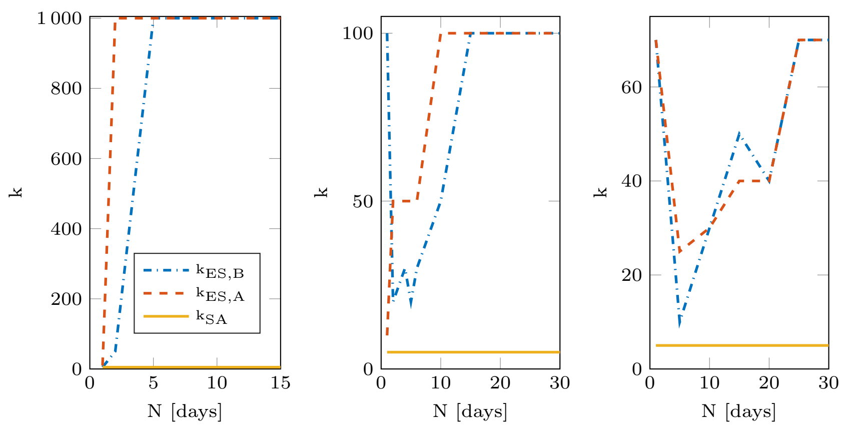

(a) 500 ms (b) 1 min (c) 5 min

Figure 5: Number of clusters k as a function of the length of the training dataset. The Summer dataset is selected for the analysis. The dashed lines refer to the value of k corresponding to maximum performance for Method A and B, respectively. It is calculated a-posteriori by applying the ES. The solid line refers to the value of k from the SA.

- For each forecast horizon, it is possible to identify a first drop of CWC after which performance tends to stabilize. The value of N that leads to performance stabilization is not sensitive to the dataset and can be fixed independently. On the contrary, as it is shown in the next sections, performance at convergence depends on the nature of the dataset and, in general, the behavior of the PIs depends on the volatility content of the dataset.

- For the sub-second time horizon and Method A we have a first drop of CWC after about one day and then performance tends to stabilize.

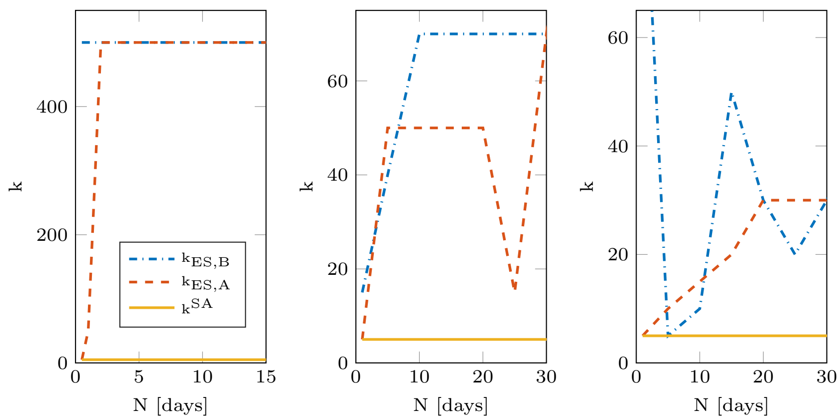

(a) 500 ms (b) 1 min (c) 5 min

Figure 6: Number of cluster k as a function of the length of the training dataset. The Autumn dataset is selected for the analysis.

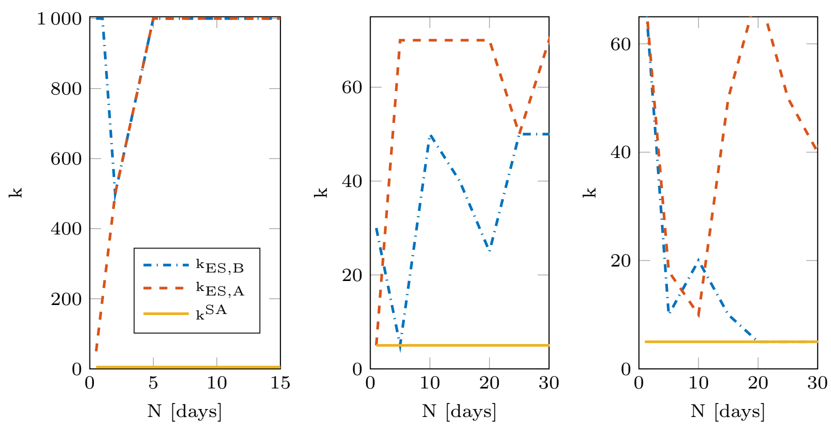

(a) 500 ms (b) 1 min (c) 5 min

Figure 7: Number of cluster k as a function of the length of the training dataset. The Winter dataset is selected for the analysis.

Method B reaches convergence after few hours of training with subsequent small CWC oscillations.

- For time horizon of 1 minute we have a first drop of CWC after about 5 days of training and then performance smooths out more slowly.

- For time horizon of 5 minutes we have a first drop of CWC after about 5 days of training and then a second drop after 10 days. Then, performance smooths out more slowly.

  In conclusion, when applying the exhaustive search we use the optimal combination of k and N found for the 55 days dataset as candidates for performance evaluation in the remaining 30 days. On the contrary, when applying the silhouette analysis approach, we select k=5 and N equal to 1, 5 and 10 days respectively for three time horizons. These values are valid for all the three datasets. A comparison between the two approaches is presented in what follows.

5.4. Ultra-short term forecasting: Performance Evaluation

We focus here on sub-second forecast horizon: one-step-ahead PIs at 500 ms. At first, we analyze the two proposed approaches to compute k and N. Then, performance of methods A and B is evaluated, and therefore compared with existing methods from the technical literature.

5.4.1. Exhaustive Search and Silhouette Analysis

In this section, we compare the performance obtained by applying the exhaustive search and the silhouette analysis to determine k and N. Furthermore, results are compared with the optimal performance found a-posteriori to evaluate how far the estimations are from the optimum. Evaluation is carried out in the testing set of 30 days, for each one of three datasets.

Results are shown in Tables 3 for methods A and B. The comparison considers metric CWC. For 500 ms forecast horizon, the exhaustive search is outperforming the silhouette analysis and returns performance very close to the optimal a-posteriori. Indeed, when dealing with high sampling frequency, the large amount of data would require a number of clusters which is much higher than the one returned by the silhouette analysis. For sub-second time horizons, we proceed our analysis by using the exhaustive search method.

The following additional conclusions can be drawn from this analysis:

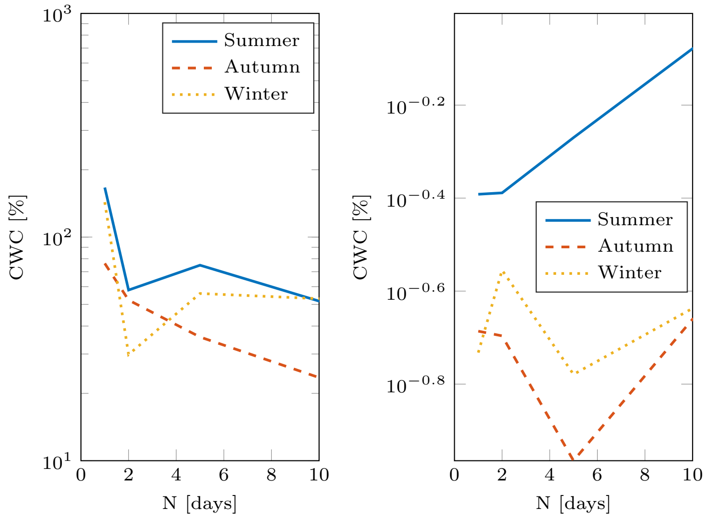

(a) Method A-500 ms

(b) Method B-500 ms

Figure 8: CWC as a function of the number od training days for 500 ms time horizon. CWC is shown in logarithmic scale.

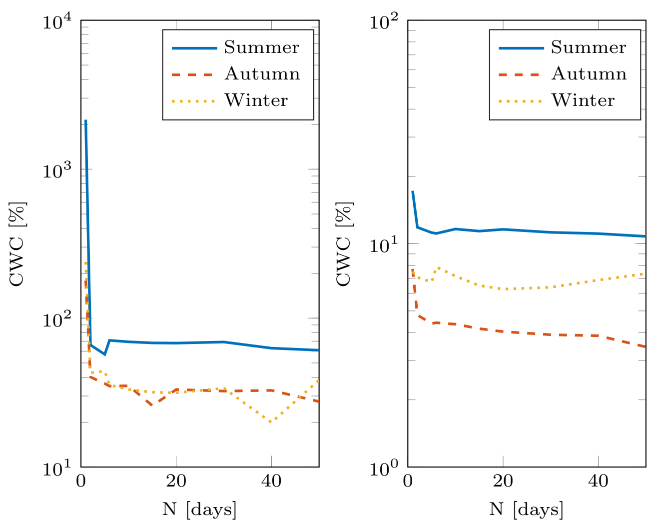

(a) Method A-1 min

(b) Method B-1 min

Figure 9: CWC as a function of the number od training days for 1 min time horizon. CWC is shown in logarithmic scale.

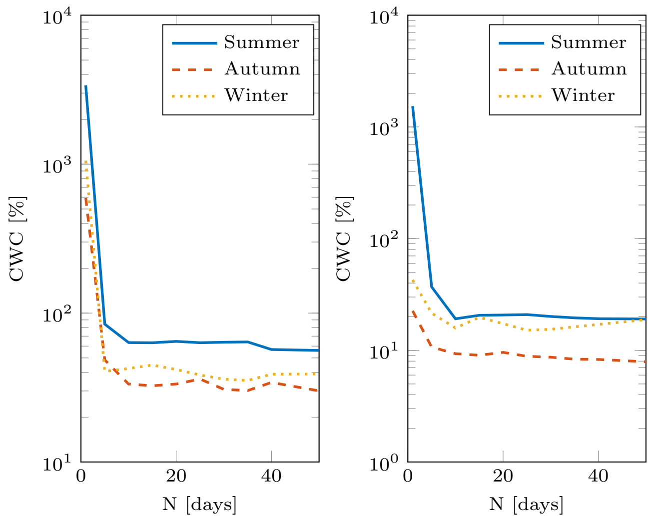

(a) Method A-5 min   (b) Method B-5 min

Figure 10: CWC as a function of the number od training days for 5 min time horizon. CWC is shown in logarithmic scale.

Table 3: CWC [%] for 500 ms, α=95%.

	Season		
Method	Summer	Autumn	Winter
Optimal	2.88	1.18	2.55
Silhouette	51	27.7	20.3
Exhaustive Search	4.69	2.41	4.03

(a) Method A

	Season		
Method	Summer	Autumn	Winter
Optimal	0.24	0.046	0.13
Silhouette	0.37	0.12	0.27
Exhaustive Search	0.27	0.047	0.15

(b) Method B

- For ultra-short term forecast, Method B outperforms Method A for each considered case.

- For both the methods, the Summer period is characterized by worse performance and this is explained by the highest volatility content, as shown in Table 2. As expected, the Autumn period returns the best performance.

5.4.2. Comparison with Benchmark Methods

Table 4 shows the performance of the proposed methods compared with the above described benchmarks. First, we refer to the simple quantiles extraction (Quantiles A and Quantiles B), where the quantiles are computed from the original and differentiated time series respectively, as described in the introduction of the results section. This comparison aims at showing the performance improving obtained by the k-mean clustering compared to the case where we extract the quantiles of the whole time series, without any clustering process. We can conclude that the k-means clustering is beneficial and leads to relevant performance improvement for all the analyzed cases.

The last row of Table 4 shows the results obtained by applying the ARIMA model and assuming a Gaussian distribution of the point forecast error (ARIMA + GAUSS). This method has to be compared with our proposed Method B since it requires a point forecast to compute the PIs. For 500 ms forecast horizon, the model is over confident with respect to the assumed normal distribution, returning PICP higher than 99% for α=95%. Thus, to allow a fair comparison with our method, we empirically adjust the target confidence level (and so η) in order to obtain values of PICP similar to those given by the k-mean algorithm. Table 4 shows that, for the same coverage probability, Method B is characterized by lower PINAW.

We refer to [13] as the only reference method for ultra-short term available in the literature. To allow a fair comparison, we compare Method B with the Dynamic Interval Predictor (DIP) coupled with the persistent point forecast, proposed in [13]. Indeed, the DIP needs to be coupled with a point forecast technique.

From Table 5 we can conclude that the proposed method shows better performance when compared to the literature with respect to ultra-short term horizons.

Table 4: PICP-PINAW-CWC [%] for a time horizon of 500 ms, α=95%.

	Season		
Method	Summer	Autumn	Winter
Method A	90.5-1.94-4.69	93.7-0.33-2.41	92.6-1.85-4.03
Quantiles A	94.6-57.4-113	93.3-29.5-62.6	95.7-35.1-35.1
Method B	97.0-0.27-0.27	96.1-0.047-0.047	98.2-0.15-0.15
Quantiles B	90.4-0.35-0.35	91.4-0.13-0.30	91.0-0.12-0.28
ARIMA+GAUSS	97.0-0.50-0.50	96.1-0.1-0.1	98.2-0.32-0.32

Table 5: PICP-PINAW-CWC [%].Performance comparison of the proposed Method B with the Dynamic Interval Predictor, [13]. α=95%.

	Season		
Method	Summer	Autumn	Winter
Method B	97.0-0.27-0.27	96.1-0.047-0.047	98.2-0.15-0.15
DIP	97.2-0.36-0.36	96.0-0.053-0.053	97.4-0.19-0.19

5.5. Short term forecasting

In this section, we extend the proposed methods to higher forecast horizons (i.e. minutes).

5.5.1. Exhaustive Search and Silhouette Analysis

Tables 6 and 7 show metric CWC obtained by applying the exhaust searching, the silhouette analysis and the optimum a-posteriori, for 1 and 5 minutes forecast horizons. The silhouette analysis coupled with Method B shows here the best performance and is used for further comparison. On the contrary, when using the original time series, the exhaustive search performs better and should be used to select k and N.

Table 6: CWC [%] for 1 min, α=95%.

	Season		
Method	Summer	Autumn	Winter
Optimal	21.7	8.50	22.9
Silhouette	58.0	31.6	33.0
Exhaustive Search	25.2	11.4	23.4

(a) Method A

	Season		
Method	Summer	Autumn	Winter
Optimal	6.82	2.83	6.20
Silhouette	10.5	3.26	9.10
Exhaustive Search	14.0	3.81	10.3

(b) Method B

Table 7: CWC [%] for 5 min, α=95%.

	Season		
Method	Summer	Autumn	Winter
Optimal	34.2	16.7	24.0
Silhouette	54.7	25.7	31.5
Exhaustive Search	37.9	20.7	24.0

(a) Method A

	Season		
Method	Summer	Autumn	Winter
Optimal	14.2	6.62	11.2
Silhouette	17.9	6.70	14.5
Exhaustive Search	16.2	15.1	23.4

(b) Method B

5.5.2. Comparison with Benchmark Methods

First, we present the advantage given by the k-means clustering for forecasting of 1 and 5 minutes ahead. Table 8 shows the comparison with the simple quantiles extraction for the original and differentiated time series (Quantiles A and B, respectively) for 1 minute forecast horizon, as explained for the ultra-short term analysis. Table 9 shows the same comparison for 5 minutes time horizon. For Method A, we apply the exhaustive search while for B we apply the silhouette analysis. For all the cases, we can see an improvement coming from the k-means clustering with respect to the simple quantile extraction.
The last row shows the results obtained by implementing the ARIMA model and assuming a Gaussian distribution of the point forecast error (ARIMA+GAUSS). For these forecast horizons, the PICP is slightly lower than the target confidence

level, with values of PINAW that are however higher than those returned by Method B.

Table 8: PICP-PINAW-CWC [%] for a time horizon of 1 minute, α=95%.

	Season		
Method	Summer	Autumn	Winter
Method A	90.1-10.2-25.2	90.8-4.81-11.4	88.8-8.87-23.4
Quantiles A	94.7-56.7- 112	93.1-29.5-61.9	95.8-34.7-34.7
Method B	96.9-10.5-10.5	97.5-3.26-3.26	97.8-9.1-9.1
Quantiles B	89.7-13.8-34.8	90.6-6.1-14.6	91.7-6.73-15.3
ARIMA+GAUSS	93.4-19.2-40.3	94.0-8.13-16.5	95.6-10.6-10.6

Table 9: PICP-PINAW-CWC [%] for a time horizon of 5 minutes, α=95%.

	Season		
Method	Summer	Autumn	Winter
Method A	91.5-16.5-37.9	86.7-6.96-20.7	96.1-24.0-24.0
Quantiles A	94.7-55.6-110	93.1-28.5-60.6	95.8-33.7-33.7
Method B	96.7-17.9-17.9	96.2-6.70-6.70	96.1-14.5-14.5
Quantiles B	89.4-25.9-66.4	89.5-13.6-34.0	91.2-16.0-36.0
ARIMA+GAUSS	91.7-28.9-65.8	92.3-14.1-31.4	95.0-19.0-19.0

It is difficult to compare the proposed method with results available in the literature due to the different GHI measurements (characterized by dissimilar climatology). Similar results are obtained in [21] for 5 minutes ahead GHI forecast, where a probability coverage of $\approx$ 95% and PINAW of $\approx$ 8%. In [21] a dataset of 1 year is considered. The percentage of periods of high volatility (i.e.

with ΔK higher than 0.5) is $\approx 0.3-0.6\%$ while in our datasets is $\approx 0.9-1.5\%$.

5.6. From Ultra-short to Short Term Forecasts

For sub-second time horizons, the best performance is obtained by applying the exhaustive search coupled with Method B. On the contrary, for higher forecast horizons, the silhouette analysis coupled with Method B is the most performing one.

For all the considered horizons, differentiating the time series has a positive effect on the final performance and allows to have a PICP higher than or equal to α. However, the improvement coming from the differentiation decreases with increasing forecast horizons. Indeed, performance of Method A worsens less than those of Method B when increasing the forecast horizon.

To complete the analysis, Fig. 11 shows our metrics as a function of the forecast horizon and for different confidence levels: 85%, 95%, and 99%.We can see that PICP (left side) is always higher or equal to the target confidence level. Furthermore, the value of PINAW (right side) increases with the forecast horizon (to account for the higher uncertainty) and increases with α, i.e. the method adapts the bound widths to guarantee the target coverage.

Fig. 12 shows the PIs and the actual realizations obtained for 500 ms, 1 and 5 minutes forecast horizons, respectively. Method B is applied and the target confidence level is 99%. A day of high variability and a clear-sky day from the Winter period are selected for the comparison. The corresponding values of PINAW are shown in Table 10, for the two days respectively, showing that the PIs are narrower for the clear-sky day where the variability is lower. For the clear-sky day at 500 ms a zoom is added since the PIs and the points are not easily distinguishable. We can see that larger intervals are associated to higher time horizons, this reflecting the higher uncertainty.

In Fig. 13 the same days of Fig. 12 are selected and PIs are plotted for different confidence level (99%, 95% and 85%). It is interesting to notice that for the considered clear-sky day a target confidence level of 85% is enough to have all the measurements inside the PI.

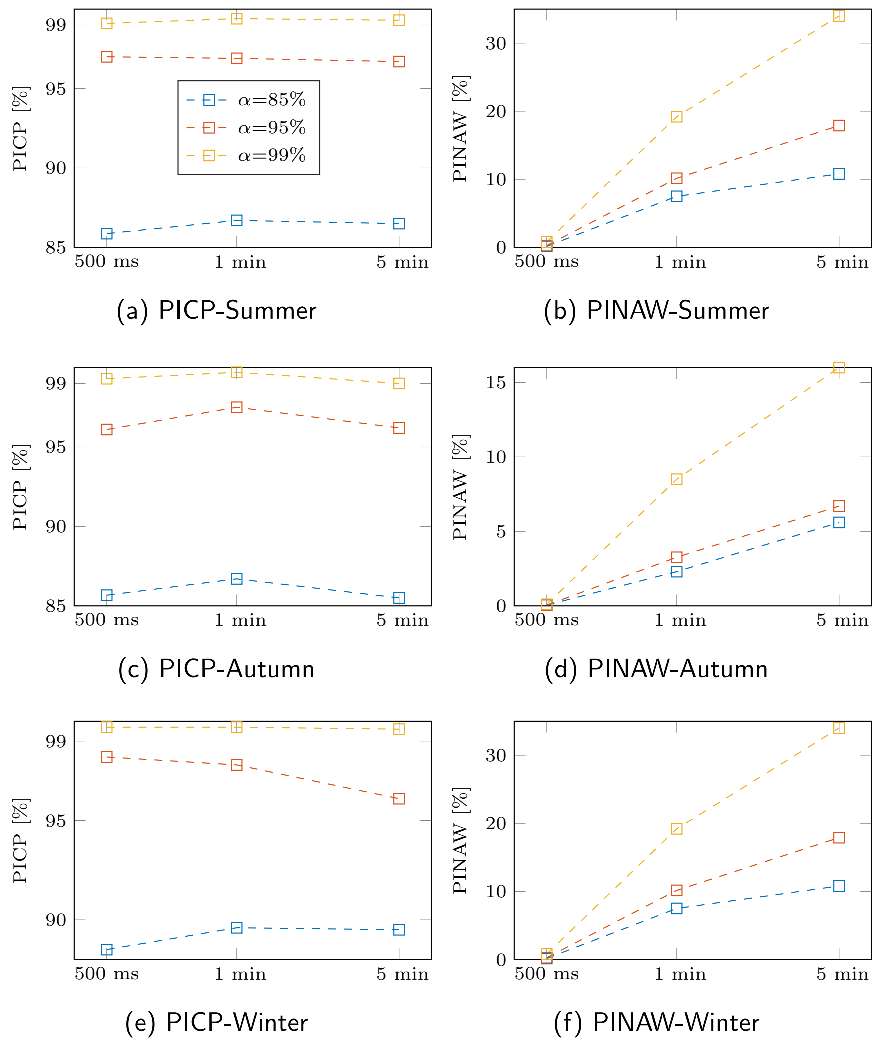

Figure 11: PICP [%] and PINAW [%] are shown for the Summer, Autumn, and Winter periods and different target confidence levels α.

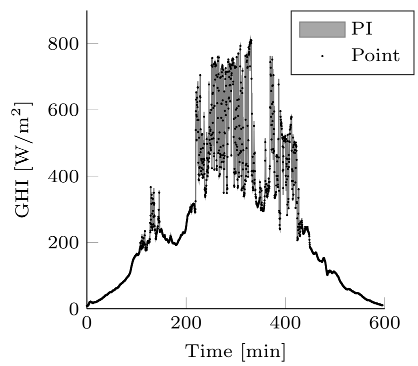

(a) High Volatility - 500 ms

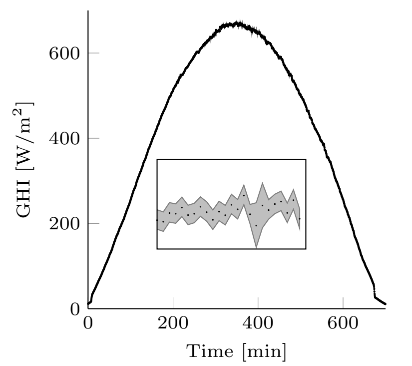

(b) Clear-sky - 500 ms

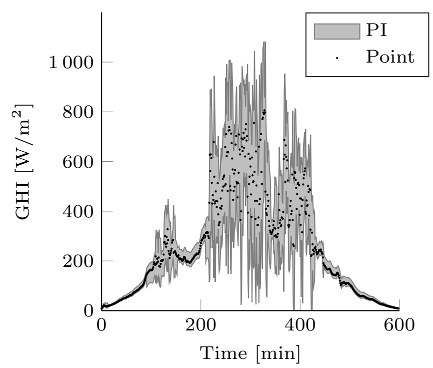

(c) High Volatility - 1 min

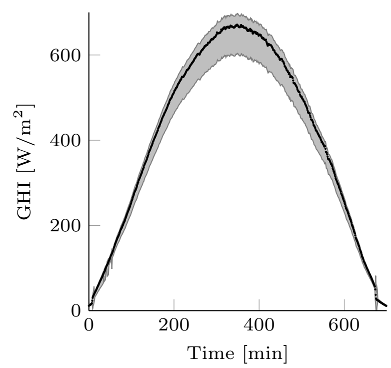

(d) Clear-sky - 1 min

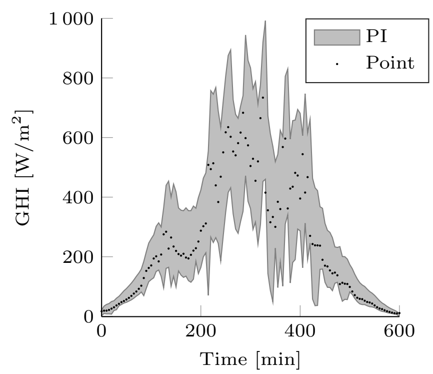

(e) High Volatility - 5 min

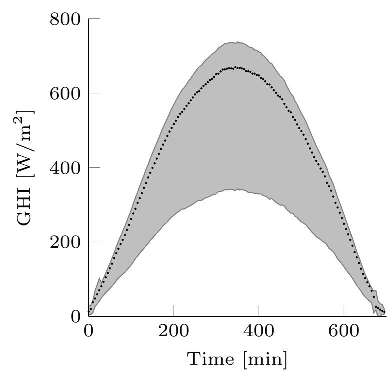

(f) Clear-sky - 5 min

Figure 12: PIs and realizations are shown for different forecast horizons considering daylight hours, $\alpha = 99\%$ and Method B is applied. Two days with different weather conditions are selected from the Winter period.

Table 10: PINAW [%] is shown for the two days of Fig. 12.

	Day	
Forecast Horizon	High Volatility	Clear-sky
500 ms	0.22	0.18
1 min	11.8	10.8
5 min	24.9	23.8

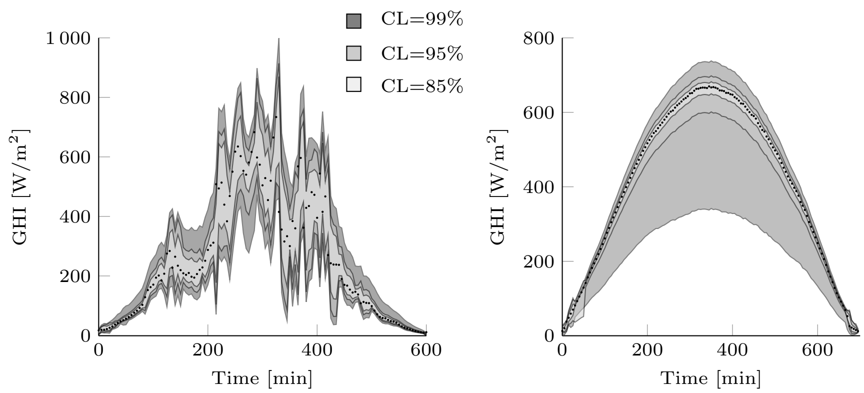

(a) High Volatility

(b) Clear-sky

Figure 13: PIs and realizations are shown for 5 minutes time horizon and different target confidence levels, represented by different levels of shadings for the same data of Fig. 12.

5.7. Reliability diagrams

The objective of this last analysis is to compare the target confidence levels with the observed ones, here represented by metric PICP.

The analysis is for the three considered periods and three forecast horizons and is shown in Fig. 14. We consider Method B and the ARIMA model with Gaussian distribution of the error and compare their performance with the ideal behaviour, namely when the target confidence level is identical to the observed one. As it can be seen, the confidence levels obtained for Method B exhibit an overall good matching with the target ones, proving the capability of the method to provide reliable predictions. In particular, Method B is always slightly over confident with the exception of the Autumn period where is under confident for low values of α. The proposed benchmark method has lower reliability, it is over confident for sub-second time horizons while its behaviour for higher horizons depends on the value of α. This mismatch suggests that parametric models, with the implicit assumption of a Gaussian distribution of the error, might not be suitable.

5.8. Execution Time Statistics

Execution times are computed adopting a Matlab 2016a implementation of the algorithm on an Intel Core i7-6600U CPU 2.60GHz machine. For the analysis, we consider the worst case scenario corresponding to the highest number of clusters. We select it equal to 1000 that is the maximum value of clusters returned by the analysis at 500 ms (above this value of k we do not see any performance improvement). At each time step, the overall operational time required to deliver the PI is always less than 0.5 ms. The mean and the standard deviation of the computational time at each time step are 0.35 ms and 1.14 ms, respectively. The method can be used for the real-time computation of PIs at sub-second time scales, and it is expected to run even faster if developed on a dedicated industrial platform and/or in a different programming language.

6. Conclusions

The problem of quantifying the uncertainty associated with solar volatility is investigated in this work, focusing on forecast horizons that are meaningful in microgrids control applications (i.e. from sub-second up to minutes).

A simple method to deliver PIs for GHI is proposed and its performance assessed. The proposed technique extracts information from a limited training set: data are clustered off-line by using the well-known k-means algorithm and the quantiles of the obtained clusters are then used for PIs computation. The method does not rely on any specific point forecast technique and does not need any information from sky imaging.

First, a clear-sky model is implemented. It is shown that the de-trending of the time-series is advantageous for time horizons higher than the minute time-scale, when the influence of the dynamics associated to solar position becomes

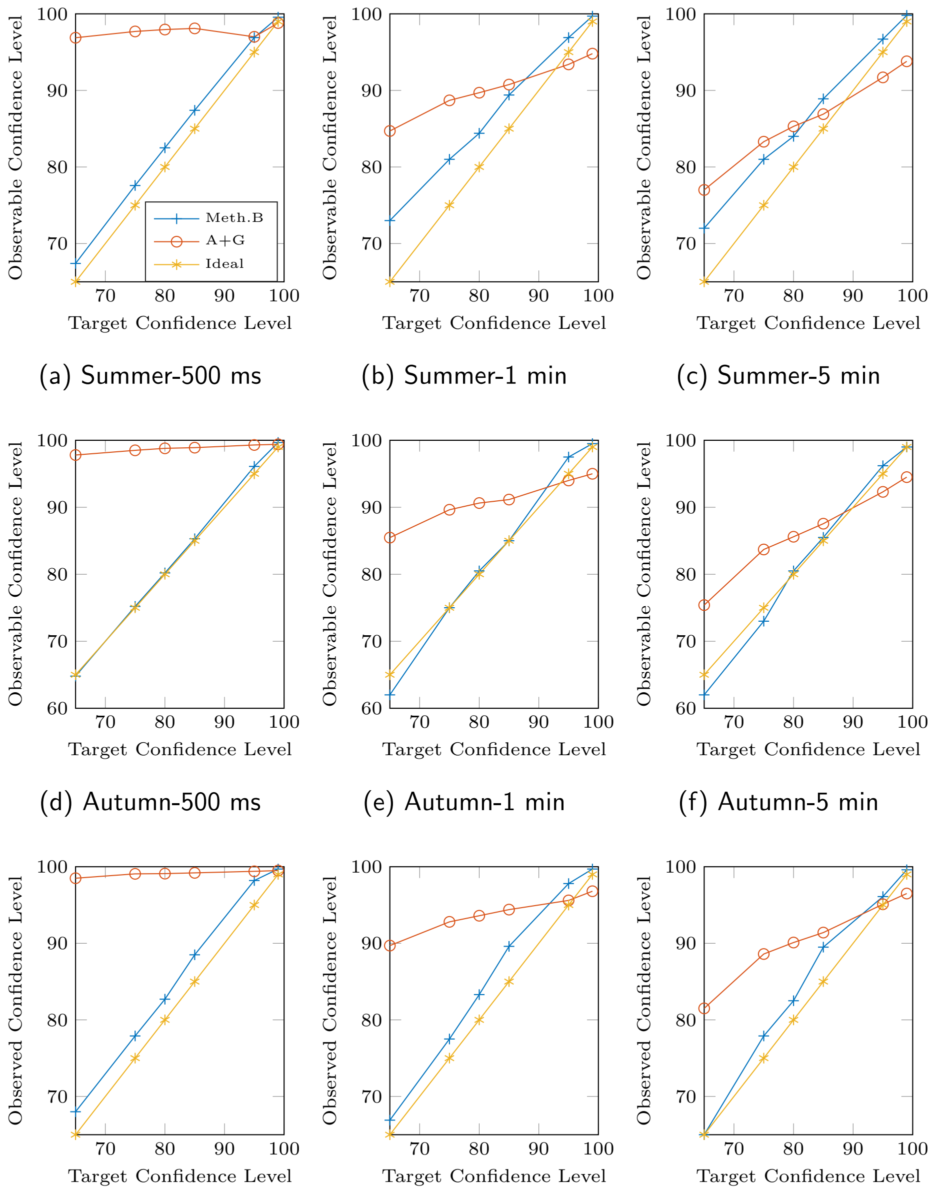

(a) Summer-500 ms (b) Summer-1 min (c) Summer-5 min

(d) Autumn-500 ms (e) Autumn-1 min (f) Autumn-5 min

(g) Winter-500 ms (h) Winter-1 min (i) Winter-5 min

Figure 14: Reliability Diagrams for the three periods and forecast horizons.

non-negligible.

We show that the algorithm outperforms the benchmark case with simple quantiles extractions and the benchmark case considering the ARIMA model with Gaussian distribution of the point forecast error. Furthermore, performance is shown to be in line or improve those available in the literature, for all the considered forecast horizons and using a shorter and limited training set. A comparison with more sophisticated methods available in literature will be part of future work.

The method is applied to the original and differentiated clear-sky index time series. Results show that the benefit coming from the time series differentiation decreases while increasing the forecast horizon.

It is shown that the proposed method is able to adapt the widths of the PIs in order to guarantee the target coverage.

Thanks to its simple formulation, computational inexpensiveness and good performance at different forecast horizons, the model can be useful for providing forecast information in the field of photovoltaic generation and in the context of real-time control of microgrids.

References

[1] J. P. Lopes, N. Hatziargyriou, J. Mutale, P. Djapic, N. Jenkins, Integrating distributed generation into electric power systems: A review of drivers, challenges and opportunities, Electric power systems research 77 (9) (2007) 1189–1203.

[2] D. E. Olivares, A. Mehrizi-Sani, A. H. Etemadi, C. Canizares, R. Iravani, M. Kazerani, A. H. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke, et al., Trends in microgrid control, Smart Grid, IEEE Transactions on 5 (4) (2014) 1905–1919.

[3] L. Xie, M. D. Ilic, Model predictive dispatch in electric energy systems with intermittent resources, in: Systems, Man and Cybernetics, 2008. SMC 2008. IEEE International Conference on, IEEE, 2008, pp. 42–47.

[4] Y. Zhang, N. Gatsis, G. Giannakis, Robust energy management for microgrids with high-penetration renewables, Sustainable Energy, IEEE Transactions on 4 (4) (2013) 944–953. doi:10.1109/TSTE.2013.2255135.

[5] F. Sossan, E. Namor, R. Cherkaoui, M. Paolone, Achieving the dispatchability of distribution feeders through prosumers data driven forecasting and model predictive control of electrochemical storage, arXiv preprint arXiv:1602.02265.

[6] N. Troy, D. Flynn, M. OMalley, Multi-mode operation of combined-cycle gas turbines with increasing wind penetration, Power Systems, IEEE Transactions on 27 (1) (2012) 484–492.

[7] A. Bernstein, L. Reyes-Chamorro, J.-Y. Le Boudec, M. Paolone, A Composable Method for Real-Time Control of Active Distribution Networks with Explicit Power Setpoints. Part I: Framework, Electric Power Systems Research 125 (2015) 254–264.

[8] L. Reyes-Chamorro, A. Bernstein, J.-Y. Le Boudec, M. Paolone, A Composable Method for Real-Ttime Control of Active Distribution Networks with Explicit Power Setpoints. Part II: Implementation and Validation, Electric Power Systems Research 125 (2015) 265–280.

[9] R. H. Inman, H. T. Pedro, C. F. Coimbra, Solar forecasting methods for renewable energy integration, Progress in Energy and Combustion Science 39 (6) (2013) 535 – 576. doi:http://dx.doi.org/10.1016/j.pecs.2013.06.002.
URL http://www.sciencedirect.com/science/article/pii/S0360128513000294

[10] A. Mellit, S. A. Kalogirou, Artificial intelligence techniques for photovoltaic applications: A review, Progress in Energy and Combustion Science 34 (5) (2008) 574 – 632. doi:http://dx.doi.org/10.1016/j.pecs.2008.01.001.

URL http://www.sciencedirect.com/science/article/pii/S0360128508000026

[11] C. W. Chow, B. Urquhart, M. Lave, A. Dominguez, J. Kleissl, J. Shields, B. Washom, Intra-hour forecasting with a total sky imager at the {UC} san diego solar energy testbed, Solar Energy 85 (11) (2011) 2881 – 2893. doi:http://dx.doi.org/10.1016/j.solener.2011.08.025. URL http://www.sciencedirect.com/science/article/pii/S0038092X11002982

[12] J. Calbo, J. Sabburg, Feature extraction from whole-sky ground-based images for cloud-type recognition, Journal of Atmospheric and Oceanic Technology 25 (1) (2008) 3–14.

[13] D. Torregrossa, J.-Y. Le Boudec, M. Paolone, Model-free computation of ultra-short-term prediction intervals of solar irradiance, Solar Energydoi:10.1016/j.solener.2015.11.017.

[14] A. Bracale, P. Caramia, G. Carpinelli, A. R. Di Fazio, G. Ferruzzi, A bayesian method for short-term probabilistic forecasting of photovoltaic generation in smart grid operation and control, Energies 6 (2) (2013) 733–747.

[15] M. Ghonima, B. Urquhart, C. Chow, J. Shields, A. Cazorla, J. Kleissl, A method for cloud detection and opacity classification based on ground based sky imagery, Atmospheric Measurement Techniques 5 (11) (2012) 2881–2892.

[16] Y. Chu, H. T. Pedro, C. F. Coimbra, Hybrid intra-hour {DNI} forecasts with sky image processing enhanced by stochastic learning, Solar Energy 98, Part C (2013) 592 – 603. doi:http://dx.doi.org/10.1016/j.solener.2013.10.020. URL http://www.sciencedirect.com/science/article/pii/S0038092X13004325

[17] R. Marquez, V. G. Gueorguiev, C. F. Coimbra, Forecasting of global horizontal irradiance using sky cover indices, Journal of Solar Energy Engineering 135 (1) (2013) 011017.

[18] S. Alessandrini, L. Delle Monache, S. Sperati, G. Cervone, An analog ensemble for short-term probabilistic solar power forecast, Applied Energy 157 (2015) 95–110.

[19] S. Sperati, S. Alessandrini, L. Delle Monache, An application of the ecmwf ensemble prediction system for short-term solar power forecasting, Solar Energy 133 (2016) 437–450.

[20] Y. Chu, M. Li, H. T. Pedro, C. F. Coimbra, Real-time prediction intervals for intra-hour dni forecasts, Renewable Energy 83 (2015) 234–244.

[21] H. T. Pedro, C. F. Coimbra, Nearest-neighbor methodology for prediction of intra-hour global horizontal and direct normal irradiances, Renewable Energy 80 (2015) 770 – 782. `doi:http://dx.doi.org/10.1016/j.renene.2015.02.061`.
URL `http://www.sciencedirect.com/science/article/pii/S0960148115001792`

[22] A. Grantham, Y. R. Gel, J. Boland, Nonparametric short-term probabilistic forecasting for solar radiation, Solar Energy 133 (2016) 465 – 475. `doi:http://dx.doi.org/10.1016/j.solener.2016.04.011`.
URL `http://www.sciencedirect.com/science/article/pii/S0038092X16300342`

[23] M. David, F. Ramahatana, P. Trombe, P. Lauret, Probabilistic forecasting of the solar irradiance with recursive {ARMA} and {GARCH} models, Solar Energy 133 (2016) 55 – 72. `doi:http://dx.doi.org/10.1016/j.solener.2016.03.064`.
URL `http://www.sciencedirect.com/science/article/pii/S0038092X16300172`

[24] S. P. Lloyd, Least squares quantization in pcm, Information Theory, IEEE Transactions on 28 (2) (1982) 129–137.

[25] P. Pinson, Estimation of the uncertainty in wind power forecasting, Ph.D. thesis, École Nationale Supérieure des Mines de Paris (2006).

[26] GRASS Development Team, Geographic Resources Analysis Support System (GRASS GIS) Software, Open Source Geospatial Foundation, USA (2015).
URL http://grass.osgeo.org

[27] J. Hofierka, M. Suri, et al., The solar radiation model for open source gis: implementation and applications, in: Proceedings of the Open source GIS-GRASS users conference, 2002, pp. 1–19.

[28] R. Marquez, C. F. Coimbra, Proposed metric for evaluation of solar forecasting models, Journal of solar energy engineering 135 (1) (2013) 011016.

[29] L. Kaufman, P. J. Rousseeuw, Finding groups in data: an introduction to cluster analysis, Vol. 344, John Wiley & Sons, 2009.

[30] F. O. of Metereology, C. MateoSwiss. [link].
URL https://gate.meteoswiss.ch/idaweb

[31] J. D. Hamilton, Time series analysis, Vol. 2.

[32] A. Khosravi, S. Nahavandi, D. Creighton, Prediction intervals for short-term wind farm power generation forecasts, Sustainable Energy, IEEE Transactions on 4 (3) (2013) 602–610.