Short-term Probabilistic Solar Forecasting via Reinforcement Learning over ECMWF

Jingyi Yan*, Cong Feng†, *Member, IEEE,* Mucun Sun‡, *Member, IEEE,* and Jie Zhang*, *Senior Member, IEEE*
*The University of Texas at Dallas, Richardson, TX, USA
†National Renewable Energy Laboratory, Golden, CO, USA
‡Idaho National Laboratory, Idaho Falls, ID, USA

Abstract—In this paper, we present an innovative reinforcement learning approach for short-term solar forecasting, leveraging data from the European Centre for Medium-Range Weather Forecasts (ECMWF). The methodology begins with the application of the System Advisor Model (SAM) to transform various ECMWF numerical weather prediction members into predictive photovoltaic power generation. To enhance the precision of deterministic forecasting, we introduce a dynamic model selection algorithm based on Q-learning. This algorithm dynamically identifies and utilizes the most accurate ensemble member for forecasting purposes. Furthermore, we employ a support vector regression surrogate model with a Gaussian distribution to generate probabilistic forecasts, providing a holistic view of solar energy generation uncertainty. To expedite the training process and make it more practical for real-world applications, we integrate a rolling update workflow. This innovative workflow reduces the training period from months to a mere 19 days, making our method highly efficient. Numerical results of the case study show that in comparison to benchmark models, the proposed method improves the deterministic and probabilistic solar forecasting accuracy by up to 40.84% and 48.42%, respectively.

Index Terms—solar forecasting, ensemble forecasting, reinforcement learning, probabilistic forecasting

I. INTRODUCTION

The increasing awareness of environmental conservation and the simultaneous reduction in the cost of solar photovoltaic (PV) technology have led to accelerated penetration of solar energy into electricity market. However, the inherent variability of solar energy poses new challenges for stakeholders in the electricity industry. This situation underscores the critical need for improved accuracy in solar forecasting to facilitate optimal planning and decision-making.

As statistically validated in [1] and [2], it's evident that no single forecasting model consistently outperforms others at every timestamp. Ensemble learning has emerged as an effective and adaptable approach for enhancing the accuracy of a suite of forecasting methods, as highlighted by [3]. For example, Carneiro et al. [4] applied Ridge regression ensemble machine learning to reduce forecast errors in intermittent solar and wind resources. Their ensemble model outperformed all machine learning (ML) techniques.

While combining results from diverse forecasting models is one method for constructing an ensemble system, an alternative approach involves consolidating outcomes from various numerical weather predictions (NWP), often referred to as ensemble weather predictions (EWP). However, few studies have delved into the post-processing of EWP disseminated by major operational weather prediction facilities for PV power prediction. Sperati et al. [5] first employed a neural network to generate 51 PV power predictions from 51 ensemble weather sets, provided by the Ensemble Prediction System (EPS) of European Centre for Medium-Range Weather Forecasts (ECMWF). Their study evaluated two statistical calibration methods, estimation of variance deficit (VD) and Ensemble Model Output Statistics (EMOS), which demonstrated comparable probabilistic forecasting performance across various evaluation metrics.

In this paper, the EPS from ECMWF is utilized to generate solar ensemble forecasts. To enhance deterministic ensemble forecasting, we employ a Q-learning-based dynamic model selection (QMS) method, as initially introduced in [2] and [6] for short-term load forecasting. Additionally, we adapt the two-step probabilistic forecasting approach originally outlined in [7] to create short-term probabilistic solar forecasts, utilizing a recursive training process. The main contributions of this paper can be summarized as follows:

- Leveraging EWP for ensemble solar power forecasting.
- Achieving substantial enhancements in both deterministic and probabilistic accuracy, by leveraging Q-learning-based dynamic model selection.
- Introducing a novel rolling update workflow that dramatically reduces the training period to a mere 19 days.

The remaining part of the paper is organized as follows: the proposed methodologies are elaborated in Section II; Section III describes data processing and experiment setup for the case study; discussions of the experiment results are presented in section IV and conclusion follows in Section V.

II. METHODOLOGY

The complete forecasting workflow is illustrated in Fig. 1. The overall methodology can be divided into three parts: 1) the construction of ensemble system, 2) the application of QMS to perform deterministic forecasting, and 3) conversion from deterministic forecasts to probabilistic forecasts using a surrogate model.

A. Construction of ensemble system

Within ECMWF, the NWP capabilities encompass two key components: high-resolution forecasts (HRES) and EPS. Notably, HRES delivers a single forecast, whereas EPS comprises
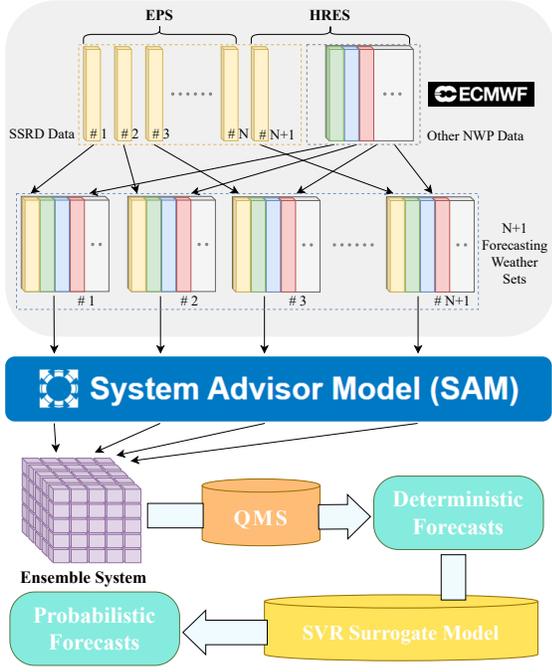
Fig. 1. The overall solar forecasting workflow

one control forecast (CNTL) and 50 forecasts generated with slightly perturbed initial conditions and model physics.

To facilitate ensemble solar PV forecasting, we employ the PVWatts model from the System Advisor Model (SAM), a tool developed by the National Renewable Energy Laboratory (NREL). Specifically, PVWatts, one of SAM's modules, is tailored for modeling grid-connected PV systems and offering irradiance-to-power calculations. For accurate PV power generation predictions, specific NWP data is fed into SAM, encompassing parameters such as global horizontal irradiance (GHI), diffuse horizontal irradiance (DHI), direct normal irradiance (DNI), wind speed, temperature, and solar zenith angle. Among these parameters, GHI, which exhibits a strong correlation with PV energy generation, is singled out as a crucial element within EWP [8], [9].

It's important to note that the primary differentiator among each set of NWP lies in GHI (from both EPS and HRES) and corresponding DHI values. As a result, inputting the $N+1$ sets of NWP data into SAM generates $N+1$ distinct forecasts for PV power generations, which comprises the ensemble system in Fig. 1. Throughout the remainder of this paper, we use $E_i, i = 1, 2, 3, ..., N$ to denote the ith PV forecasting power generation based on the specific NWP set with EPS "Member i". Furthermore, we will use E_{hres} or E_{N+1} to denote the PV power generation forecast based on all HRES weather data.

B. QMS for deterministic solar forecasting

Once an ensemble system is established, QMS, serving as an ensemble algorithm, is employed to select the best ensemble member as the ultimate deterministic forecast. Since the QMS training and processing procedures are described comprehensively in [2], only the QMS algorithm is shown

in Algorithm 1 here, and hyperparameter decisions will be discussed in Section III-B.

Algorithm 1 Q-learning based Dynamic Model Selection (QMS)

Require:
    Q-learning training dataset $\boldsymbol{T}_Q$
    QMS processing dataset $\boldsymbol{P}_Q$
    learning rate α, discount factor γ, number of episodes N_e

Result:
    At each time stamp, select the best ensemble member a_t^* from current state s_t in QMS processing period

1: **Training Procedure:**
2:     Initialize $\boldsymbol{Q} = \vec{\boldsymbol{0}}_{(N+1) \times (N+1)}$, $\epsilon = 1$
3:     **for** $e = 1$ **to** N_e **do**
4:         With the probability of ϵ select a random action a_e from $\{\boldsymbol{a}\}$, otherwise, select $a_e = \arg\max_{a \in A} Q_e(s_e, a)$
5:         Calculate $\boldsymbol{R}$ by

$$R_t(s_i, a_j) = \mathrm{rank}(E_{i,t}) - \mathrm{rank}(E_{j,t+1}) \quad (1)$$

6:         Update $\boldsymbol{Q}$ by

$$\begin{aligned} Q_{e+1}(s_e, a_e) = &(1 - \alpha)Q_e(s_e, a_e) + \\ &\alpha[R_e(s_e, a_e) + \gamma \max_{a \in A} Q_e(s_{e+1}, a)] \end{aligned} \quad (2)$$

7:         $\epsilon \leftarrow \epsilon - \frac{1}{N_e}$
8:     **end for**
9: **Processing Procedure:**
10:     Take action $a_t^* = \arg\max_{a \in A} \boldsymbol{Q}^*(s_t, a)$

C. Probabilistic forecasting based on CRPS optimization

To effectively account for the inherent high uncertainty in solar forecasting, we employ a two-step probabilistic forecasting method. This approach involves conducting probabilistic forecasts on top of deterministic forecasts.

To gauge the accuracy of these probabilistic forecasts, we utilize the Continuous Ranked Probability Score (CRPS). CRPS is a mathematical function that quantifies the dissimilarity between the predicted and observed cumulative distributions [10]. For specific timestamp t, the mathematical expression of CRPS is presented as follows:

$$CRPS_t = \int (F(x_t) - O(x_t, y_t))^2 dx_t \quad (3)$$

where $F(x_t)$ is the cumulative distribution function (CDF) of the forecast quantity x_t at time stamp t, and $O(x_t, y_t)$ is the observed CDF, which is also a Heaviside function associated with the forecast quantity x_t and the observed value y_t,

$$O(x_t, y_t) = \begin{cases} 0, & x_t < y_t \\ 1, & x_t > y_t \end{cases} \quad (4)$$

To implement the CRPS optimization based probabilistic solar forecasting [7], we consider four different predictive distribution types, namely Gaussian, Gamma, Laplace, and logistic distributions, all characterized by probability density

functions that depend on two parameters: the mean (μ) and the standard deviation (σ), expressed as $f(q|\mu, \sigma)$, with 'q' representing the forecast quantile.

During the training process, all deterministic forecasts obtained from the QMS serve as the means of the forecast distribution. To optimize the CRPS, we employ a genetic algorithm (GA) to find the optimal sigma, σ^*, within predefined boundaries, specifically, $\sigma_l \leq \sigma^* \leq \sigma_u$. After all pair of (μ, σ^*) are generated by GA, a support vector regression (SVR) surrogate model is utilized to learn the relationship between these parameters. In the forecasting phase, this surrogate model can estimate the predictive standard deviations of the probabilistic forecasts based on the means, which are essentially the deterministic forecasts provided by the QMS.

For a more in-depth exploration of the two-step forecasting method, please refer to [7]. However, it's important to note that, in contrast to pre-training the probabilistic forecasting model, this paper implements a recursive training process. In this approach, the forecasting model continuously adapts and fine-tunes itself as new observed data becomes available.

III. CASE STUDY

A. Data summary

In order to protect business privacy and facilitate experiment replication, the proposed short-term solar forecasting method is applied within a scenario that uses an alternative source for weather data and power conversion. Specifically, we've utilized weather data from the National Solar Radiation Database (NSRDB). A hypothetical PV power plant located in Austin, TX, is constructed. The hypothetical system configurations are listed in Table I. The actual power generation data from the years 2018 and 2019 serve as the training and testing datasets.

To construct the NWP set for our ensemble system, 20 members of surface solar radiation downwards (SSRD) predictions, which have been converted to represent GHI, are obtained from the EPS system. Direct solar radiation (DSRP, converted to represent DNI), wind speed at 10 meters in the U and V directions, temperature and another set of SSRD are from the HRES system. It's important to note that all of these parameters are measured at the surface level. However, the daily issuing time of the data from EPS is 6:00 UTC, whereas HRES data become available at 0:00 UTC, due to their distinct operational schedules.

After converting EWP into power generation, two-year data are available for QMS and probabilistic forecasting. The first two months of data, ranging from 01-01-2018 to 02-28-2018, are designated for training and fine-tuning model hyperparameters. Subsequently, the complete two-year dataset is employed for testing purposes. Due to the diurnal variations in solar power generation, our analysis focuses on data recorded from 6:00 am to 8:00 pm local time. Aligning with the mechanism of QMS, which performs forecast selection based on the performance of the previous hour's forecasts, our solar forecasting methodology provides hourly forecasts.

TABLE I
HYPOTHETICAL PV SYSTEM CONFIGURATIONS

Attributes	Value	Unit
Longitude, latitude	(30.25,-97.75)	degrees
Timezone	-6	hours offset from GMT
Site elevation	199	meters above see level
System capacity	1	MW
DC to AC ratio	1.1	\
Tilt angle	28	degrees
Inverter efficiency	96	%
Azimuth angle	180	degrees from the north
System losses	14	%
Array type	fixed	\
Ground coverage ratio	0.4	\
Constant loss adjustment	0	%

B. Hyperparameters

Following the hyperparameter tuning process, the optimal hyperparameters for both the QMS model and the probabilistic forecasting model are determined. It's important to note that both models feature a moving window, encompassing both training and forecasting components. These training sets are refreshed daily, with the training set length for the QMS spanning 10 days and for the probabilistic model, it spans 9 days. Additionally, the Q-table within the QMS is updated every 5 hours during the processing period.

Furthermore, 12 best ensemble members with the lowest normalized Mean Absolute Error (nMAE) are selected from the training set before enabling the QMS for each day, which reduces the computation complexity. Other crucial hyperparameters for Q-learning are decided as follows: $N_e = 300$, $\alpha = 0.5$, and $\gamma = 0.6$. Following a comparative analysis of the results during the tuning period, the Gaussian distribution is selected as the predictive distribution for our probabilistic forecasting model.

IV. RESULTS

To demonstrate the effectiveness of the ensemble algorithm, we consider all results generated within the ensemble system as a basis for deterministic forecasting comparisons. Moreover, two essential ensemble forecasting methods are used for comparison in the case study, i.e., the model averaging method and the persistence ensemble method, labeled as $AveEn$ and $PerEn$, respectively. The former one outputs the average results from the ensemble system as the ensemble forecasts, whilst the later selects the best model, based on absolute percentage error performance, to perform deterministic forecasting in the upcoming step. Besides the direct quantifying predictive PV generation based on the complete dataset of 50 EPS ensemble members and the HRES forecasts, another daily updated quantile regression (QR) method based on the same 9-day training set for the proposed model is treated as a probabilistic benchmark.

A. Deterministic solar forecasting results

Four different evaluation metrics are used to evaluate the deterministic forecasting performance, which are nMAE, mean

absolute percentage error (MAPE), mean bias error (MBE), and mean square error (MSE). The overall performance of QMS and other deterministic forecasts are compared in Table II. To maintain brevity, we showcase only the benchmark ensemble members that exhibit the lowest or highest absolute values of evaluation metrics. Among ensemble members, the lowest absolute errors are highlighted in green, while the largest in pink.

Notably, the proposed QMS forecasts excel in terms of nMAE, MAPE, and MSE, achieving the smallest values among the comparisons. And the HRES system in ECMWF does prove itself with the least values in three out of four metrics compared to other forecasting members from EPS. The MBE value, which accumulates bias error, reflects an overall negative bias in the ensemble system, indicating a tendency toward overestimation. The QMS alleviates this overestimation as decreasing the absolute MBE value. Surprisingly, none of the comparative ensemble methods demonstrated superior performance compared to deterministic forecasting based on HRES weather data, except in the case of the MAPE criterion, underscoring the significance of the ensemble algorithm. It highlights that an ineffective ensemble method has the potential to compromise the performance of an otherwise distinguished model within the ensemble system.

The maximum improvement of the QMS deterministic forecasts over the ensemble system is achieved under MAPE with a value of 40.84%. The average improvements of nMAE, MAPE, MBE, and MSE are 16.89%, 30.28%, 16.63%, and 30.04%, respectively.

Fig. 2 illustrates the deterministic PV generation forecasting of proposed and comparative methods. Specifically, we have selected forecasts for the periods from 2018-07-01 to 2018-07-02 and from 2019-10-01 to 2019-10-02 to illustrate the forecasting performance under low and high uncertainty scenarios. It's important to emphasize that we use the same calculator, SAM, for both generating forecasts and calculating the target PV generation from the weather data. Consequently, any disparities between the target and prediction curves primarily arise from the weather forecasts obtained from ECMWF, reflecting the inherent unpredictability of meteorological conditions. However, even in the face of these meteorological uncertainties, the QMS model substantially improves the overall deterministic forecasting accuracy, as the QMS curves closely align with the target curves, outperforming other forecasting methods, particularly under high uncertainty situations. In Fig. 2(c), QMS model performs the best among all comparative models from 6 am to 12 pm. After two hours missing the best forecasting models under chaotic weather, the QMS quickly selects the high-accuracy model back in the third hour.

B. Probabilistic solar forecasting results

Fig. 3 illustrates the probabilistic forecasts of the proposed method, QR method, and quantified ensemble method, labeled as Prob, QR, and ENS, respectively. For each method, two solid lines circle out the central 90% probabilistic prediction interval, and the darker ribbon covers the 30% prediction

TABLE II
EVALUATION METRICS OF DETERMINISTIC FORECASTS

Models	nMAE	MAPE[%]	MBE	MSE
E_3	0.0925	2.04	-0.0175	0.0143
E_7	0.0926	2.97	-0.0168	0.0144
E_{10}	0.0932	2.44	-0.0182	0.0147
E_{12}	0.0938	2.58	-0.0173	0.0149
E_{hres}	0.0886	2.61	-0.0143	0.0129
$AveEn$	0.0891	2.43	-0.0149	0.0131
$PerEn$	0.0923	2.20	-0.0161	0.0144
QMS	**0.0770**	**1.76**	**-0.0139**	**0.0101**
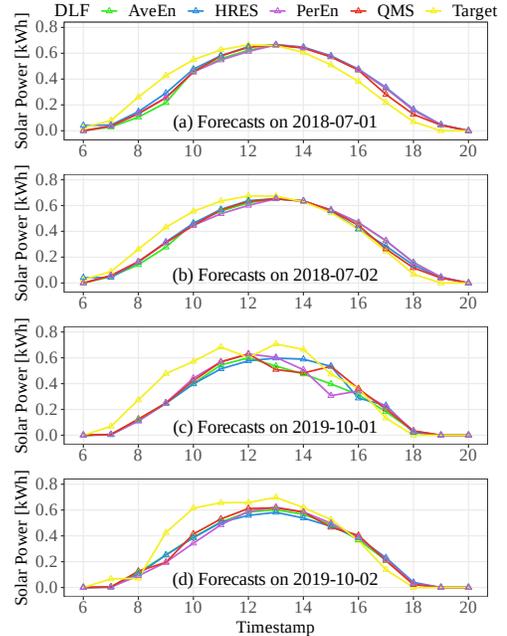
Fig. 2. Time series of deterministic solar forecasts under low (a-b) and high (c-d) uncertainty weather conditions. The superior performance of the proposed method is evidenced by the enhanced alignment between QMS prediction (depicted by red curves) and target values (illustrated by yellow curves), surpassing that of other benchmark methods.

interval. From both Fig. 3(a) and 3(b), it becomes evident that the benchmark forecasting, which is based on quantified predictive PV generation from the EPS and HRES systems, lacks the appropriate level of spread. This observation is also observed in [5] and [11]. While the QR probabilistic forecasting excels in providing widespread prediction intervals, its performance diminishes under conditions of high uncertainty, such as during the time period between 8 - 10 am on 2019-10-01. The proposed probabilistic method shrinks it's prediction interval under less uncertainty weather, while enlarges enough interval confronting high uncertainty weather.

The probabilistic forecasting performance is validated with the average pinball loss and CRPS metrics, as summarized in Table III. Both criteria are evaluated across a range of quantiles, spanning from 5% to 95%. While the pinball loss focuses on the deviation of a specific quantile q within probabilistic forecast, the CRPS takes into account the whole distribution in a continuous manner. In both cases, superior performance is indicated by smaller values. As demonstrated in Table III, our proposed probabilistic forecasting method outperforms the
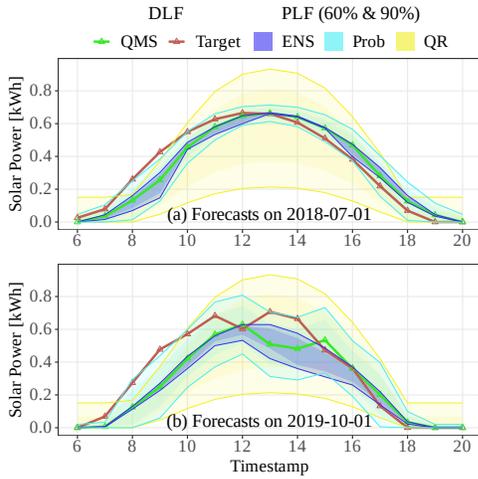
Fig. 3. Time series of probabilistic solar forecasts under low (a) and high (b) uncertainty weather conditions. For each probabilistic forecasting method, two solid lines circle out the central 90% probabilistic prediction interval, and the darker ribbon covers the 30% prediction interval. The blue ribbon (90% probabilistic forecasts of the proposed method) better covers the target curve considering the uncertainty nature of meteorology.

benchmarks under both criteria. The maximum improvement of the proposed method over the benchmarks in terms of pinball loss and CRPS are 47.42% and 48.42%, respectively.

A probabilistic forecasting model is deemed reliable if, and only if, for any forecast interval between the upper and lower quantiles, the distribution of target values falling within that interval across the entire testing dataset precisely matches the length of the forecast interval [12], [13]. The correlations between the nominal proportion and the estimated coverage are presented in Fig. 4. The reliability criterion is met when the estimated coverage aligns with the nominal proportion, as indicated by the dashed line in Fig. 4. The better overall convergence of the green curve with the black dashed curve illustrates the better reliability of the proposed method.

TABLE III
EVALUATION METRICS OF PROBABILISTIC FORECASTS

	Average Pinball Loss	Average CRPS
Prob	0.0227	0.0464
QR	0.0432	0.0899
ENS	0.0330	0.0528
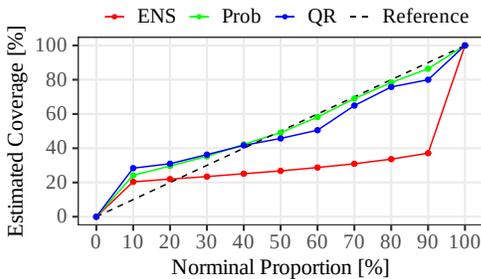
Fig. 4. Reliability comparison of probabilistic solar forecasts

V. CONCLUSIONS

This paper developed a method of reinforced learning based short-term deterministic and probabilistic solar ensemble forecasts. To enhance deterministic forecasts obtained from ECMWF, a Q-learning Model Selection (QMS) method was employed. Subsequently, a surrogate model was employed to conduct probabilistic forecasting, optimized based on CRPS. Numerical simulations confirm that the proposed method outperforms the benchmark approaches. The QMS improves the deterministic forecasting accuracy, reducing the absolute errors of ensemble members by an average of 16.89%, 30.28%, 16.63%, and 30.04% for nMAE, MAPE, MBE, and MSE, respectively. The success of the probabilistic forecasting is underscored by up to 47.42% improvement in pinball loss and 48.42% improvement in CRPS compared to the benchmarks.

Future research will focus on (1) evaluating forecasting accuracy, the length of the training set, and computational load with different numbers of EPS members for ensemble system construction; (2) conducting comparisons with EWP from other operational weather prediction facilities, exploring variations both vertically with different ensemble member counts and horizontally with data from diverse facilities.

REFERENCES

[1] G. M. Yagli, D. Yang, and D. Srinivasan, "Automatic hourly solar forecasting using machine learning models," *Renewable and Sustainable Energy Reviews*, vol. 105, pp. 487–498, May 2019.

[2] C. Feng, M. Sun, and J. Zhang, "Reinforced Deterministic and Probabilistic Load Forecasting via Q-Learning Dynamic Model Selection," *IEEE Transactions on Smart Grid*, vol. 11, pp. 1377–1386, Mar. 2020.

[3] C. Zhang and Y. Ma, eds., *Ensemble Machine Learning: Methods and Applications*. New York, NY: Springer New York, 2012.

[4] T. C. Carneiro, P. A. C. Rocha, P. C. M. Carvalho, and L. M. Fernández-Ramírez, "Ridge regression ensemble of machine learning models applied to solar and wind forecasting in Brazil and Spain," *Applied Energy*, vol. 314, p. 118936, May 2022.

[5] S. Sperati, S. Alessandrini, and L. Delle Monache, "An application of the ECMWF Ensemble Prediction System for short-term solar power forecasting," *Solar Energy*, vol. 133, pp. 437–450, Aug. 2016.

[6] C. Feng and J. Zhang, "Reinforcement Learning based Dynamic Model Selection for Short-Term Load Forecasting," in *2019 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, (Washington, DC, USA), pp. 1–5, IEEE, Feb. 2019.

[7] M. Sun, C. Feng, E. K. Chartan, B.-M. Hodge, and J. Zhang, "A two-step short-term probabilistic wind forecasting methodology based on predictive distribution optimization," *Applied Energy*, vol. 238, pp. 1497–1505, Mar. 2019.

[8] Y. Huang, J. Lu, C. Liu, X. Xu, W. Wang, and X. Zhou, "Comparative study of power forecasting methods for PV stations," in *2010 International Conference on Power System Technology*, pp. 1–6, Oct. 2010.

[9] U. K. Das, K. S. Tey, M. Seyedmahmoudian, S. Mekhilef, M. Y. I. Idris, W. Van Deventer, B. Horan, and A. Stojcevski, "Forecasting of photovoltaic power generation and model optimization: A review," *Renewable and Sustainable Energy Reviews*, vol. 81, pp. 912–928, Jan. 2018.

[10] H. Hersbach, "Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems," *Weather and Forecasting*, vol. 15, pp. 559–570, Oct. 2000. Publisher: American Meteorological Society Section: Weather and Forecasting.

[11] W. Wang, D. Yang, T. Hong, and J. Kleissl, "An archived dataset from the ECMWF Ensemble Prediction System for probabilistic solar power forecasting," *Solar Energy*, vol. 248, pp. 64–75, Dec. 2022.

[12] J. Bröcker and L. A. Smith, "Increasing the Reliability of Reliability Diagrams," *Weather and Forecasting*, vol. 22, pp. 651–661, June 2007.

[13] I. T. Jolliffe and D. B. Stephenson, eds., *Forecast verification: a practitioner's guide in atmospheric science*. Chichester, West Sussex, England ; Hoboken, NJ: J. Wiley, 2003.