

Computation of Ultra-Short-Term Prediction Intervals of the Power Prosumption in Active Distribution Networks

Plouton Grammatikos
Jean-Yves Le Boudec
Mario Paolone
School of Engineering
EPFL
Lausanne, Switzerland

{plouton.grammatikos, jean-yves.leboudec, mario.paolone}@epfl.ch

Fabrizio Sossan
School of Engineering
HES-SO Valais - Wallis
Sion, Switzerland
fabrizio.sossan@hevs.ch

Abstract—Microgrids and, in general, active distribution networks require ultra-short-term prediction, i.e., for sub-second time scales, for specific control decisions. Conventional forecasting methodologies are not effective at such time scales. To address this issue, we propose a non-parametric method for computing ultra short-term prediction intervals (PIs) of the power prosumption of generic electrical-distribution networks. The method groups historical observations into clusters according to the values of influential variables. It is applied either to the original or to the differentiated power-prosumption time series. The clusters are considered statistically representative pools of future realizations of power prosumption (or its derivative). They are used to determine empirical PDFs and, by extracting the quantiles, to deliver PIs for respective arbitrary confidence levels. The models are validated a posteriori by carrying out a performance analysis that uses experimentally observed power-prosumption for different building types, thus allowing the identification of the dominant model.

Index Terms—prosumption, forecast, prediction intervals, electrical load, microgrids.

I. INTRODUCTION

After being the mainstream framework for the integration and coordination of distributed generation, the concepts of an active distribution network (ADN) and a microgrid recently came to prominence to tackle the challenges caused by the large-scale integration of variable renewable generation. ADNs comprise low-voltage (LV) or medium-voltage (MV) electrical grids with systems in place to control a combination of distributed energy resources (DERs), such as generators, loads, and storage devices [1].

This research is carried out in the frame of the “Optimal integration of electric vehicles fast charging stations into medium voltage power distribution grids” project, with the financial support of the Swiss Federal Office of Energy (SFOE Pilot and Demonstration Program) (SI/502045-01) as a contribution to the ERA-NET project, “MESH4U - Multi Energy Storage Hub For reliable and commercial systems Utilization) (MICall19 - 100522).

Due to the low level of aggregation, the ADN requirements for electrical-power prosumption¹ forecasting are different than for conventional large interconnected grids. One example relates to the possible violation of the ampacity rating of transformers, power converters, and lines due to sudden changes in the prosumption associated with the highly stochastic nature of prosumers². A significant change in the prosumers’ renewable-power generation, or a spike in load, can create power-flow variations that can exceed the rating of transformers. Whereas, a spike in current could cause the line relays to trip. This example is particularly relevant in the presence of photovoltaics (PV) and electrical-vehicle (EV) charging stations (CS) in the grid (e.g., [2]). The former can exhibit power variations of even 60% of their capacity in under a second [3], whereas the latter can cause significant load changes of hundreds of kW within a few seconds [4].

The prediction of the prosumption can be integrated into various real-time (RT) (e.g., [5]) and model-predictive control (MPC) (e.g., [6]) frameworks to ensure the safe operation of the ADNs and the optimal usage of the hosted resources. For example, in a typical MPC, the goal is to minimize a given objective over a certain horizon, which could be as short as 30 s in [6], subject to the grid operational constraints. The problem is solved repeatedly whenever updated forecasts of the system’s uncertainties become available. Inaccuracies in these forecasts might lead to violating the grid’s constraints.

Even shorter prediction horizons are needed when a microgrid needs to switch to operating as a standalone system. When connected to the external grid, microgrids can provide ancillary services to the upper grid layer [7]. Meanwhile, in case of contingencies, they can operate islanded to the main grid enhancing the resiliency of the supply to the local load. The islanding maneuver (i.e., the operation sequence for

¹In power systems, the prosumption indicates the aggregated power provided or consumed by users that have the capability to generate electricity by means of user-owned distributed generation locally.

²A node that can both absorb or inject power due to prosumption.

bringing a microgrid from connected to off-grid) needs to be performed by a myopic control which utilizes predictions over a short horizon in the range of the fundamental frequency period (e.g., 20 ms for a 50 Hz system)) to prevent steering the system to an undesirable state. This can be achieved, for example, by correctly setting the gains of the slack resource droop control (e.g., [8]).

A third example, which is of high importance both in distribution and transmission networks, is that of voltage sags [9], [10]. Voltage sags are defined as a sudden reduction of the voltage between 90% and 10% of the nominal value and can last from 10 ms up to 1 minute. They are caused primarily by power-system faults, such as short circuits, or by the start-up of large motors and can cause system outages if not treated in time. For distribution systems, where voltage sags typically last between 90-2000 ms [11], in order to mitigate the voltage drop, a control framework equipped with a forecasting tool acting in the sub-second range could have a timely reaction to the voltage sag by injecting an optimally computed active/reactive power into the grid.

As ultra-short-term power-prosumption forecasts are actionable for fundamental decisions in the context of ADN/microgrid operation and their RT control, we note that well-established forecasting methodologies (e.g., developed using several techniques, such as regression-based model, artificial neural network [12]–[19]) are not suited to this purpose because, besides referring to point predictions, they were developed considering a high level of aggregation and forecasting horizon from 15 minutes and up. Moreover, in certain applications, such as robust optimization (e.g., [20]), worst-case analysis is required; therefore, point predictions are inadequate.

As stated in [21] and further supported by the experimental measurements of this paper, when decreasing the aggregation level and measurements sampling time, the power-prosumption volatility and noise level become prominent because consumer behaviors tend not to cancel out.

As the current state of the art appears inadequate to deliver ultra-short-term power-prosumption prediction intervals (PIs) of ADN prosumers, we propose an adaptive non-parametric method based on pattern recognition. Indeed, it is reasonable to assume that for ultra-short-term horizons and low aggregation level, the forecast distribution will change model over time. Therefore, the adoption of a parametric forecasting model may be inappropriate. For this reason, we opted to empirically learn the cdf, bypassing this problem. The algorithm is designed to be computationally efficient, thus allowing for the delivery of high-time resolution probabilistic PIs in RT and at a high sampling rate with low computational overhead. The model is initially trained using a time series of the aggregated-power prosumption without requiring any knowledge of the nature and number of loads/generators present in the network. With an efficient updating and aging procedure, it is then continuously updated as new measurements become available.

The paper is organized as follows. In Section II, the problem of estimating the power prosumption PI is stated, along with a

review of methods already developed in the context of power system applications. In Section III, we describe the proposed PI models whose performance is analyzed in Section IV by using experimental data for different building types. Finally, in Section V, we summarize the findings.

II. PROBLEM STATEMENT

As stated in the previous section, most of the existing literature on power-prosumption forecasting is concerned with point predictions, specifically the problem of estimating the expected realization of the power prosumption for a given look-ahead time. Whereas, we target the computation of PIs; in other words, we predict, with a given confidence level, the interval where the future power-prosumption realization is expected to lie. Denoting the PI at the target confidence level α as the couple $(P^{\downarrow\alpha}, P^{\uparrow\alpha})$ composed by the lower and upper bound of the interval, we address the problem that consists in finding the one-step-ahead PI as a function of a sequence of n historical power-prosumption measurements until the time instant i , specifically:

$$(P_{i+1|i}^{\downarrow\alpha}, P_{i+1|i}^{\uparrow\alpha}) = f(P_i, \dots, P_{i+1-n}) \quad (1)$$

where i is the current time interval, and f is a PI estimation model.

When using parametric point predictors (such as autoregressive integrated moving average (ARIMA) models), we can determine the PIs by estimating the variance of the model residuals and computing the quantiles for the prescribed confidence level. This procedure can be performed under the hypothesis of Gaussian *iid* (independent and identically distributed) model residuals.

In cases where this hypothesis does not hold, non-parametric methods could be considered. For example, in order to determine the PIs of the power output of a wind farm, the authors of [22] apply quantile regression to characterize the historical residuals of a state-of-the-art point-prediction model.

The same concept, but developed using fuzzy inference instead of quantile regression, is described in [23]. As far as forecasting the electrical-power prosumption is concerned, non-parametric methods have been proposed lately in [24] and [25]. In the former work, an artificial neural network (ANN) with an empirically chosen number of layers was trained using historical data to provide a 30-minute-ahead PI for a given confidence level according to the values of selected data features. The latter work is concerned with predicting the minimum and maximum bounds of the power consumption by applying empirical-mode decomposition and support-vector regression to an interval-valued time signal obtained from a one-hour historical sample of power consumption. Both methods target a prediction horizon that is too long for the requirements of RT ADN/microgrid operation discussed in the introduction.

Even though ANN-based algorithms could be used for computing PIs at shorter time scales, to the best of our knowledge, there are no works using these methods in resolutions as small as the ones targeted by our method, i.e., 20 ms.

III. COMPUTATION OF PREDICTION INTERVALS

As it will be exhaustively described in the rest of this section, the estimation methodology for PIs consists in grouping historical power-prosumption measurements into clusters according to the value of selected influential variables. At the time of delivering a PI, the values of the influential variables are determined, thus allowing for the selection of the appropriate cluster that is finally used as the empirical PDF (probability distribution function) of the next realization. The algorithm is designed to deliver a PI in rolling RT with a minimum report rate of 20ms (longer values are also analyzed).

The algorithm operation sequence is sketched in Fig. 1. The first phase, called batch training, consists of the off-line training of the estimation model that uses historical data. In the second phase, the one-step-ahead PI is delivered. And finally, in the third phase, the new progressively available measurements are used for the on-line model training. This methodology is applied in two flavors, specifically on the original and differentiated time-series, as discussed in the following two sections.

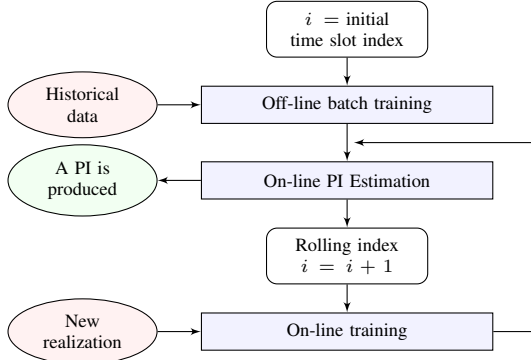


Fig. 1: Operation sequence of the PI estimation models. The batch training is performed off-line, whereas the PIs computation and on-line training are performed in rolling RT. Ellipses denote the input and output of each phase.

A. PI Model A

1) *Off-Line Batch Training*: We consider n historical power-prosumption measurements P_{i+1-n}, \dots, P_i and the respective time stamp t_{i+1-n}, \dots, t_i . Power-prosumption measurements are discretized in time and amplitude with user-defined discretization steps ΔT and ΔP , respectively. We define c as the operator that determines a non-negative integer l , said *label*, according to the power-prosumption value and timestamp in the following way:

$$l_j = c(P_j, t_j), \quad l_j \in \{0, 1, \dots, L-1\}. \quad (2)$$

where j is a generic time slot, and L is the total number of labels. The process used to operate the classification c is described in detail in Section III-C. The label in (2) is utilized to group the historical power-prosumption measurements $\mathcal{P}_{i,n}$ into L clusters, denoted as $\mathcal{G}_i^0, \dots, \mathcal{G}_i^{L-1}$. Each cluster contains all the historical power-prosumption measurements for which the previous observation was of the respective given

label. For example, the cluster \mathcal{G}_i^0 contains the measurements until the time slot i for which the respective previous realization was with label 0, \mathcal{G}_i^1 those for which the respective previous realization was with label 1, and so on. Formally, the clusters are defined as:

$$\mathbb{G}_i^l = \{P_{j+1} : c(P_j, t_j) = l, j = i-n, \dots, i-1\}, \quad (3)$$

$$l = 0, \dots, L-1.$$

Let $\mathcal{G}_i^0, \dots, \mathcal{G}_i^{L-1}$ be the normalized histogram of each cluster computed as

$$\mathcal{G}_i^l(x) = \frac{1}{|\mathbb{G}_i^l|} \sum_{p \in \mathbb{G}_i^l} \delta(x-p), \quad l = 0, \dots, L-1 \quad (4)$$

where $|\cdot|$ denotes the set cardinality (i.e., the number of elements it contains) and δ is the Dirac measure:

$$\delta(x) = \begin{cases} 1 & x = 0 \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

As the power prosumption is bounded, say between P_{\min} and P_{\max} , histograms in (4) are defined over a finite domain. Specifically, the domain is as:

$$x \in \mathbb{X} = \{P_{\min}, P_{\min} + \Delta P, P_{\min} + 2\Delta P, \dots, P_{\max}\}. \quad (6)$$

The value of ΔP is chosen as a trade-off between accuracy and computational efficiency. Indeed, the smaller the ΔP is, the more accurate the prediction of the PI is. This aspect is made clear below. However, choosing a very small step will require more memory to store all the measurements, thus resulting in a slower computation of the PI.

2) *On-Line PI Estimation*: At time i , the objective of the PI estimator is to determine the PI for the time slot $i+1$ at a given arbitrary target confidence level, said α . The underlying idea is to assume the clusters (3) as a statistically representative pool of possible realizations of the one-step-ahead power-prosumption realization. Therefore, the normalized histograms (4) are assumed to be discrete PDFs and used to extract the symmetric quantiles corresponding to the α confidence level. Let

$$l_i = c(P_i, t_i) \quad (7)$$

be the label calculated with the information at the current time instant. The PI lower and upper bounds are determined as

$$P_{i+1|i}^{\downarrow\alpha} = (1-\alpha)/2 \text{ quantile of } \mathcal{G}_i^{l_i}, \quad (8)$$

$$P_{i+1|i}^{\uparrow\alpha} = (1+\alpha)/2 \text{ quantile of } \mathcal{G}_i^{l_i}. \quad (9)$$

For the sake of clarity, the lower and upper quantiles in the expressions above are approximated by, respectively:

$$\inf_{x \in \mathbb{X}} \{x : F(x) \geq (1-\alpha)/2\} \quad (10)$$

and

$$\sup_{x \in \mathbb{X}} \{x : F(x) \leq (1+\alpha)/2\}, \quad (11)$$

where F denotes the discrete CDF (cumulative distribution function) of $\mathcal{G}_i^{l_i}$ calculated by computing its cumulative sum.




Fig. 2: Exemplification of the procedure (10)-(11) to approximate the quantiles for (8)-(9) for a normal distribution
. In this case, the target confidence level α is 80%.

An example of the PI estimation is shown in Fig. 2 for a normal distribution.

There are two advantages to this approach. First, histograms carry the complete information over the empirical PDFs, thus allowing for computing PIs at arbitrary confidence levels by training only one model. Therefore, we do not need to train another model if the performance requirements change. Second, it overcomes the problem of quantile crossing that arises, for example, in [22], from treating the bounds at a given confidence level as two different time series. As a matter of fact, quantile crossing produces a mistake in the model that has to be handled by a suitable correction mechanism in the controller. However, in time-critical applications where high reliability is desired, such as the islanding maneuver, the controller needs to ensure the validity of its decisions and should not make any assumptions about the causes of the quantile crossing.

3) *On-Line Training*: As time passes, new measurements become available and can be included to improve future PI estimates. Once the outcome P_{i+1} is known, the normalized histogram associated with the label l_i is updated with the new information, and the other histograms stay the same. Formally, the training procedure is as follows:

$$\mathcal{G}_{i+1}^l(x) = \begin{cases} \phi \mathcal{G}_i^l(x) + (1 - \phi) \delta(x - P_{i+1}) & l = l_i \\ \mathcal{G}_i^l(x) & l \neq l_i, \end{cases} \quad (12)$$

where

$$\phi = \frac{T_\phi / T}{T_\phi / T + 1} \quad (13)$$

is called the *forgetting factor*, T is the measurement period and T_ϕ is called the *forgetting time constant*. The forgetting factor controls how much past measurements influence the computation of PIs. Specifically, each new measurement has the same weight in the computation as all the measurements in the past T_ϕ seconds. The adoption of such a forgetting factor is important in order to track changes in the composition of prosumers' load/generation patterns.

B. PI Model B

1) *Off-Line Batch Training* : We apply the same principles described for Model A but on the once differentiated power-prosumption training data-set. The differentiated time series is denoted as:

$$B_j = P_j - P_{j-1}, \quad j = i - n + 1, \dots, i. \quad (14)$$

The observations clusters are now calculated as follows:

$$\mathbb{H}_i^l = \{B_{j+1} : c(P_j, t_j) = l, j = i - n + 1, \dots, i - 1\}, \\ l = 0, \dots, L - 1, \quad (15)$$

and are used to determine the normalized histograms of the differentiated power prosumption time series, which are denoted as $\mathcal{H}_i^0, \dots, \mathcal{H}_i^{L-1}$.

2) *On-Line PI Estimation*: The one-step-ahead PI bounds are computed as:

$$P_{i+1|}^{\downarrow \alpha} = P_i + (1 - \alpha)/2 \text{ quantile of } \mathcal{H}_i^{l_i}, \quad (16)$$

$$P_{i+1|}^{\uparrow \alpha} = P_i + (1 + \alpha)/2 \text{ quantile of } \mathcal{H}_i^{l_i}, \quad (17)$$

i.e., the current power-prosumption plus two back-off terms representing the expected power prosumption variation with respect to the current realization P_i .

3) *On-Line Training*: Once the prosumption P_{i+1} is known and the power difference B_{i+1} is computed, the normalized histogram corresponding to the current label l_i is updated by adding the new differenced value, whereas the others stay the same, i.e.

$$\mathcal{H}_{i+1}^l(x) = \begin{cases} \phi \mathcal{H}_i^l(x) + (1 - \phi) \delta(x - B_{i+1}) & l = l_i \\ \mathcal{H}_i^l(x) & l \neq l_i. \end{cases} \quad (18)$$

C. Classification According to Influential Variables

The assignment performed through the function $c(\cdot)$ in (2) is realized by first clustering the historical measurements into groups based on similarities with respect to chosen influential variables and, then, assigning a cluster label to each new measurement. Needless to say, influential (or explanatory) variables are quantities that have an influence on the power prosumption. In general, they can be discovered using numerical methods (like analysis of variance, correlation analysis, or other procedures [26]) or identified by exploiting any empirical knowledge on the observed process (in our case the structure of prosumption of a given node).




Fig. 3: Visualization of k-means clustering for a node with an EV charging station with 8 clusters and using the power level and the time of day as influential variables. Positive power is consumption.

In this work, influential variables are chosen by adopting the latter approach. The chosen variables are (i) power magnitudes and (ii) time of day (in seconds). The first variable accounts that, in a limited capacity feeder, prosumption variations depend on the same power-prosumption magnitude. For example, when the consumption is large, a load disconnection is more likely than a load insertion because many loads are already active, and vice-versa. The second variable is supposed to capture the different power-prosumption patterns that might occur during the day.

The historical measurements are grouped, by using a k-means approach, into L clusters, according to the values of the influential variables. An example of the clustering method is shown on Fig. 3 for a node with an office building and a level-3 EV charging station. The k-means algorithm was trained with two weeks of historical measurements by using both the power level and the time of day as influential variables and by grouping the measurements into $L = 8$ clusters, each of which is labeled with a different color in the figure. In this example, the measurements of the node when there are no connected EVs, which is the most frequent case, have been split into 6 clusters, dividing the day into 6 time periods. The rare occasions of an EV connection, which is characterized by a total consumption larger than 50kW, have been split into two clusters depending on the power level (50-100kW and 100-200kW approximately). The value of L is chosen by the user and is fixed a priori. Its influence on the algorithm's performance is studied in Section IV.

During the off-line and on-line training phases of the algorithm, the label of measurement (P_i, t_i) is computed as follows:

$$l_i = \arg \min_{l=0, \dots, L-1} d(e(P_i, t_i), c_l) \quad (19)$$

where $e(P_i, t_i)$ is the point in the $m - d$ space whose coordinates are the m influential variables chosen in the classification scheme, c_l is the center of cluster l and $d(\cdot, \cdot)$ is the Euclidean distance between two points.

The different classification schemes are introduced with the objective of performing an a-posteriori validation of the selection process of influential variables and the number of clusters. Indeed, as stated earlier in this section, influential variables are assigned by exploiting the empirical knowledge of the process: By comparing the performance of different classification schemes (in Section IV), it is possible to infer whether the progressively more complex classification schemes are meaningful or not.

D. Implementation Aspects and Complexity

The main design requirement of the proposed algorithm is to deliver PIs in RT in order to, for example, activate enough capacity in inertia-less microgrids or to assist in the decision process of setting the droop controller of slack generators in the islanding maneuver [8]. Given the large PIs reporting rate, computational complexity is a central aspect and hence is addressed in this section. The batch-training phase does not have RT requirements because it is performed off-line. This

consists in labeling each observation of the training data-set by applying the discussed classification algorithm. As the number of labels L is fixed by design, the overall complexity of the classification procedure, given by Eq. (19), for one observation is constant time, or $O(1)$. Iterating it over a set of N training data is an operation with linear time complexity, or $O(N)$.

The computation of the PIs and the on-line training are performed in rolling RT. The former operation requires computing the label in (7), which is $O(1)$, and the PI bounds by (9)-(11), which involve a minimum and maximum search over the set \mathbb{X} , a problem with log time complexity with respect to the set cardinality, $O(\log_2 |\mathbb{X}|)$, e.g., using a binary search that can exploit the monotonicity of the discrete CDFs. However, as the cardinality of \mathbb{X} is fixed by design in (6), the complexity of the problem can be regarded as constant time. Therefore, delivering PIs and performing on-line training are procedures whose complexities do not scale with the size of the problem.

The training data and progressively incoming measurements are encoded in L normalized histograms. Each of them is stored using $2 \times |\mathbb{X}|$ doubles, specifically the height and value of each bin. For example, assuming a discretization of 1024 levels (10 bit), preserving the information for a 1 year at 20 ms of resolution with the proposed method requires 128 kB per label (considering a double representation of 64-bit), while storing the individual values would require approximately 14 Gb.

IV. PERFORMANCE EVALUATION

A. A real case application: university buildings

To test the performance of the proposed models, we consider four sequences of power-prosumption measurements that were recorded from different points inside the MV, 21 kV distribution network of the EPFL campus.

The first three prosumers are very different regarding maximum power and power variations and were chosen to test our algorithm in the most challenging case for forecasting models: the lowest aggregation level. The first one consists of an office building with a maximum consumption of 80 kW and is equipped with a 30 kVA roof PV; the second one includes an office load with a maximum consumption of 30 kW and a 150 kW level 3 EV charging station; and the third one is a heat pump with a maximum consumption of 3 MW. Finally, to evaluate the algorithm at a higher aggregation level, we have considered the measurements on a primary substation supplying 8 MV feeders interfacing 20 heterogeneous prosumers, which comprise about half of the EPFL campus. The prosumption includes multiple office buildings, 1 MWp of PV production, and two heat pumps with a total maximum consumption of 4 MW. Even though the rated power of the primary substation transformer is 40 MVA, the maximum active power flow observed in the concerned period was 8 MW, which is the value we consider as P_{\max} in the implementation of the algorithm (Eq. (6)).

The measurements are with a resolution of 20 ms and are provided by a PMU-based metering infrastructure that has been deployed on the university campus (see [27]). For the

first three prosumers, we consider 45 days of historical power-prosumption measurements that span the period of September-October 2022, whereas for the MV feeders, we consider 30 days in January-February 2024. In each case, we consider two weeks of training data; they are used to construct the clusters (as explained in Section III-C) and to perform the off-line training of the algorithm. Then, the proposed PI estimation models are operated for the rest of the available measurements (with on-line training), and the estimated PIs are validated against the latter data-set, at 20 ms resolution. Each month consists of approximately 130 million data points. The evaluation is performed in a simulated environment coded in C++. The simulations are executed in a Windows Server with 128GB RAM and an Intel Xeon Gold 6130 CPU at 2.10GHz.

B. Performance Metrics

We introduce the following metrics to allow for a quantitative comparison between the performance of models and classification schemes. The first is the PI normalized averaged width (PINAW), which is as follows:

$$\text{PINAW} = \frac{1}{N} \sum_{j=1}^N (P_j^{\uparrow\alpha} - P_j^{\downarrow\alpha}) / P_{\text{nom}}. \quad (20)$$

The second metric is the PI coverage probability (PICP), i.e. the percentage of power-prosumption realization that falls inside the predicted PI. It is as follows:

$$\text{PICP} = \frac{1}{N} \sum_{j=1}^N b_j^\alpha \quad (21)$$

where

$$b_j^\alpha = \begin{cases} 1, & P_j^{\downarrow\alpha} \leq P_j \leq P_j^{\uparrow\alpha} \\ 0, & \text{otherwise.} \end{cases} \quad (22)$$

Because there is a trade-off between the width of the PI and the accuracy of the model, it is imperative that we define a third metric to quantify it. The metric we chose for this purpose is a modification of the coverage width-based criterion (CWC) proposed in [28] and is defined as follows:

$$\text{CWC} = \text{PINAW} \max(1, e^{-\mu \frac{\text{PICP} - \alpha}{1 - \alpha}}) \quad (23)$$

where μ is a user-defined parameter that quantifies the trade-off between PICP and PINAW. For our experiments, we chose $\mu = \frac{\log(10)}{10}$. This means that a deviation of one order of magnitude in the error rate penalizes the width of the interval by 10 times. The same result can be achieved with a linear penalty with slope 0.9, as can be seen in Fig. 4. The x -axis is the percentage deviation of the error rate from the target error rate, whereas the y -axis is the ratio $\frac{\text{CWC}}{\text{PINAW}}$. We observe that the exponential function penalizes less severely the error rates within one order of magnitude in comparison to the linear function. However, with our choice, higher error rates are more penalized.




Fig. 4: Visual representation of the CWC. $\text{error} = 1 - \text{PICP}$ is the error rate of the algorithm, and $\text{target} = 1 - \alpha$ is the target error rate. The red line indicates Eq. (23) with the chosen value of μ , and the blue line indicates a linear penalty.

TABLE I
PERFORMANCE EVALUATION FOR THE OFFICE BUILDING AT 20ms
RESOLUTION (3 BEST CONFIGURATIONS)

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	8	604800	0.021477	0.99129	0.021477
B	1	604800	0.021479	0.99021	0.021479
B	256	86400	0.021582	0.9916	0.021582

(a) $\alpha = 0.99$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	1	604800	0.030539	0.99909	0.030539
B	8	604800	0.030844	0.99913	0.030844
B	1024	21600	0.031535	0.99924	0.031535

(b) $\alpha = 0.999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	1	604800	0.0495	0.99978	0.11225
B	8	604800	0.050864	0.99976	0.11787
B	256	86400	0.051083	0.99976	0.11876

(c) $\alpha = 0.9999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	1	604800	0.097604	0.99999	0.19928
B	8	604800	0.087544	0.99996	0.2362
B	64	86400	0.086138	0.99996	0.25202

(d) $\alpha = 0.99999$

C. Clustering Using Power Levels

To evaluate the algorithm, we consider four target confidence levels α , namely 99, 99.9, 99.99, and 99.999%. The parameters to be analyzed are (i) the model A or B, (ii) the number of clusters L , and (iii) the forgetting factor ϕ or, equivalently, the forgetting constant T_ϕ . Each combination of model, L and T_ϕ is called a *configuration*. To choose the best configuration for each building, we run the algorithm for both models A and B, with all possible combinations of values of L in the set $\{1, 8, 64, 256, 512, 1024\}$ and T_ϕ in the set $\{1, 60, 3600, 21600, 86400, 604800\}$ seconds³. For all the experiments, the value of the quantization step ΔP was

³The values correspond to 1 second, 1 minute, 1 hour, 6 hours, 1 day and 1 week respectively.

TABLE II
PERFORMANCE EVALUATION FOR THE CHARGING STATION AT 20ms
RESOLUTION (3 BEST CONFIGURATIONS)

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	8	86400	0.0070666	0.99068	0.0070666
B	64	86400	0.0078829	0.99058	0.0078829
B	256	21600	0.0081288	0.99077	0.0081288

(a) $\alpha = 0.99$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	8	86400	0.019295	0.99868	0.039879
B	1	86400	0.019835	0.9988	0.040468
B	256	86400	0.019518	0.99832	0.041887

(b) $\alpha = 0.999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	8	86400	0.032019	0.99958	0.093142
B	64	86400	0.03352	0.99956	0.099959
B	256	21600	0.033447	0.99953	0.10298

(c) $\alpha = 0.9999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	8	86400	0.051835	0.99992	0.27484
B	64	86400	0.054395	0.9999	0.35737
B	64	21600	0.053203	0.9999	0.3654

(d) $\alpha = 0.99999$

TABLE III
PERFORMANCE EVALUATION FOR THE HEAT PUMP AT 20ms
RESOLUTION (3 BEST CONFIGURATIONS)

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	1	60	0.0094042	0.98832	0.019179
A	1	1	0.010129	0.98714	0.020948
B	64	60	0.0092479	0.97969	0.020973

(a) $\alpha = 0.99$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	1	60	0.012571	0.99709	0.032104
B	64	60	0.012433	0.99495	0.044051
B	8	3600	0.019926	0.9973	0.049417

(b) $\alpha = 0.999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
A	1	3600	0.15949	0.99945	0.61341
B	8	86400	0.035472	0.99868	0.61803
B	8	604800	0.034094	0.99865	0.63687

(c) $\alpha = 0.9999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
A	1	3600	0.18531	0.99996	0.59946
A	1	21600	0.38217	0.99994	1.5316
A	1	86400	0.59664	0.99994	2.6857

(d) $\alpha = 0.99999$

chosen such that the size of the domain \mathbb{X} (see Eq. (6)) is equal to 2000 points. It was experimentally observed that the values of PIs computed using larger domains did not change within three significant digits. Also, the average time needed to compute the PI and do one cycle of on-line training was around $50\mu s$; thus, the method meets the RT requirements.

For each combination of models, of a number of clusters, and time constants, and for each confidence level α , we compute the three metrics PINAW, PICP, and CWC. Tables I, II, III and IV show the 3 best configurations, i.e., those that achieve the smallest CWC, for each target confidence level, for the four prosumers, respectively.

Concerning the office building and the one hosting the EV

TABLE IV
PERFORMANCE EVALUATION FOR THE MV FEEDERS AT 20ms
RESOLUTION (3 BEST CONFIGURATIONS)

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	8	21600	0.0043777	0.99021	0.0043777
B	256	3600	0.0045151	0.99053	0.0045151
B	512	3600	0.0045814	0.99118	0.0045814

(a) $\alpha = 0.99$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	64	21600	0.0061943	0.999	0.0061943
B	256	21600	0.0062623	0.99908	0.0062623
B	512	21600	0.0062895	0.99911	0.0062895

(b) $\alpha = 0.999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	64	21600	0.0080307	0.9999	0.0080307
B	256	21600	0.0081893	0.99991	0.0081893
B	64	86400	0.0082057	0.99991	0.0082057

(c) $\alpha = 0.9999$

Model	L	T_ϕ (sec)	PINAW	PICP	CWC
B	1	86400	0.018185	0.99999	0.018185
B	8	604800	0.018747	0.99999	0.018747
B	1	604800	0.018985	0.99999	0.018985

(d) $\alpha = 0.99999$

charging station, model B outperforms model A for every confidence level. Also, for confidence levels up to 99.9%, the algorithm is able to achieve the target confidence. For higher confidence levels, however, the algorithm misses the target confidence by at most one order of magnitude. Given our choice of the parameter μ , this means that the computed CI is penalized by a maximum factor of 10 in the computation of CWC.

Regarding the optimal configuration, fewer clusters, together with larger forgetting time constants, are better choices. More specifically, the optimal number of clusters L does not need to be larger than 8, whereas a value of T_ϕ between one day and one week is the optimal choice. By taking a closer look at Tables I and II, we observe a trade-off between the values of L and T_ϕ in the performance of the algorithm. If we focus on Table Ia, for example, we notice that we could choose, without affecting the value of CWC by more than 1%, either a combination of 8 clusters and a forgetting time constant of one week, or 256 clusters and one day. In fact, this table showcases that the choice of L and T_ϕ has a minor effect on the computation of PIs for the office building. The performance of the algorithm is mainly influenced by the differentiation of the measurements performed by model B. For the charging station, however, more clusters perform up to 30% worse compared to the optimal case of only 8 clusters.

The results are quite different for the heat pump in Table III. First of all, the algorithm can only approach the target confidence levels within one order of magnitude. Unlike the other two buildings, the optimal results are achieved by a combination of few clusters (ideally only one) and a small forgetting time constant (less than an hour). This implies that the consumption of the heat pump changes more rapidly than that of the office building and the charging station, and that the measurements older than one hour do not influence the

TABLE V
BEST CONFIGURATION AT 20ms RESOLUTION - CLUSTERING BASED ON POWER (P) AND POWER+TIME OF DAY (P+TOD)

α	Model		L		T_ϕ (sec)		PINAW		PICP		CWC	
	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD
0.99	B	B	8	8	604800	86400	0.02148	0.02127	0.99129	0.99038	0.02148	0.02127
0.999	B	B	1	1	604800	604800	0.03054	0.03054	0.99909	0.99909	0.03054	0.03054
0.9999	B	B	1	1	604800	604800	0.04950	0.04950	0.99978	0.99978	0.11455	0.11455
0.99999	B	B	1	1	604800	604800	0.09760	0.09760	0.99999	0.99999	0.19992	0.19992

(a) OFFICE BUILDING

α	Model		L		T_ϕ (sec)		PINAW		PICP		CWC	
	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD
0.99	B	B	8	8	86400	86400	0.00707	0.00749	0.99068	0.99138	0.00707	0.00749
0.999	B	B	8	8	86400	21600	0.01930	0.01843	0.99868	0.99848	0.04008	0.03918
0.9999	B	B	8	8	86400	86400	0.03202	0.03233	0.99958	0.99957	0.09942	0.10076
0.99999	B	B	8	8	86400	86400	0.05183	0.05399	0.99992	0.99992	0.32993	0.29896

(b) CHARGING STATION

α	Model		L		T_ϕ (sec)		PINAW		PICP		CWC	
	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD
0.99	B	B	1	1	60	60	0.00940	0.00940	0.98832	0.98832	0.01918	0.01918
0.999	B	B	1	1	60	60	0.01257	0.01257	0.99709	0.99709	0.03210	0.03210
0.9999	A	A	1	1	3600	3600	0.15949	0.15949	0.99945	0.99945	0.61341	0.61341
0.99999	A	A	1	1	3600	3600	0.18531	0.18531	0.99996	0.99996	0.59946	0.59946

(c) HEAT PUMP

α	Model		L		T_ϕ (sec)		PINAW		PICP		CWC	
	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD	P	P+TOD
0.99	B	B	8	256	21600	21600	0.00438	0.00465	0.99021	0.99078	0.00438	0.00465
0.999	B	B	64	256	21600	21600	0.00619	0.00625	0.99900	0.99908	0.00619	0.00625
0.9999	B	B	64	64	21600	21600	0.00803	0.00804	0.99990	0.99990	0.00803	0.00804
0.99999	B	B	1	1	86400	86400	0.01818	0.01818	0.99999	0.99999	0.01818	0.01818

(d) FEEDERS

computation of PIs. We also observe that model A outperforms model B for large confidence levels. Even though the value of PINAW computed by model A is up to 5 times larger than the one computed by model B (not shown in the table for the sake of brevity), model B fails to achieve the target confidence level by up to two orders of magnitude, which results in an exponential increase in the metric CWC.

Finally, the algorithm performs better when applied to the large heterogeneous aggregated prosumption of the EPFL campus (see Table IV) than the three individual prosumers. The algorithm can achieve all four target confidence levels, and the relative widths of the PIs are around 5 times smaller than the respective ones for the office building, which was the best-performing case of the first three prosumers. This result showcases that it is easier to forecast electrical prosumption at higher aggregation levels and that the proposed PI computation method can be reliably applied to different prosumers' aggregated levels.

Overall, we conclude that the optimal choice of model, L and T_ϕ depends on the characteristics of the prosumer. Nodes with low volatility, such as an office building, benefit from a long memory of up to one week, whereas nodes characterized by rapid power changes require the use of short memory. The clustering of power measurements seems to benefit mainly nodes with clearly distinct power-levels, such as those at a charging station. But in any case, the number of clusters does not need to be more than 8. Finally, the differentiation of power (i.e., model B) improves the computation of PIs, provided that

the power differentials have low volatility, as is the case for the office building and the charging station.

D. Clustering Using Power Levels and Time of Day

In this section, we consider whether adding the time of day (TOD) as a feature of the algorithm, in addition to the power level (P), would improve the performance of the algorithm. The idea is that the clustering based on TOD will capture the patterns that the prosumption exhibits during the day and that this might affect the computation of the PIs. We re-run the experiments of the previous section, with the addition of the TOD. The best configuration is again the one that achieves the smallest value of CWC for each target confidence level. The comparison between the two cases, specifically (i) using only power level (P) and (ii) using power level and the time of day (P+TOD), is shown on Table V.

We observe that the best configuration for the two clustering strategies is the same in most cases. For the office building and the MV feeders, the introduction of the TOD does not affect the value of CWC by more than 1%, which confirms once again that the clustering method does not affect the performance of the algorithm. For the charging station, the value of CWC differs by 5 – 10% between the two clustering methods. However, the introduction of the new feature could either improve or worsen the performance, depending on the confidence level. Therefore, we cannot conclude whether it consistently performs better. As far as the heat pump is concerned, the usage of the TOD in the clustering has no

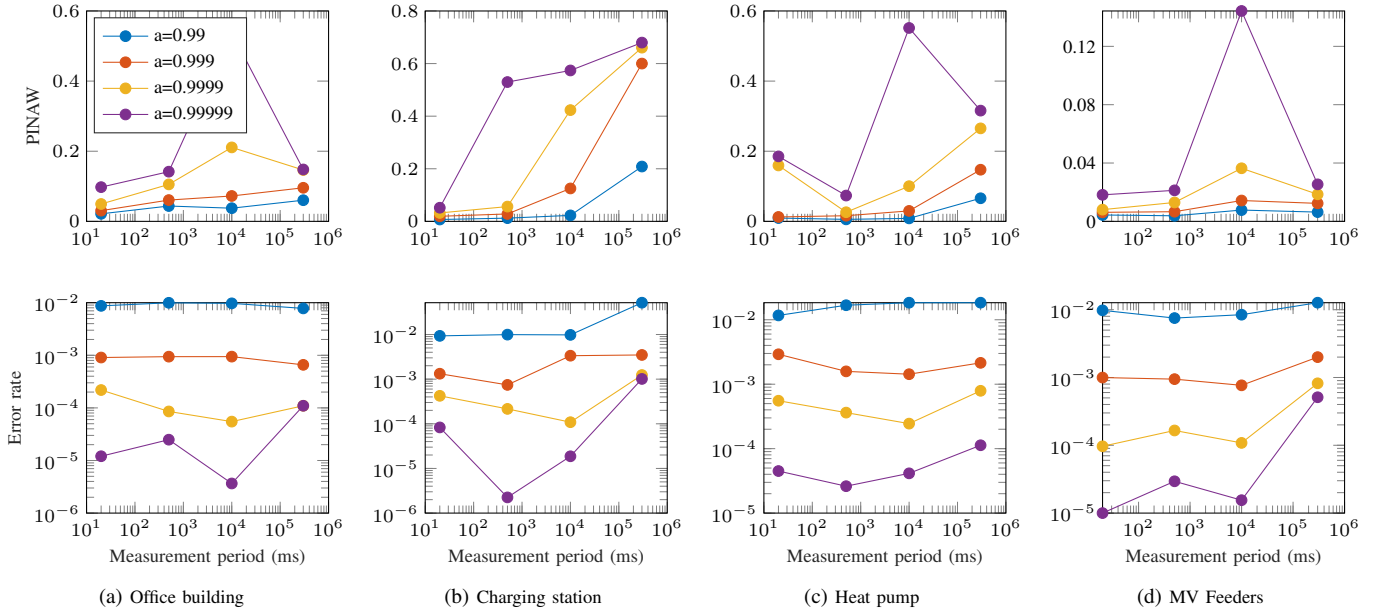


Fig. 5: Performance evaluation as a function of the measurement period

effect on the performance, because the algorithm performs better when all measurements are put in one cluster.

The results indicate that this feature does not play a critical role in the computation of PIs. Perhaps other features, apart from the power level and the time of day, could influence the performance of the algorithm. The effect of additional features could be studied in future research.

E. Effect of the Measurement Period

A crucial objective of this work is to find out how the proposed algorithm scales as the measurement period increases. To test the algorithm on different measurement periods, we integrate the available data-set. In particular, given the measurements $P_j^{20}, j = 1..N$ at $20ms$ resolution, the measurements at resolution T (ms) are recomputed as:

$$P_i^T = \frac{20}{T} \sum_{j=(i-1)\frac{T}{20}+1}^{i\frac{T}{20}} P_j^{20}, i = 1.. \frac{20}{T}N \quad (24)$$

where T is assumed to be a multiple of $20ms$. The forgetting factor ϕ is scaled according to T , as in Eq. (13). For this and the following section, we do clustering based only on the power level.

We perform simulations again with varying cluster numbers L and forgetting time constants T_ϕ for different confidence levels. In Fig. 5, we plot the values of PINAW and the error rate as a function of the measurement period from $20ms$ up to $300s$ for the four prosumers considered. Each point on the graphs corresponds to a different configuration (the one that achieves the smallest CWC), as shown in Table VI.

The graph for the office building showcases that the algorithm performs well even for larger measurement periods, provided that the target confidence level is less than 99.99%. Indeed, for low confidence levels, the error rate is kept below

the target, and the average width of the PI is less than 20% of the nominal value. For larger confidence levels, however, the algorithm fails because either the target level is not achieved or the width of the PI is so high that it becomes useless for grid control. It is worth noting that at a resolution of five minutes, the result of the algorithm is identical for all target confidence levels above 99.99%. This implies that we might need a larger history to meet high confidence levels when the measurements are sparse. The same results can be deduced for the MV feeders but with much smaller relative PI widths than the office building.

Similar observations can be made for the charging station. The main difference is that the maximum measurement period with an acceptable performance depends on the target confidence level. With target confidence of 99%, the algorithm computes a low PI width for measurement periods up to $10s$; whereas, with 99.9 – 99.999%, it has acceptable performance but only when the resolution is sub-second. For higher confidence, the algorithm cannot predict accurate PIs for periods larger than $20ms$.

Looking at the results for the heat pump on Fig. 5, we notice a break in the trend of increasing PINAW. To understand why this happens, we look at Table VI. The best configuration for confidence levels 99.99 – 99.999% changes from using model A at $20ms$ resolution to model B at higher resolutions. Model A has been shown in Section IV-C to generate, in general, larger PIs than model B. This increase in the width results in a lower error rate, which in turn might reduce the value of CWC, which is the sole metric used to compare configurations to one another. The value of CWC ultimately depends on the choice of the parameter μ , which might affect the trends in the graphs of Fig. 5.

TABLE VI
BEST CONFIGURATION AS A FUNCTION OF THE MEASUREMENT PERIOD

$T(ms)$	α	Office building			Charging station			Heat pump			MV Feeders		
		Model	L	$T_\phi(sec)$	Model	L	$T_\phi(sec)$	Model	L	$T_\phi(sec)$	Model	L	$T_\phi(sec)$
20	0.99	B	8	604800	B	8	86400	B	1	60	B	8	21600
	0.999	B	1	604800	B	8	86400	B	1	60	B	64	21600
	0.9999	B	1	604800	B	8	86400	A	1	3600	B	64	21600
	0.99999	B	1	604800	B	8	86400	A	1	3600	B	1	86400
500	0.99	B	1	21600	B	1	86400	B	1	60	B	1	604800
	0.999	B	1	604800	B	1	604800	B	1	3600	B	1	86400
	0.9999	B	1	86400	B	8	86400	B	1	3600	B	1	3600
	0.99999	B	1	21600	A	1	604800	B	1	604800	B	1	604800
10000	0.99	B	64	86400	B	1	604800	B	8	604800	B	1	604800
	0.999	B	8	86400	B	8	604800	B	1	604800	B	1	604800
	0.9999	B	1	604800	B	1	604800	B	8	604800	B	1	604800
	0.99999	B	1	604800	B	1	604800	B	1	604800	B	1	604800
300000	0.99	B	1	604800	B	1	3600	B	1	21600	B	1	86400
	0.999	B	1	604800	B	1	604800	B	1	86400	B	1	604800
	0.9999	B	1	604800	B	1	604800	B	1	604800	B	1	604800
	0.99999	B	8	604800	B	1	604800	B	1	604800	B	1	604800

F. Confidence Level Uncertainty

The metrics of Section IV-B evaluate, over the full test period, the average performance of the algorithm. However, it would be interesting to see how the algorithm performs over time. For this purpose, we split the one-month evaluation period (15 days for the MV feeders) into six-hour windows and compute the error rate achieved by the algorithm in each window. Hence, for each confidence level, we compute a histogram of 120 estimations (60 for the MV feeders) of the error rate. We then depict the statistical measures of the histograms using the box plots of Fig. 6. The blue box indicates the 25th and 75th percentiles, the red line is the median, the red crosses are outliers, and the small square indicates the target error rate. We computed the box plots only for the best configuration for each building and for confidence level at 20ms resolution, which are shown in Table VI.

If the algorithm is consistent in predicting accurate PIs, then the median of the box plot should be close to the target error-rate. We see that this is indeed the case for the office building. For the largest confidence level, our algorithm performs even better than expected, albeit with many outliers. For the charging station, the performance is sometimes worse than expected, but the target error-rate is nevertheless contained within the 25th and the 75th percentile of the box plot. This is not always the case for the heat pump, hence the algorithm cannot consistently estimate accurate PIs if there is high volatility in the measurements. Finally, the case of the MV feeders is once again the best-performing out of the four, as the median of all box plots is very close to the target level.

The results showcase that our algorithm computes accurate PIs, irrespective of the building, provided that the target confidence level is less than 99.9%. For higher confidence levels, the performance is still acceptable, but the consistency of the predictions depends on the dynamics of the node.

V. CONCLUSION

Motivated by the requirements of ADNs real-time control, we have presented a non-parametric method for computing ultra-short-term PIs (prediction intervals) of the power

prosumption in generic ADNs (e.g. buildings). The method consists in grouping historical measurements into clusters, according to the value of selected influential variables. The clusters are considered statistically representative pools of future power-prosumption realizations and are used to extract PIs at arbitrary confidence levels by calculating the quantiles from the respective PDF.

The proposed method has been applied to the original and once-differentiated power-prosumption time series, and different influential variables have been considered. The performance of the method was tested for different types of prosumers by using experimental measurements from an MV distribution network. The performance analysis enabled us to make an a-posteriori selection of the parameters of the algorithm. The algorithm was shown to compute relatively narrow PIs for the studied prosumers, for time resolutions from 20ms up to a few minutes in some cases, provided that the target confidence level is below 99.9%.

A final statement concerns the computational complexity, which becomes a relevant concern especially when considering densely sampled time series and the high reporting rate for the predictions. We have shown that the proposed algorithm performs the PI computation and on-line training in constant time hence is scalable.

REFERENCES

- [1] "Pilo f, jupe s, silvestro f et al (2014) planning and optimization methods for active distribution systems. wg c6.19: Tb 591, cigre, paris, france."
- [2] R. Rudnik, C. Wang, L. Reyes-Chamorro, J. Achara, J.-Y. L. Boudec, and M. Paolone, "Real-time control of an electric vehicle charging station while tracking an aggregated power setpoint," *IEEE Transactions on Industry Applications*, vol. 56, no. 5, pp. 5750–5761, 2020.
- [3] E. Scolari, D. Torregrossa, J.-Y. Le Boudec, and M. Paolone, "Ultra-short-term prediction intervals of photovoltaic ac active power," in *2016 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, 2016, pp. 1–8.
- [4] W. U. Rehman, A. Moeini, O. Oboreh-Snapps, R. Bo, and J. Kimball, "Deadband voltage control and power buffering for extreme fast charging station," in *2021 IEEE Madrid PowerTech*, 2021, pp. 1–6.
- [5] A. Bernstein, L. E. Reyes Chamorro, J.-Y. Le Boudec, and M. Paolone, "A composable method for real-time control of active distribution networks with explicit power set points. part I: Framework," *Electric Power Systems Research*, vol. 125, no. August, pp. 254–264, 2015.

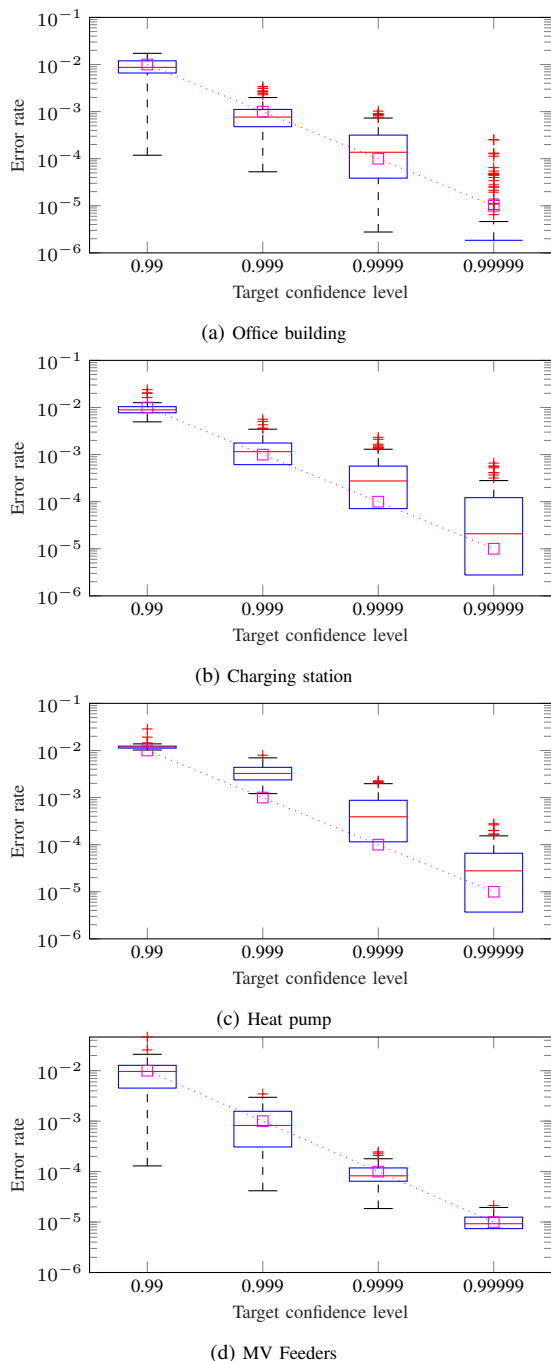


Fig. 6: Box plot of model accuracy versus target confidence level at 20ms resolution

[6] G. Valverde and T. Van Cutsem, "Model predictive control of voltages in active distribution networks," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 2152–2161, 2013.

[7] N. Hatzigiorgiou, *Microgrids: Architectures and Control*, ser. Wiley - IEEE. Wiley, 2014.

[8] L. Reyes-Chamorro, W. Saab, R. Rudnik, A. M. Kettner, M. Paolone, and J. Le Boudec, "Slack selection for unintentional islanding: Practical validation in a benchmark microgrid," in *2018 Power Systems Computation Conference (PSCC)*, 2018, pp. 1–7.

[9] R. Posree and S. Sirisumrannukul, "Voltage sag assessment in distribution system with neutral grounding resistance by methods of fault position and monte carlo simulation," in *2021 5th International Conference*

on Power and Energy Engineering (ICPEE), 2021, pp. 26–31.

[10] "Voltage sag prediction for network planning," *Electric Power Systems Research*, vol. 140, pp. 976–983, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378779616300803>

[11] M. Bollen and I. Gu, *Signal Processing of Power Quality Disturbances*, ser. IEEE Press Series on Power and Energy Systems. Wiley, 2006. [Online]. Available: <https://books.google.ch/books?id=tOhSAAAAMAAJ>

[12] M. Cho, J. Hwang, and C. Chen, "Customer short term load forecasting by using arima transfer function model," in *Energy Management and Power Delivery, 1995. Proceedings of EMPD'95., 1995 International Conference on*, vol. 1. IEEE, 1995, pp. 317–322.

[13] R. Mamlook, O. Badran, and E. Abdulhadi, "A fuzzy inference model for short-term load forecasting," *Energy Policy*, vol. 37, no. 4, pp. 1239 – 1248, 2009.

[14] F. Sossan, V. Lakshmanan, G. Costanzo, M. Marinelli, P. Douglas, and H. Bindner, "Grey-box modelling of a household refrigeration unit for energy consumption prediction and optimization using time series data," 2013, currently unpublished.

[15] F. Sossan, H. Bindner, H. Madsen, L. Reyes, D. Torregrossa, and M. Paolone, "A MPC replacement strategy for electric space heating including cogeneration of a fuel cell-electrolyzer system," *International Journal of Electrical Power & Energy Systems*, 2013.

[16] A. Azadeh, S. Ghaderi, and S. Sohrabkhani, "Annual electricity consumption forecasting by neural network in high energy consuming industrial sectors," *Energy Conversion and Management*, vol. 49, no. 8, pp. 2272–2278, 2008.

[17] H. K. Alfares and M. Nazeeruddin, "Electric load forecasting: literature survey and classification of methods," *International Journal of Systems Science*, vol. 33, no. 1, pp. 23–34, 2002.

[18] H. Hahn, S. Meyer-Nieberg, and S. Pickl, "Electric load forecasting methods: Tools for decision making," *European Journal of Operational Research*, vol. 199, no. 3, pp. 902–907, 2009.

[19] M. Tan, S. Yuan, S. Li, Y. Su, H. Li, and F. He, "Ultra-short-term industrial power demand forecasting using lstm based hybrid ensemble learning," *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 2937–2948, 2020.

[20] Z. Wang and A. Perera, "Robust optimization of power grid with distributed generation and improved reliability," *Energy Procedia*, vol. 159, pp. 400–405, 2019, renewable Energy Integration with Mini/Microgrid. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S187661021831364X>

[21] B. Hayes, J. Gruber, and M. Prodanovic, "Short-term load forecasting at the local level using smart meter data," in *PowerTech, 2015 IEEE Eindhoven*, June 2015, pp. 1–6.

[22] H. A. Nielsen, H. Madsen, and T. S. Nielsen, "Using quantile regression to extend an existing wind power forecasting system with probabilistic forecasts," *Wind Energy*, vol. 9, no. 1-2, pp. 95–108, 2006.

[23] P. Pinson, "Estimation of the uncertainty in wind power forecasting," 2006.

[24] M. Rana, I. Koprinska, A. Khosravi, and V. Agelidis, "Prediction intervals for electricity load forecasting using neural networks," in *Neural Networks (IJCNN), The 2013 International Joint Conference on*, Aug 2013, pp. 1–8.

[25] T. Xiong, Y. Bao, and Z. Hu, "Interval forecasting of electricity demand: A novel bivariate emd-based support vector regression modeling framework," *International Journal of Electrical Power & Energy Systems*, vol. 63, no. 0, pp. 353 – 362, 2014.

[26] H. Chernoff, S.-H. Lo, and T. Zheng, "Discovering influential variables: a method of partitions," *The Annals of Applied Statistics*, pp. 1335–1369, 2009.

[27] M. Pignati, M. Popovic, S. Barreto Andrade, R. Cherkaoui, D. Flores, J.-Y. Le Boudec, M. M. Maaz, M. Paolone, P. Romano, S. Sarri, T. T. Tesfay, D.-C. Tomozei, and L. Zanni, "Real-Time State Estimation of the EPFL-Campus Medium-Voltage Grid by Using PMUs," in *The Sixth Conference on Innovative Smart Grid Technologies (ISGT2015)*, 2014.

[28] A. Khosravi, S. Nahavandi, and D. Creighton, "Prediction intervals for short-term wind farm power generation forecasts," *IEEE Transactions on Sustainable Energy*, vol. 4, no. 3, p. 602 – 610, 2013.