

Deep learning-based multi-output quantile forecasting of PV generation

Jonathan Dumas, Colin Cointe, Xavier Fettweis, Bertrand Cornélusse

Departments of computer science, electrical engineering, and geography

University of Liège, Belgium

{jdumas, xavier.fettweis, bertrand.cornelusse}@uliege.be, colin.cointe@mines-paristech.fr

Abstract—This paper develops probabilistic PV forecasters by taking advantage of recent breakthroughs in deep learning. A tailored forecasting tool, named *encoder-decoder*, is implemented to compute intraday multi-output PV quantiles forecasts to efficiently capture the time correlation. The models are trained using quantile regression, a non-parametric approach that assumes no prior knowledge of the probabilistic forecasting distribution. The case study is composed of PV production monitored on-site at the University of Liège (ULiège), Belgium. The weather forecasts from the regional climate model provided by the Laboratory of Climatology are used as inputs of the deep learning models. The forecast quality is quantitatively assessed by the continuous ranked probability and interval scores. The results indicate this architecture improves the forecast quality and is computationally efficient to be incorporated in an intraday decision-making tool for robust optimization.

Index Terms—Quantile forecasting, probabilistic PV forecasting, LSTM, deep learning, encoder-decoder

I. INTRODUCTION

The Intergovernmental Panel on Climate Change special report¹ on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways² presents several scenarios of decarbonisation of the electricity sector with a renewable share that should reach in 2030 the interquartile range [47, 65] ([69, 86] in 2050). Therefore, the development of renewable generations, typically from wind and photovoltaic (PV) sources has been facilitated by policy makers. However, the intermittent and uncertain nature of these sources is challenging the traditional operation of electricity networks. Market players require reliable decision-making tools to deal with uncertainty based on forecasts of renewable generation.

In contrast to *point* predictions, *probabilistic* forecasts aim at providing decision-makers with full information about potential future outcomes [1]. The various types of probabilistic forecasts range from quantile to density forecasts, and through prediction intervals. This paper focuses on *quantile* forecasts that provide probabilistic information about future renewable power generation, in the form of a threshold level associated with a probability [1]. Quantile regression [2] is one of the most famous non-parametric approaches. It does not assume

the shape of the predictive distributions and is implemented with neural networks, linear regression, gradient boosting, or any other regression techniques.

The following papers have gained our attention in PV probabilistic forecasting. At the Global Energy Forecasting Competition 2014 [3] solar forecasts were to be expressed in the form of 99 quantiles with various nominal proportions between zero and one. A systematic framework for generating PV probabilistic forecasts is developed by [4]. A non-parametric density forecasting method based on Extreme Learning Machine is adopted to avoid restrictive assumptions on the shape of the forecast densities. A combination of bidirectional Recurrent Neural Networks (RNNs) with Long Short-Term Memory (LSTM) resulting in Bidirectional LSTM (BLSTM) is proposed by [5]. It has the benefits of both long-range memory and bidirectional processing. The BLSTM is trained by minimizing the quantile loss to compute quantile forecasts of aggregated load, wind and PV generation, and electricity prices on a day-ahead basis. Finally, an innovative architecture, referred to as encoder-decoder (ED), is developed by [6] to generate reliable predictions of the future system imbalance used for robust optimization.

In this study, the forecast quality of the models is evaluated. It corresponds to the ability of the forecasts to genuinely inform of future events by mimicking the characteristics of the processes involved. We follow the framework proposed by [7], [8] for evaluating the quality of solar irradiance probabilistic forecasts based on visual diagnostic tools and a set of scoring rules. Overall they indicate that two main attributes, reliability and resolution, characterize the quality of quantile forecasts.

This work exploits recent breakthroughs in the field of data science by using advanced deep learning structures, such as the encoder-decoder architecture [6], and quality metrics [8], [9] to develop a tailored deep learning-based multi-output quantile forecaster. The goal is to capture the time correlation between time periods and to use this forecaster as input of a robust optimization model. For instance, to address the energy management system of a grid-connected renewable generation plant coupled with a battery energy storage device [10].

Overall, the contributions can be summarized as follows. First, a deep learning-based multi-output quantile architecture is used to compute prediction intervals of PV generation on a day-ahead and intraday basis. Specifically, the goal is to implement an improved probabilistic intraday forecaster,

¹<https://www.ipcc.ch/sr15/>

the encoder-decoder, to benefit from the last PV generation observations. This architecture is compared to a feed-forward neural network that is used as the benchmark model. Second, the weather forecasts of the MAR climate regional model [11] are used. It allows to directly take into account the impact of the weather forecast updates generated every six hours. Finally, a proper assessment of the quantile forecasts is conducted by using a k -fold cross-validation methodology and probabilistic metrics. It allows computing average scores over several testing sets and mitigating the dependency of the results to specific days of the dataset.

The remainder of this paper is organized as follows. Section II provides the non-parametric quantile forecasting framework considered. Section III presents the forecasting techniques used to compute the quantile forecasts. Section IV details the different metrics used to evaluate the quality of PV quantile forecasts. Section V describes the case study and presents the results. Finally, Section VI summarizes the main findings and highlights ideas for further work.

II. NON-PARAMETRIC QUANTILE FORECASTING

Let y_t be the PV power generation measured at time t , which corresponds to a realization of the random variable Y_t . Then let f_t and F_t be the probability density function (PDF) and related cumulative distribution function (CDF) of Y_t , respectively. Following the definition of probabilistic forecasting of [1], a probabilistic forecast issued at time t for time $t+k$ consists of a prediction of the PDF (or equivalently, the CDF) of Y_{t+k} , or of some summary features. Various types of probabilistic forecasts have been developed, from quantile to density forecasts, and through prediction intervals. This study focuses on quantile regression [2], which is the most widely used type of probabilistic forecasting method. It is a non-parametric approach with no restrictive assumption on the shape and features of the predictive distributions. Indeed, empirical investigations [4] showed that PV power forecast errors do not follow common, *e.g.* Gaussian, Beta, distributions. Following the definition of [1], a (model-based) forecast $\hat{y}_{t+k|t}$ of PV power generation is an estimate of some of the characteristics of the stochastic process Y_{t+k} given a model g , its estimated parameters $\hat{\Theta}_t$ and the information set Ω_t gathering all data and knowledge about the processes of interest up to time t .

A. Point forecasting

A point prediction $\hat{y}_{t+k|t}$ is a single-valued issued at time t for $t+k$. It corresponds to the conditional expectation of Y_{t+k} given g , $\hat{\Theta}_t$, and the information set Ω_t

$$\hat{y}_{t+k|t} = \mathbb{E}[Y_{t+k|t}|g, \Omega_t, \hat{\Theta}_t]. \quad (1)$$

A multi-output point forecast computed at t for $t+k_1$ to $t+k_T$ is the vector $[\hat{y}_{t+k_1|t}, \dots, \hat{y}_{t+k_T|t}]^T$ of size T .

B. Quantile forecasting

A quantile forecast $\hat{y}_{t+k|t}^{(q)}$ with nominal level q is an estimate, issued at time t for time step $t+k$ of the quantile

$y_{t+k|t}^{(q)}$ for the random variable $Y_{t+k|t}$ given a model g , its estimated parameters $\hat{\Theta}_t$ and the information set Ω_t

$$P[Y_{t+k|t} \leq \hat{y}_{t+k|t}^{(q)} | g, \Omega_t, \hat{\Theta}_t] = q, \quad (2)$$

or equivalently $\hat{y}_{t+k|t}^{(q)} = \hat{F}_{t+k|t}^{-1}(q)$, with \hat{F} the estimated cumulative distribution function of the continuous random variable Y . Finally, a multi-output quantile forecast computed at t for $t+k_1$ to $t+k_T$ is the matrix \hat{Z}_t of dimensions $T \times Q$

$$\hat{Z}_t = \begin{bmatrix} \hat{y}_{t+k_1|t}^{(q=1)} & \dots & \hat{y}_{t+k_1|t}^{(q=Q)} \\ \vdots & \vdots & \vdots \\ \hat{y}_{t+k_T|t}^{(q=1)} & \dots & \hat{y}_{t+k_T|t}^{(q=Q)} \end{bmatrix}, \quad (3)$$

with Q quantiles per time period.

C. Quantile loss function

Quantile regression consists of estimating quantiles by applying asymmetric weights to the mean absolute error. Following [2], the quantile loss function is

$$\rho_q(x, y) = \begin{cases} q \times (x - y) & x > y \\ (1 - q) \times (y - x) & x \leq y \end{cases}. \quad (4)$$

For a given time period t , ρ_q is summed over the forecasting time periods $k_1 \leq k \leq k_T$ and quantiles $q \in \mathcal{Q}$ to compute multi-output quantile forecasts at t for $t+k_1$ to $t+k_T$

$$L_t = \sum_{k_1 \leq k \leq k_T} \sum_{q \in \mathcal{Q}} \rho_q(y_{t+k}, \hat{y}_{t+k|t}^{(q)}). \quad (5)$$

Finally, the model g is trained by minimizing $\frac{1}{|\Omega_t|} \sum_{t' \in \Omega_t} L_{t'}$. Note in the case of perfect prediction, the quantile loss cannot be differentiated, and a smooth approximation of (4) using the Huber norm [6] is built.

D. Prediction intervals (PIs)

Prediction intervals (PIs) define the range of values within which the observation is expected to be with a certain probability, *i.e.*, its nominal coverage rate [9]. Formally, a prediction interval $\hat{I}_{t+k|t}^{(\alpha)}$ issued at t for $t+k$, defines a range of potential values for Y_{t+k} , for a certain level of probability $(1 - \alpha)$, $\alpha \in [0, 1]$. Its nominal coverage rate is

$$P[Y_{t+k} \in \hat{I}_{t+k|t}^{(\alpha)} | g, \Omega_t, \hat{\Theta}_t] = 1 - \alpha. \quad (6)$$

A central PI consists of centering the PI on the median where there is the same probability of risk below and above the median. A central PI with a coverage rate of $(1 - \alpha)$ is estimated by using the quantiles $q = (\alpha/2)$ and $q = (1 - \alpha/2)$. Its nominal coverage rate is

$$\hat{I}_{t+k|t}^{(\alpha)} = [\hat{y}_{t+k|t}^{(q=\alpha/2)}, \hat{y}_{t+k|t}^{(q=1-\alpha/2)}]. \quad (7)$$

For instance, central PIs with a nominal coverage rate of 90 %, *i.e.*, $(1 - \alpha) = 0.9$, are defined by quantile forecasts with nominal levels of 5 and 95 %.

III. FORECASTING TECHNIQUES

A. Gradient boosting regression (GBR)

Gradient boosting builds an additive model in a forward stage-wise fashion [12]. It allows for the optimization of arbitrary differentiable loss functions. In each stage, a regression tree is fit on the negative gradient of the given loss function. The Scikit-learn [13] Python library is used to implement a gradient boosting regressor (GBR) with the quantile loss function. The learning rate is set to 10^{-2} , the max depth to 5, and the number of estimators to 500. There is a GBR model per quantile as the library does not support multi-output quantile regression.

B. Multilayer perceptron (MLP)

A description of the most widely used "vanilla" neural network, the *Multilayer perceptron* (MLP), is provided by [12]. A MLP with a single hidden layer is considered for the day-ahead forecasts and as the benchmark for the intraday forecasts. MLPs with two and three hidden layers did not provide any significant improvement. The activation function is the Rectified Linear Unit (ReLU). The number of neurons of the hidden layer is $n_{\text{input}} + (n_{\text{output}} - n_{\text{input}})/2$, with n_{input} and n_{output} the number of neurons of the input and output layers, respectively. The learning rate is set to 10^{-2} and the number of epoch to 500 with a batch size of 8. It is implemented using the PyTorch Python library [14].

C. Encoder-decoder (ED)

Several technical information about recent advances in neural networks is provided by [5], [6]. In particular, recurrent neural networks, have shown a high potential in processing and predicting complex time series with multi-scale dynamics. However, RNNs are known to struggle in accessing time dependencies more than a few time steps long due to the vanishing gradient problem. Indeed, back-propagated errors during the training stage either fades or blows up over time. Long Short-Term Memory and Gated Recurrent Units networks tackle this problem by using internal memory cells [6]. A neural network composed of a LSTM and feed-forward layers, referred to as LSTM in the rest of the paper, is implemented for the day-ahead and intraday forecasts. The number of LSTM units is $n_{\text{input}} + (n_{\text{output}} - n_{\text{input}})/3$, and the number of neurons of the feed-forward layer $n_{\text{input}} + 2 \times (n_{\text{output}} - n_{\text{input}})/3$.

An innovative architecture, referred to as encoder-decoder [6], is composed of two different networks and has recently shown promising results for translation tasks and speech recognition applications and imbalance price forecasting. The encoder-decoder processes features from the past, such as past PV observations, to extract the relevant historical information that is contained into a reduced vector of fixed dimensions, based on the last hidden state. Then, the decoder processes this representation along with the known future information such as weather forecasts. A version of the encoder-decoder architecture (ED-1) is implemented with a LSTM as the

encoder and a MLP as the decoder. In a second version (ED-2) the decoder is a LSTM followed by an additional feed-forward layer. Both versions of the encoder-decoder are used as intraday forecasters. In ED-1, the encoder has $2 \times n_{\text{input}}$ units with n_{input} the number of neurons of the encoder input layer, features from the past. Then, the encoder output is merged with the weather forecasts becoming the decoder input layer that has $n_{\text{output}}/2$ neurons. In ED-2, the decoder has the same number of cells as the encoder, and the feed-forward layer is composed of $n_{\text{output}}/2$ neurons. The LSTM, ED-1, and ED-2 models are implemented using the TensorFlow Python library [15]. The activation functions are the ReLU, the learning rate is set to 10^{-3} , the number of epoch to 500 with a batch size of 64 for the three models.

A sensitivity analysis has been conducted to select the hyperparameters: number of hidden layers, neurons, epochs, and learning rate. Overall, increasing the number of hidden layers and neurons increases the model complexity. It can enhance the accuracy, but only up to a limited number of layers and neurons due to overfitting issues. In addition, the hyperparameter solution is closely related to the size of the historical database [5]. A deep learning model with a larger amount of hidden layers and neurons requires a large amount of data to accurately estimate the parameters. In the case study considered, there are only 157 days of data with a 15 minutes resolution. Thus, we decided to restrict the number of layers and neurons to select a smaller model that performs better with the available information.

IV. PROBABILISTIC FORECASTING QUALITY ASSESSMENT

For predictions in any form, one must differentiate between their quality and their value [1]. Forecast quality corresponds to the ability of the forecasts to genuinely inform of future events by mimicking the characteristics of the processes involved. Forecast value relates, instead, to the benefits from using forecasts in a decision-making process such as participation in the electricity market. This section proposes quality metrics based on the framework proposed by [8]. The value assessment is not in the scope of this paper as it would require to consider a decision-making process.

A. Continuous rank probability score (CRPS)

A score is said to be proper if it ensures that the perfect forecasts should be given the best score value [16]. It is the case of the Continuous Rank Probability Score (CRPS) that penalizes the lack of resolution of the predictive distributions as well as biased forecasts. For deterministic forecasts, the CRPS turns out to be the Mean Absolute Error (MAE). The energy form [16] of the CRPS (NRG) is selected in this study

$$\text{CRPS}_{t,k}^{\text{NRG}}(\hat{F}_{t+k|t}, y_{t+k}) = \mathbb{E}_{\hat{F}_{t+k|t}} |X - y_{t+k}| - \frac{1}{2} \mathbb{E}_{\hat{F}_{t+k|t}} |X - X'|, \quad (8)$$

where X and X' are two independent copies of a random variable with distribution function $\hat{F}_{t+k|t}$ and finite first moment. A CRPS estimator (eNRG) of the energy form is provided

by [17] when the CDF is only known at $t+k$ through a Q -ensemble of quantile forecasts $\mathcal{Q}_{t,k} = \{\hat{y}_{t+k|t}^{(q)}\}_{q \in \mathcal{Q}}$

$$\begin{aligned} \text{CRPS}_{t,k}^{\text{eNRG}}(\mathcal{Q}_{t,k}, y_{t+k}) &= \frac{1}{Q} \sum_{q \in \mathcal{Q}} |\hat{y}_{t+k|t}^{(q)} - y_{t+k}| \\ &\quad - \frac{1}{2Q^2} \sum_{q \in \mathcal{Q}} \sum_{q' \in \mathcal{Q}} |\hat{y}_{t+k|t}^{(q)} - \hat{y}_{t+k|t}^{(q')}|. \end{aligned} \quad (9)$$

Finally, $\text{CRPS}^{\text{eNRG}}(k)$ is the mean over the evaluation set for a given forecasting time period k , and $\text{CRPS}^{\text{eNRG}}$ is the average over all k with $k_1 \leq k \leq k_T$.

B. Interval score (IS)

The Interval Score (IS) is a proper score proposed by [16] to specifically assess the quality of central $(1-\alpha)$ prediction interval forecasts. The IS rewards narrow prediction intervals but penalizes, with the penalty term that depends on α , the forecasts for which the observation is outside the interval. The averaged IS over an evaluation set of length M and $\forall k, k_1 \leq k \leq k_T$ is

$$\begin{aligned} \text{IS}_\alpha &= \frac{1}{M} \sum_{t=1}^M \frac{1}{T} \sum_{k=k_1}^{k_T} (\hat{y}_{t+k|t}^{(1-\alpha/2)} - \hat{y}_{t+k|t}^{(\alpha/2)}) \\ &\quad + \frac{2}{\alpha} (\hat{y}_{t+k|t}^{(\alpha/2)} - y_{t+k}) \mathbf{1}(y_{t+k} \leq \hat{y}_{t+k|t}^{(\alpha/2)}) \\ &\quad + \frac{2}{\alpha} (y_{t+k} - \hat{y}_{t+k|t}^{(1-\alpha/2)}) \mathbf{1}(y_{t+k} \geq \hat{y}_{t+k|t}^{(1-\alpha/2)}). \end{aligned} \quad (10)$$

V. THE ULIÈGE CASE STUDY

A. Case study description

The ULiège case study is composed of a PV generation plant with an installed capacity of 466.4 kW. The PV generation has been monitored on a minute basis from *April 4, 2020* to *September 14, 2020*, 157 days in total, and the data is resampled to 15 minutes. The set of quantiles is $\mathcal{Q} = \{0.1, 0.2, \dots, 0.9\}$ for both the day-ahead and intraday forecasts. Numerical experiments are performed on an Intel Core i7-8700 3.20 GHz based computer with 12 physical CPU cores and 32 GB of RAM running on Ubuntu 18.04 LTS.

B. Numerical settings

The MAR regional climate model [11] provided by the Laboratory of Climatology of the Liège University is forced by GFS (Global Forecast System) to compute weather forecasts on a six hours basis, four-time gates per day at 00:00, 06:00, 12:00, and 18:00 with a 10 day horizon and a 15 minutes resolution. The solar irradiance and air temperature at 2 meters are normalized by a standard scaler and used as inputs to the forecasting models.

A k -fold cross-validation strategy is used to compute average scores over several testing sets to mitigate the dependency of the results to specific days of the dataset. The dataset is divided into k parts of equal length, and there are k possible testing sets $1 \leq i \leq k$. For a given testing set i , the models are trained over the $k-1$ parts of the dataset. Eleven pairs

of fixed lengths of 142 and 15 days are built. One pair is used to conduct the hyperparameters sensitivity analysis, and the ten others for testing where the scores are averaged. The Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) are introduced to evaluate the point forecasts. The MAE, RMSE, CRPS, and IS are normalized by the PV total installed capacity with NMAE and NRMSE the normalized MAE and RMSE.

The day-ahead models, MLP, LSTM, and GBR compute forecasts at 12:00 for the next day. Four intraday time gates are considered at 00:00, 06:00, 12:00, and 18:00. The intraday forecasts of time gate 00:00 are computed by the day-ahead models using only the weather forecasts. Then, the next three intraday forecasts are computed by intraday models where the MLP, ED-1, and ED-2, models use the weather forecasts and the last three hours of PV generation.

The day-ahead and the first intraday predictions are delivered for the 96 quarters of the next day from 00:00 to 23:45 indexed by time steps $0 \leq k \leq 95$. The prediction horizons span from 12 to 36 hours, for the day-ahead gate 12:00, and 0 to 24 hours, for the intraday gate 00:00. The prediction horizon is cropped to $11 \leq k \leq 80$ because the PV generation is always 0 for time steps $0 \leq k \leq 10$ and $81 \leq k \leq 95$ on the ULiège case study. The next three intraday predictions are performed for the 72, 48, and 24 next quarters of the day corresponding to the gates 06:00, 12:00, and 18:00. Therefore, the prediction horizons span from 0 to 18 hours, 0 to 12 hours, and 0 to 6 hours. The intraday forecasting time periods are $24 \leq k \leq 80$, $48 \leq k \leq 80$, and $72 \leq k \leq 80$. Table I compares the mean and the standard deviation of the computation times, over the ten learning sets, to train the point and quantile forecast models².

day-ahead	MLP	LSTM	GBR
point	5.3 (0.1)	23.7 (0.3)	3.4 (0.1)
quantile	7.6 (0.2)	69.0 (0.6)	44.6 (0.4)
intraday	MLP	ED-1	ED-2
point	5.0 (0.1)	5.2 (0.1)	17.2 (0.2)
quantile	17.9 (0.2)	6.4 (0.2)	18.0 (0.3)

TABLE I: Training computation time (s).

C. Day-ahead results

Figure 1a compares the NMAE (plain lines), NRMSE (dashed lines), and Figure 1b the CRPS per forecasting time periods k of the day-ahead models of gate 12:00. Table II provides the mean and standard deviation of the NMAE, NRMSE, and CRPS. The LSTM achieved the best results for both point and quantile forecasts. Figures 3a, 3c, and 3e compare the MLP, LSTM, and GBR day-ahead quantile and point forecasts (black line named dad 12) of gate 12:00 on *August 2, 2020* with the observation in red. One can see that the predicted intervals of the LSTM model better encompass

²The day-ahead and intraday LSTM training times are identical for both point and quantile forecasts as they only take the weather forecasts as inputs.

the actual realizations of uncertainties than the MLP and GBR.




Fig. 1: Day-ahead models results with the NMAE (plain lines), NRMSE (dashed lines), and CRPS.

Score	Gate	MLP	LSTM	GBR
NMAE	12	8.2 (1.2)	7.6 (1.5)	9.2 (0.9)
	24	7.9 (1.2)	7.7 (1.6)	9.0 (0.8)
NRMSE	12	10.2 (1.4)	9.2 (1.6)	11.2 (0.9)
	24	9.7 (1.2)	9.4 (1.8)	10.9 (0.8)
CRPS	12	6.2 (1.1)	4.4 (0.2)	6.4 (0.7)
	24	6.2 (1.0)	4.4 (0.2)	6.3 (0.6)

TABLE II: Day-ahead models results.

D. Intraday results

Table III provides the averaged NMAE, NRMSE, and CRPS per gate of intraday models. The LSTM achieved the best NMAE and NRMSE for the 06:00 gate and the ED-1 achieved the best NMAE and NRMSE for the 12:00 gate and the best CRPS for both gates. Figure 2 compares the CRPS per forecasting time periods k of the intraday models. The ED-1 benefits from the last PV generation observations. Indeed, some CRPS values for both 06:00 and 12:00 gates are below the ones of 00:00 gate. Table IV provides the Interval score of intraday models for 80 %, 60 %, 40 %, and 20 % width of central intervals. The ED-1 model achieved the best results for both 06:00 and 12:00 gates and all prediction intervals except for the 06:00 gate and the prediction interval width of 80 % where it is ED-2. The LSTM achieved close results to the ED-1. Figures 3b, 3d, and 3f compare the ED-1, LSTM, and ED-2 intraday quantile and point forecasts (black line named intra 6) of 06:00 gate on August 2, 2020 with the observation in red. Generally, one can see that the predicted intervals of ED-1 and LSTM models better encompass the actual realizations of uncertainties than ED-2.

VI. CONCLUSION

An encoder-decoder architecture is implemented on the intraday scale to produce accurate forecasts. It efficiently captures the contextual information composed of past PV observations and future weather forecasts, while capturing the temporal dependency between forecasting time periods

Score	Gate	MLP	ED-1	ED-2	LSTM
NMAE	6	8.9 (1.0)	8.5 (1.4)	9.4 (1.0)	7.6 (1.5)
	12	6.7 (1.4)	6.4 (1.3)	7.1 (1.1)	7.2 (1.1)
NRMSE	6	10.9 (0.9)	10.3 (1.3)	11.3 (1.1)	7.7 (1.6)
	12	8.7 (1.3)	7.8 (1.2)	8.5 (1.2)	9.4 (1.8)
CRPS	6	8.1 (0.7)	5.9 (0.9)	6.6 (0.7)	6.2 (0.7)
	12	5.8 (1.2)	4.5 (0.7)	5.6 (1.8)	4.7 (0.5)

TABLE III: Intraday models NMAE, NRMSE, CRPS results.




Fig. 2: Intraday models CRPS results.

Width	Gate	MLP	ED-1	ED-2	LSTM
80 %	6	24.4 (2.9)	14.9 (4.0)	13.9 (4.9)	19.3 (4.2)
	12	17.4 (3.5)	10.6 (1.8)	11.6 (10.1)	9.6 (2.0)
60 %	6	37.6 (3.2)	29.9 (5.0)	32.2 (4.2)	30.7 (4.6)
	12	27.2 (4.3)	22.4 (4.2)	27.5 (10.8)	22.6 (3.1)
40 %	6	58.0 (4.5)	50.1 (6.5)	57.2 (6.0)	51.6 (5.8)
	12	42.4 (6.9)	37.7 (5.9)	48.1 (16.8)	39.2 (4.9)
20 %	6	111.8 (8.4)	97.1 (11.7)	112.1 (10.3)	99.5 (10.4)
	12	81.5 (13.8)	72.7 (10.0)	94.8 (32.1)	76.5 (8.0)

TABLE IV: Intraday models IS results.

over the entire forecasting horizon. The models are compared by using a k -fold cross-validation methodology and quality metrics on a real case study composed of the PV generation of the parking rooftops of the Liège University. The best day-ahead model for both point and quantile forecasts is a neural network composed of a LSTM cell and an additional feed-forward layer. Then, the encoder-architecture composed of a LSTM-MLP yields accurate and calibrated forecast distributions learned from the historical dataset in comparison with the MLP and LSTM-LSTM models for the intraday point and quantile forecasts. However, the LSTM produced similar

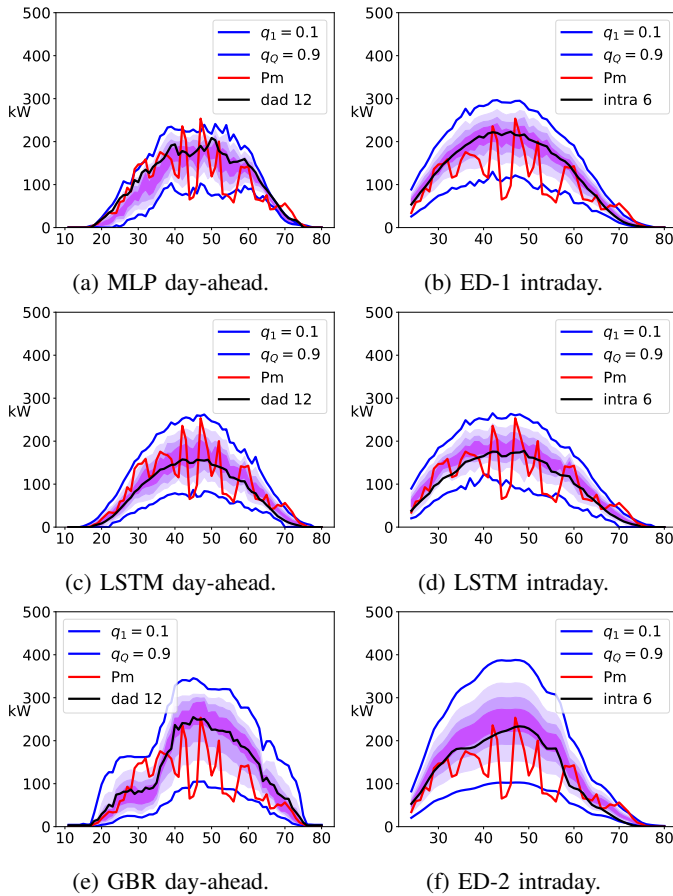


Fig. 3: Quantiles vs point forecasts of day-ahead models of gate 12:00 (left), and intraday models of gate 06:00 (right) on August 2, 2020, the observations are in red.

results. Several extensions are under investigation. First, considering a larger dataset of at least one full year to take into account the entire PV seasonality. Second, developing a PV scenario approach based on the encoder-decoder architecture.

REFERENCES

- [1] J. M. Morales, A. J. Conejo, H. Madsen, P. Pinson, and M. Zugno, *Integrating renewables in electricity markets: operational problems*, vol. 205. Springer Science & Business Media, 2013.
- [2] R. Koenker and G. Bassett Jr, "Regression quantiles," *Econometrica: journal of the Econometric Society*, pp. 33–50, 1978.
- [3] T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R. J. Hyndman, "Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond," 2016.
- [4] F. Golestaneh, P. Pinson, and H. B. Gooi, "Very short-term nonparametric probabilistic forecasting of renewable energy generation—with application to solar energy," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3850–3863, 2016.
- [5] J.-F. Toubeau, J. Bottieau, F. Vallée, and Z. De Grève, "Deep learning-based multivariate probabilistic forecasting for short-term scheduling in power markets," *IEEE Transactions on Power Systems*, vol. 34, no. 2, pp. 1203–1215, 2018.
- [6] J. Bottieau, L. Hubert, Z. De Grève, F. Vallée, and J.-F. Toubeau, "Very-short-term probabilistic forecasting for a risk-aware participation in the single price imbalance settlement," *IEEE Transactions on Power Systems*, vol. 35, no. 2, pp. 1218–1230, 2019.

- [7] P. Lauret, M. David, and H. T. Pedro, "Probabilistic solar forecasting using quantile regression models," *energies*, vol. 10, no. 10, p. 1591, 2017.
- [8] P. Lauret, M. David, and P. Pinson, "Verification of solar irradiance probabilistic forecasts," *Solar Energy*, vol. 194, pp. 254–271, 2019.
- [9] P. Pinson, H. A. Nielsen, J. K. Møller, H. Madsen, and G. N. Kariniotakis, "Non-parametric probabilistic forecasts of wind power: required properties and evaluation," *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, vol. 10, no. 6, pp. 497–516, 2007.
- [10] J. Dumas, B. Cornélusse, A. Giannitrapani, S. Paoletti, and A. Vicino, "Stochastic and deterministic formulations for capacity firming nominations," in *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pp. 1–7, IEEE, 2020.
- [11] X. Fettweis, J. Box, C. Agosta, C. Amory, C. Kittel, C. Lang, D. van As, H. Machguth, and H. Gallée, "Reconstructions of the 1900–2015 greenland ice sheet surface mass balance using the regional climate MAR model," *Cryosphere (The)*, vol. 11, pp. 1015–1033, 2017.
- [12] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
- [13] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, *et al.*, "Scikit-learn: Machine learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.
- [14] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *NIPS-W*, 2017.
- [15] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous systems," 2015.
- [16] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *Journal of the American statistical Association*, vol. 102, no. 477, pp. 359–378, 2007.
- [17] M. Zamo and P. Naveau, "Estimation of the continuous ranked probability score with limited information and applications to ensemble weather forecasts," *Mathematical Geosciences*, vol. 50, no. 2, pp. 209–234, 2018.