

# Tackling Missing Values in Probabilistic Wind Power Forecasting: A Generative Approach

Honglin Wen\*, Pierre Pinson†, Jie Gu\*, Zhijian Jin\*,

\* Department of Electrical Engineering, Shanghai Jiao Tong University, Shanghai, China.

† Dyson School of Design Engineering, Imperial College London, London, United Kingdom.

**Abstract**—Machine learning techniques have been successfully used in probabilistic wind power forecasting. However, the issue of missing values within datasets due to sensor failure, for instance, has been overlooked for a long time. Although it is natural to consider addressing this issue by imputing missing values before model estimation and forecasting, we suggest treating missing values and forecasting targets indifferently and predicting all unknown values simultaneously based on observations. In this paper, we offer an efficient probabilistic forecasting approach by estimating the joint distribution of features and targets based on a generative model. It is free of preprocessing, and thus avoids introducing potential errors. Compared with the traditional “impute, then predict” pipeline, the proposed approach achieves better performance in terms of continuous ranked probability score.

**Index Terms**—forecasting, wind power, missing values, generative model.

## I. INTRODUCTION

Though renewable energy is commonly acknowledged as the workhorse for carbon neutrality, its inherent uncertainty challenges power systems operation and electricity markets. Therefore, forecasting is deemed an indispensable tool for system operators and has developed for decades; see a recent review [1]. Of particular interest is probabilistic wind power forecasting [2], which leverages information such as weather and lagged observations up to the current time to communicate the probability of wind power generation at a future time in terms of densities, quantiles, prediction intervals, etc.

Typically, forecasting models can be developed using either parametric or nonparametric approaches with state-of-the-art machine learning techniques. The parametric approach relies on distributional assumptions such as Gaussian, Beta, and Logit-normal distributions, with shape parameters estimated through statistical and machine learning methods [3], [4]. In contrast, non-parametric approaches are free of such assumptions. Among these, quantile regression (QR) [5] is the most popular due to its success in forecasting competitions and ease of use. Particularly, by using novel machine learning models such as gradient boosting machine and N-BEATS [6], [7], the performance of QR could be considerably improved. However, the notorious “quantile crossing” phenomena in QR can be hardly avoided. Consequently, increasing works are focusing on non-parametric density forecasting approaches [8], [9].

Along with advancements in probabilistic wind power forecasting research, data are becoming the fuel for powerful data-driven models. However, missing values in data are prevalent in the energy sector, posing challenges for both

model estimation and forecasting [10]. An intuitive approach is to select samples based on observed missingness patterns, and retrain models for each pattern. However, this requires prohibitive computational effort, as the number of potential missing patterns grows exponentially with the number of features. Though it is natural to impute missing values prior to model estimation and forecasting (referred to as “impute, then predict”), the imputation may introduce errors that ultimately degrade the quality of forecasts. Consequently, it remains an open issue to design probabilistic wind power forecasting approaches in the presence of missing values.

The seminal work in time series forecasting with missing observations can date back to [11], which represents ARMA models in the state-space form. In this approach, the state update equation is bypassed at times when observations are missing. This was later extended to ARIMA models in [12], [13]; however, the focus is still confined to linear models and point forecasts. A robust optimization approach for the regression framework is proposed in [14], assuming some features are complete. Parameters are optimized by minimizing the worst-case loss over the missingness. However, it only addresses model training issues aroused by missing values, leaving operational forecasting challenges unsolved.

Indeed, both the missing features and the forecasting targets are unknown, and their prediction can be indifferently deduced from observations. Specifically, we can incorporate the missing features and targets into a distributional learning framework. This allows us to estimate a joint distribution using partially observed data, which is termed the “universal imputation” strategy. During the operational forecasting stage, forecasts can be obtained through the marginalization of missing features. Building upon this concept, we formulated a probabilistic forecasting model using the fully conditional specification (FCS) approach in our prior research [15]. Nevertheless, implementing the FCS necessitates training a model independently for each probabilistic scenario in an iterative fashion. As a result, its computational complexity poses a significant obstacle, impeding its practical applicability.

In this work, we tackle the computational challenge discussed in [15] using a generative approach. Specifically, we develop a generative model employing an auto-encoder [16] alongside deep learning techniques. This model assumes observations to be mappings of latent variables from an unknown distribution. During the training phase, we concatenate feature-target pairs as instances and utilize them to estimate the parameters of the generative model. The evidence lower bound

is obtained by using the method proposed in [17]. Specifically, we employ a normalizing flow model [18] to represent the variational distribution, enabling a more flexible approximation of the posterior distribution. As such, it is feasible to approximately maximize the likelihood of observations. During the forecasting stage, we randomly sample scenarios from the the latent variable distribution and map them to values that then act as observation proposals. Then, we assign weights to each proposal sample based on the likelihood of actual observations and employ these weights for resampling. After that, the forecasts for targets are generated through marginalization over missing variables. We demonstrate the superiority of the proposed approach based on an open dataset. It surpasses the routine ‘‘impute, then predict’’ strategy in the context of continuous ranked probability score (CRPS) and significantly enhances computational efficiency as compared to the FCS model in [15].

The contexts are organized as follows. We describe the preliminaries in Section II and formulate the problem in Section III. The generative model is introduced in Section IV, whereas the case study is performed in Section V. Results and discussion are presented in Section VI. We conclude the paper in Section VII.

## II. PRELIMINARIES

In this section, we will introduce the concept of probabilistic wind power forecasting, the definition of missingness mechanisms, and the theorem for inference with missing values.

### A. Probabilistic Wind Power Forecasting

For simplicity, let us consider the univariate wind power forecasting with a special focus on very short-term cases, where missingness often occurs. Let  $Y_t$  denote the random variable for wind power generation value at time  $t$ , and  $y_t$  its realization. Probabilistic wind power forecasting aims at communicating the probability of wind power generation at time  $t + k$  (also referred to as targets) with the information up to time  $t$  (also referred to as features), based on model  $\mathcal{M}$  with parameters  $\theta$ . The information is often composed of previous values of length  $h$ , i.e.,  $y_{t-h+1}, y_{t-h+2}, \dots, y_{t-1}, y_t$ . For convenience, we write them as  $\mathbf{y}_t$  and the corresponding random variable as  $\mathbf{Y}_t$ . Probabilistic forecasting is described as

$$\hat{p}_{Y_{t+k}}(y_{t+k}) = p(y_{t+k} | y_{t-h+1}, y_{t-h+2}, \dots, y_{t-1}, y_t; \mathcal{M}, \theta). \quad (1)$$

The model  $\mathcal{M}$  could be either parametric or non-parametric (w.r.t distributional assumption). Particularly, we assume the underlying process is stationary, therefore parameters  $\theta$  can be estimated based on historical data  $y_1, y_2, \dots, y_T$ . The estimate for  $\theta$  is denoted as  $\hat{\theta}$ .

### B. Missingness Mechanism

Let random variable  $M_t$  represent whether  $y_t$  is missing,  $m_t \in \{0, 1\}$  its realization.  $m_t = 1$  implies  $y_t$  is missing, whereas  $m_t = 0$  implies  $y_t$  is observed. Accordingly, the corresponding mask variable for the variable  $\mathbf{Y}_t$  is written as  $\mathbf{M}_t$ ,

whose realization is  $\mathbf{m}_t \in \{0, 1\}^h$ . In the modern statistical theory [19], missing mechanisms can be classified into three categories: missing completely at random (MCAR), missing at random (MAR), and missing not at random (MNAR). Taking variables  $\mathbf{Y}_t$  and  $\mathbf{M}_t$  as an example. The parametric model for the joint distribution of the data sample and its mask is described as

$$p(\mathbf{y}_t, \mathbf{m}_t; \theta, \psi) = p(\mathbf{y}_t; \theta) p(\mathbf{m}_t | \mathbf{y}_t; \psi), \quad (2)$$

where  $\theta$  and  $\psi$  represent the parameters of the distribution of data and masks respectively. The data sample can be split into an observed part  $\mathbf{y}_t^o$  and a missing part  $\mathbf{y}_t^m$ , i.e.,  $\mathbf{y}_t = (\mathbf{y}_t^o, \mathbf{y}_t^m)$ . Then the missingness mechanisms are respectively defined as

- MCAR:  $p(\mathbf{m}_t | \mathbf{y}_t; \psi) = p(\mathbf{m}_t; \psi)$ ,
- MAR :  $p(\mathbf{m}_t | \mathbf{y}_t; \psi) = p(\mathbf{m}_t | \mathbf{y}_t^o; \psi)$ .
- MNAR:  $p(\mathbf{m}_t | \mathbf{y}_t; \psi) = p(\mathbf{m}_t | \mathbf{y}_t^o, \mathbf{y}_t^m; \psi)$ ,

That is, the MCAR mechanism is independent of the data sample. In contrast, the MAR mechanism is dependent on  $\mathbf{y}_t^o$  yet independent of  $\mathbf{y}_t^m$ , whereas the MNAR mechanism is dependent on both the observed and missing values.

As  $\mathbf{y}_t$  can be decomposed into an observed part  $\mathbf{y}_t^o$  and a missing part  $\mathbf{y}_t^m$ , The likelihood can be derived by integrating over missing values, i.e.,

$$p(\mathbf{y}_t^o, \mathbf{m}_t; \theta, \psi) = \int_{\mathbf{y}_t^m} p(\mathbf{y}_t^o, \mathbf{y}_t^m; \theta) p(\mathbf{m}_t | \mathbf{y}_t^o, \mathbf{y}_t^m; \psi) d\mathbf{y}_t^m. \quad (3)$$

In addition, the likelihood of observed data is

$$p(\mathbf{y}_t^o, \mathbf{m}_t; \theta) = \int_{\mathbf{y}_t^m} p(\mathbf{y}_t^o, \mathbf{y}_t^m; \theta) d\mathbf{y}_t^m. \quad (4)$$

Rubin proved inference for  $\theta$  on the condition of MAR mechanism [20], which is stated as:

**Theorem 1 (Theorem 7.1 in [20])** *Let  $\psi$  such that for all  $t$ ,  $p(\mathbf{m}_t | \mathbf{y}_t; \psi) > 0$ . Assuming data are MAR,  $p(\mathbf{y}_t^o, \mathbf{m}_t; \theta, \psi)$  is proportional to  $p(\mathbf{y}_t^o, \mathbf{m}_t; \theta)$  w.r.t.  $\theta$ , so that inference for  $\theta$  can be obtained by maximizing  $p(\mathbf{y}_t^o, \mathbf{m}_t; \theta)$  while ignoring the mechanism.*

Here we assume wind power data are MAR, as the missingness is often caused by sensor failures and communication errors, which is irrelevant to sample values. As a result, we can only focus on the distribution of data and leave the modeling of missingness aside.

## III. PROBLEM FORMULATION

Missingness may occur in both features and targets, which raises challenges to the parameter estimation stage for the model (1), as the model predictions and gradients cannot be calculated in the presence of missing values. In fact, the conditional distribution can be derived through marginalization if the joint distribution  $p(\mathbf{y}_t, y_{t+k}; \mathcal{M}, \theta)$  is available. Here, we denote the model as  $\mathcal{M}'$ , and the parameters as  $\theta$ . The forecast for  $Y_{t+k}$  is described as

$$\hat{p}_{Y_{t+k}} = \frac{p(\mathbf{y}_t, y_{t+k}; \mathcal{M}', \theta)}{\int_{\mathbf{y}_t} p(\mathbf{y}_t, y_{t+k}; \mathcal{M}', \theta) d\mathbf{y}_t}. \quad (5)$$




Fig. 1. Transition from conditional distribution modeling to joint distribution modeling, where blue blocks indicate observed values and yellow blocks indicate unobserved values.

Then, the forecasting problem boils down to a distribution modeling problem in the presence of missing values. With Theorem 1, we know that we are allowed to estimate the distribution  $p(\mathbf{y}_t, y_{t+k}; \mathcal{M}', \theta)$  with missing values when the missingness mechanism is MAR. We illustrate the model framework in Fig. 1, where blue boxes indicate observed values and yellow boxes indicate unobserved values.

For notational simplicity, we concatenate  $\mathbf{Y}_t$  and  $Y_{t+k}$ , and write it as  $\mathbf{Z}_t = [\mathbf{Y}_t^\top, Y_{t+k}]^\top$ , its realization as  $\mathbf{z}_t = [\mathbf{y}_t^\top, y_{t+k}]^\top \in \mathbb{R}^{h+1}$ , and the mask as  $\mathbf{s}_t = [\mathbf{m}_t^\top, m_{t+k}]^\top \in \{0, 1\}^{h+1}$ . Accordingly, the observed and missing parts of  $\mathbf{z}_t$  are written as  $\mathbf{z}_t^o$  and  $\mathbf{z}_t^m$  respectively. While at the forecasting stage, we derive the conditional distribution of  $\mathbf{z}_t^m$ , i.e.,

$$p(\mathbf{z}_t^m | \mathbf{z}_t^o; \mathcal{M}', \hat{\theta}) = \frac{p(\mathbf{z}_t^m, \mathbf{z}_t^o; \mathcal{M}', \hat{\theta})}{\int_{\mathbf{z}_t^m} p(\mathbf{z}_t^m, \mathbf{z}_t^o; \mathcal{M}', \hat{\theta}) d\mathbf{z}_t^m}.$$

In fact,  $\mathbf{z}_t^m$  is composed of  $\mathbf{y}_t^m$  and  $y_{t+k}$ , i.e.,  $\mathbf{z}_t^m = [\mathbf{y}_t^m, y_{t+k}]$ . Then, the forecast for  $y_{t+k}$  can be derived via marginalization:

$$p(y_{t+k} | \mathbf{z}_t^o; \mathcal{M}', \hat{\theta}) = \int_{\mathbf{y}_t^m} p(\mathbf{z}_t^m | \mathbf{z}_t^o; \mathcal{M}', \hat{\theta}) d\mathbf{y}_t^m. \quad (6)$$

To sum up, the challenge at the model estimation stage is to learn the distribution  $p(\mathbf{z}_t; \mathcal{M}', \hat{\theta})$  based on incomplete data  $\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_T\}$ . While at the forecasting stage, the challenge lies in the efficient calculation of marginalization.

#### IV. METHODOLOGY

In this section, we describe the model architecture at first, followed by the estimation procedure. Then, we show how to use the estimated model for operational forecasting by using the techniques derived in [17].

##### A. Model Architecture

Particularly, we establish the model based on the variational auto-encoder (VAE) [16], which assumes data are generated from a latent variable  $\mathbf{u}_t \sim p(\mathbf{u})$  via a *decoder* model  $p(\mathbf{z}_t | \mathbf{u}_t; \theta)$ , i.e.,

$$\mathbf{u}_t \sim p(\mathbf{u}), \quad \mathbf{z}_t \sim p(\mathbf{z}_t | \mathbf{u}_t; \theta). \quad (7)$$

Usually, the prior distribution of the latent variable is set as a normal distribution  $\mathcal{N}(\mathbf{z} | \mathbf{0}, \mathbf{I})$  where  $\mathbf{0}$  represents a vector of 0 and  $\mathbf{I}$  represents a diagonal matrix of 1. And, it has become a mainstream approach to leverage variational inference to estimate the parameters  $\theta$ . For that, a posterior approximation namely the *encoder* model  $q(\mathbf{u}_t | \mathbf{z}_t; \phi)$  is required, which is assumed to approximate  $p(\mathbf{u}_t | \mathbf{z}_t)$ . In this work, of particular interest is approximating the posterior distribution via the observed part, i.e.,  $q(\mathbf{u}_t | \mathbf{z}_t^o; \phi)$ . Then, we can recover the unknown part  $\mathbf{z}_t^m$  via  $p(\mathbf{z}_t | \mathbf{u}_t; \theta)$ . The graphical model of data recovering is illustrated in Fig. 2. For computational ease, we use a function  $g$  to convert  $\mathbf{z}_t$  into a value in the Real number space, i.e.,

$$g(\mathbf{z}_t) \in \mathbb{R}^{h+1}, \quad g(\mathbf{z}_t^o) = \mathbf{z}_t^o.$$

The function could be implemented as imputing missing values with 0.




Fig. 2. Illustration of how missing values are recovered by using the encoder and decoder models.

In this work, we use a Student's t-distribution to model the decoder, i.e.

$$p(\mathbf{z}_t | \mathbf{u}_t; \theta) = St(\mathbf{z}_t | \boldsymbol{\mu}(\mathbf{u}_t; \theta), \boldsymbol{\Sigma}(\mathbf{u}_t; \theta), \boldsymbol{\nu}(\mathbf{u}_t; \theta)), \quad (8)$$

where  $\boldsymbol{\mu}(\mathbf{u}_t; \theta)$ ,  $\boldsymbol{\Sigma}(\mathbf{u}_t; \theta)$ ,  $\boldsymbol{\nu}(\mathbf{u}_t; \theta)$  are the shape parameters estimated via neural networks. To obtain more faithful approximation for the posterior distribution, we use the strategy in [18], i.e., using a base distribution together with a normalizing flow to model the posterior. Specifically, the base distribution is set as a Gaussian distribution. To make a distinction, we here write the base variable for  $\mathbf{u}_t$  as  $\mathbf{u}_t^{(0)}$ . Then  $\mathbf{u}_t^{(0)}$  follows:

$$\mathbf{u}_t^{(0)} \sim \mathcal{N}(\boldsymbol{\mu}(g(\mathbf{z}_t); \phi), \boldsymbol{\Sigma}(g(\mathbf{z}_t); \phi)) \quad (9)$$

The normalizing flow operates as a chain of  $N$  transforms, i.e.,

$$\mathbf{u}_t^{(N)} = f_N \circ \dots \circ f_2 \circ f_1(\mathbf{u}_t^{(0)}), \quad (10)$$

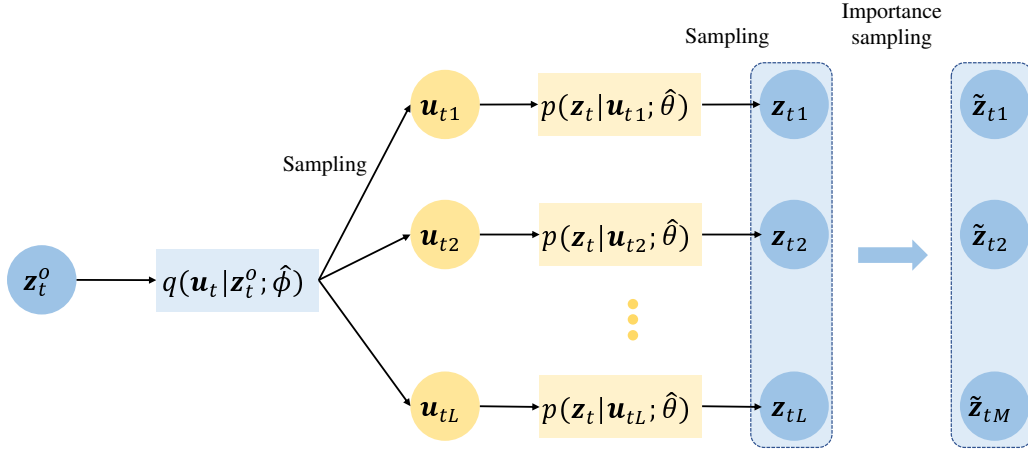


Fig. 3. Sampling based on the decoder and encoder models for operational forecasting.

where  $\circ$  represents the composition of functions. For the  $n$ -th transform, it operates on  $\mathbf{u}_t^{(n-1)}$  and yields  $\mathbf{u}_t^{(n)}$ , i.e.,

$$\mathbf{u}_t^{(n)} = f_n(\mathbf{u}_t^{(n-1)}), \quad (11a)$$

$$\mathbf{u}_t^{(n-1)} = f_n^{-1}(\mathbf{u}_t^{(n)}), \quad (11b)$$

where  $f_n^{-1}$  is the inverse function of  $f_n$ . Specially, the final output  $\mathbf{u}_t^{(N)}$  is  $\mathbf{u}_t$ . The log-likelihood is then derived as

$$\log q(\mathbf{u}_t | \mathbf{z}_t^o) = \log q(\mathbf{u}_t^{(0)}) - \sum_n \log \left| \det \frac{\partial f_n}{\partial \mathbf{u}_t^{(n-1)}} \right|, \quad (12)$$

where  $\det$  represents the determinant operator. Specifically, we use autoregressive flow as in [9].

### B. Model Estimation Stage

In statistical learning, the parameters in  $p(\mathbf{z}_t; \mathcal{M}', \theta)$  are often learned via the maximum likelihood estimation, i.e.,

$$\hat{\theta} = \arg \max_{\theta} \sum_t \ell(\mathbf{z}_t; \theta),$$

where  $\ell(\cdot)$  represents the log-likelihood function. In the presence of missing values within data samples, the likelihood is only associated with observations under the assumption of MAR. The estimate then converts into:

$$\hat{\theta} = \arg \max_{\theta} \sum_t \ell(\mathbf{z}_t^o; \theta).$$

For the generative model, we derive the log-likelihood via the chain rule:

$$\hat{\theta} = \arg \max_{\theta} \sum_t \log \int_{\mathbf{u}_t} p(\mathbf{z}_t^o | \mathbf{u}_t; \theta) p(\mathbf{u}_t) d\mathbf{u}_t. \quad (13)$$

With the variational distribution, we write the log-likelihood as

$$\begin{aligned} & \sum_t \log \int_{\mathbf{u}_t} \frac{p(\mathbf{z}_t^o | \mathbf{u}_t; \theta) p(\mathbf{u}_t)}{q(\mathbf{u}_t | g(\mathbf{z}_t); \phi)} q(\mathbf{u}_t | g(\mathbf{z}_t); \phi) d\mathbf{u}_t \\ &= \sum_t \log \mathbb{E}_q \frac{p(\mathbf{z}_t^o | \mathbf{u}_t; \theta) p(\mathbf{u}_t)}{q(\mathbf{u}_t | g(\mathbf{z}_t); \phi)} \stackrel{def}{=} \mathcal{L}(\theta, \phi). \end{aligned} \quad (14)$$

As  $\mathcal{L}(\theta, \phi)$  is intractable, we derive the evidence lower bound by using the Jensen inequality, i.e.,

$$\begin{aligned} \mathcal{L}(\theta, \phi) &\geq \sum_t \mathbb{E}_q \log \frac{p(\mathbf{z}_t^o | \mathbf{u}_t; \theta) p(\mathbf{u}_t)}{q(\mathbf{u}_t | g(\mathbf{z}_t); \phi)} \\ &= \sum_t \mathbb{E}_q [\log p(\mathbf{z}_t^o | \mathbf{u}_t; \theta) + \log p(\mathbf{u}_t) - \log q(\mathbf{u}_t | g(\mathbf{z}_t); \phi)] \\ &\stackrel{def}{=} ELBO(\theta, \phi). \end{aligned} \quad (15)$$

Particularly, we derive an importance weighted lower bound by sampling from  $q(\mathbf{u}_t | g(\mathbf{z}_t); \phi)$ , i.e.

$$ELBO(\theta, \phi) = \sum_t \log \frac{1}{K} \sum_{k=1}^K \frac{p(\mathbf{z}_t^o | \mathbf{u}_{tk}; \theta) p(\mathbf{u}_{tk})}{q(\mathbf{u}_{tk} | g(\mathbf{z}_t); \phi)}, \quad (16)$$

where  $\mathbf{u}_{t1}, \mathbf{u}_{t2}, \dots, \mathbf{u}_{tK} \sim q(\mathbf{u}_t | g(\mathbf{z}_t); \phi)$ . Consequently, we estimate the parameters by minimizing  $ELBO(\theta, \phi)$ , and denote the final estimate as  $\hat{\theta}, \hat{\phi}$ .

### C. Forecasting Stage

At the forecasting stage,  $y_{t+k}$  is missing by nature, that is  $\mathbf{z}_t^m = [\mathbf{x}_t^m \top, y_{t+k}] \top$ . Then, given the observed part, we are allowed to generate samples from  $q(\mathbf{u}_t | g(\mathbf{z}_t^o); \hat{\phi}) p(\mathbf{z}_t | \mathbf{u}_t; \hat{\theta})$ . Assume we obtain  $L$  samples, which are denoted as

$$(\mathbf{u}_{t1}, \mathbf{z}_{t1}), (\mathbf{u}_{t2}, \mathbf{z}_{t2}), \dots, (\mathbf{u}_{tL}, \mathbf{z}_{tL}).$$

Each sample  $(\mathbf{u}_{ti}, \mathbf{z}_{ti})$  is associated with a weight

$$w_i = \frac{p(\mathbf{z}_t^o | \mathbf{u}_{ti}; \hat{\theta}) p(\mathbf{u}_{ti})}{q(\mathbf{u}_{ti} | g(\mathbf{z}_t); \hat{\phi})}. \quad (17)$$

Then we resample  $M$  of the  $L$  samples, with the probability defined by their corresponding weights. The procedure is sketched in Fig. 3. The samples for  $\mathbf{z}_t$  are rewritten as

$$\tilde{\mathbf{z}}_{t1}, \tilde{\mathbf{z}}_{t2}, \dots, \tilde{\mathbf{z}}_{tM},$$

which are generated by  $p(\mathbf{z}_t | \mathbf{z}_t^o; \hat{\theta}, \hat{\phi})$ . Accordingly, the probabilistic scenarios for  $Y_{t+k}$  are derived by fetching the last value of each sample, i.e.,  $\tilde{y}_{t+k,i} = \tilde{\mathbf{z}}_{ti}[h+1]$ , and denoted as

$$\tilde{y}_{t+k,1}, \tilde{y}_{t+k,2}, \dots, \tilde{y}_{t+k,M}. \quad (18)$$

## V. CASE STUDY

This section focuses on the validation of the proposed approach using an open dataset from the USA, specifically the Wind Integration National Dataset [21]. The dataset encompasses hourly data spanning 7 years, from 2007 to 2013. 80% of the data is allocated as the training set, with the remaining 20% designated for the test set. Subsequently, we commence by detailing the experimental configurations and proceed to introduce benchmark models alongside verification metrics.

### A. Experimental Setups

As the original dataset is complete, we introduce simulated missingness to validate the efficacy of the proposed approach. Data samples are systematically extracted from the dataset in a random manner, leading to the dispersion of missing values across the data. Various missing rates are explored to validate the applicability of the proposed approach. As for the features for use, we consider whether the information from neighboring wind farms is available, and design two cases, i.e., forecasting at a selected site using solely its local information, and incorporating nearby data in addition to its own. Importantly, both cases adhere to the MAR assumption, given that the occurrence of missingness is irrelevant to the specific missing values. The setups are detailed below:

- **Case 1:** Forecasting at a site by using its own information.
- **Case 2:** Forecasting at a site by using nearby information together with its own.

We explore various lead times (namely, 1, 2, and 3) in this context. The feature length is determined through the process of cross-validation. Specifically, we employ the logit-normal transformation to map normalized power values into the Real space  $\mathbb{R}$ .

### B. Benchmark Models

In general, we consider a naive model that estimate unconditional distributions via historical samples and models based on the “impute, then predict” strategy as benchmarks. Besides, we use a fully conditional specification model based on the “universal imputation” strategy as a benchmark, enabling comparison in forecast accuracy and complexity. Furthermore, we regard a model trained on complete data as a reference point. They are described in detail as below:

- **Climatology:** The method conveys forecasts by estimating the unconditional distribution using historical values. Consequently, this method yields identical distributional forecasts for all prediction targets.
- **QR-IM:** Missing values within features are imputed through a regression-based approach i.e., MissForest [22] during both the training and forecasting stages. Subsequently, a QR model is trained using the imputed data during the training stage and then applied during the forecasting stage.
- **Gaussian-IM:** Missing values within features are imputed through the MissForest during both the training and

forecasting stages. A Gaussian model is trained based on the imputed data during the training stage and leveraged during the forecasting stage.

- **DeepAR [23]:** Instead of employing an extra imputation procedure, the intermediate outcomes of the recurrent neural network model are harnessed for imputing missing values during both model estimation and forecasting stages.
- **FCS [15]:** The joint distribution of features and targets is estimated by using data with missing values. Specifically, a distinct conditional distribution for each variable is estimated by iteratively sampling and predicting. Forecasts are derived via multiple imputation by chained equations.
- **Reference:** A QR model trained based on the complete dataset.

Ideally, we assume proficient forecasting methods should exhibit performance comparable to the reference model when confronted with missing values.

### C. Verification Metrics

The quality of probabilistic forecasts is assessed via sharpness, calibration, and a proper score, i.e., the continuous ranked probability score (CRPS), which are briefed below. For further details, readers are referred to [24].

- **Sharpness:** The concentration of the predictive distributions. Specifically, it is expressed as the widths of central prediction intervals at several nominal levels.
- **Calibration:** Statistical compatibility of probabilistic forecasts and observations. Here, it is expressed as the empirical coverage of central prediction intervals at several nominal levels.
- **CRPS:** Given the lead time  $k$ , write  $\hat{F}_{t+k}$  the predicted cumulative distribution function, the CRPS is defined as

$$\text{CRPS}(\hat{F}_{t+k}, y_{t+k}) = \int_y (\hat{F}_{t+k}(y) - \mathcal{I}(y - y_{t+k}))^2 dy,$$

where  $\mathcal{I}(\cdot)$  is a step function at  $y_{t+k}$ . We report the average CRPS of all observations in the testing set.

## VI. RESULTS AND DISCUSSION

We will present the results in each case and analyze the merits and caveats of the proposed approach. In the two cases, we set  $K$  as 50.

### A. Case 1

CRPS values for 1-step forecasts generated by the proposed model are presented in Table I across varying missing rates. As expected, the CRPS value demonstrates an increase with the rising missing rate. For illustrative purposes, we display the 90% prediction intervals for a 7-day period using the proposed approach under a 20% missing rate in Fig. 4. The displayed intervals effectively encompass the observed data, as demonstrated. The comparison in performance with benchmark models will be presented in subsequent sections.

TABLE I  
CRPS VALUES OF 1-STEP FORECASTS ISSUED BY THE PROPOSED MODEL AT DIFFERENT MISSING RATES (IN TERMS OF NORMALIZED CAPACITY).

Missing rate	5%	10%	15%	20%	25%
CRPS	6.9	7.0	7.2	7.3	7.5

### B. Case 2

In this subsection, we present results achieved through the incorporation of neighboring data along with the site’s own information. Our emphasis lies in forecasting for the selected site under a 20% missing rate. Meanwhile, varying missing rates are simulated for two adjacent wind farms. Table II displays the CRPS values for 1-step forecasts, utilizing auxiliary features (AFs), across varying missing rates. The percentage represents the AFs’ respective missing rates. Despite the presence of missing values in AFs, the utilization of nearby data can still enhance the quality of forecasts. Furthermore, this observation validates the capacity of our proposed model’s features to encompass additional relevant information.

TABLE II  
CRPS VALUES OF 1-STEP FORECASTS WITH THE ASSISTANCE OF AFs AT DIFFERENT MISSING RATES (IN TERMS OF NORMALIZED CAPACITY).

	No AFs	AFs	AFs 5%	AFs 10%	AFs 20%
CRPS	7.3	6.9	7.0	7.1	7.1




Fig. 4. 90% prediction intervals of 1-step forecasts for 7 days, issued by the proposed approach.

### C. Comparison with Benchmarks

To facilitate comparison, we display the CRPS values for forecasts generated by benchmark models in Case 1 with a 20% missing rate in Table III. Broadly speaking, the DeepAR model’s performance falls short of both the “impute, then predict” and “universal impute” based models we utilized. In contrast to our proposed method, which estimates parameters by maximizing observation likelihood, DeepAR employs intermediate outcomes of a recurrent neural network to impute

missing values within the sequence and then employs a conventional parameter estimation technique. Consequently, its performance even lags behind that of simple models trained on imputed data. To elaborate, despite the convenience of being preprocessing-free, certain “end-to-end” methods such as DeepAR mainly resort to ad-hoc strategies for addressing missing values, lacking solid theoretical foundations and ultimately diminishing forecast quality.

Models based on the “impute, then predict” approach exhibit relatively inferior performance compared to those rooted in “universal imputation”. The “impute, then predict” concept is intuitive. If we can effectively restore missing data, the errors stemming from imputation in forecasting would remain minor. Nevertheless, despite employing the state-of-the-art MissForest technique in QR-IM and Gaussian-IM, errors are inevitably introduced to the forecasting model. Furthermore, it lacks a solid theoretical foundation; in contrast, the “universal imputation” approach, operating under the assumption of MAR, is capable of deriving Bayes optimal parameter estimates. The proposed model’s performance is relatively inferior to that of FCS. This discrepancy might stem from the fact that the decoder in our model is solely characterized by a Student’s t-distribution. Additionally, reliability and sharpness diagrams for 1-step forecasts across all models are illustrated in Fig. 5 and Fig. 6, respectively. The reliability diagrams of FCS, Gaussian-IM, and our proposed VAE model is closely aligned, comparable to the reference model. On the other hand, our model’s forecasts exhibit greater sharpness compared to those of FCS and Gaussian-IM.




Fig. 5. Reliability diagrams of 1-step forecasts for all models.

### D. Influence of $K$ on ELBO

As indicated by formula (16), the lower bound is approximated through  $K$  samples drawn from the latent distribution. In this subsection, we examine the impact of  $K$  by altering its value in Case 1 under a lead time of 1. We set  $K$  to the values 1, 5, 10, 20, and 50, respectively. Table IV

TABLE III  
CRPS VALUES WITH DIFFERENT LEAD TIMES AT THE MISSING RATE OF 20% (IN TERMS OF NORMALIZED CAPACITY).

Lead Time	Climatology	QR-IM	Gaussian-IM	DeepAR	FCS	Reference	Proposed
1	18.6	7.8	7.6	7.8	6.9	6.9	7.3
2	18.6	10.1	10.2	10.2	9.1	9.3	9.7
3	18.6	11.9	11.9	12.1	10.9	11.2	11.5

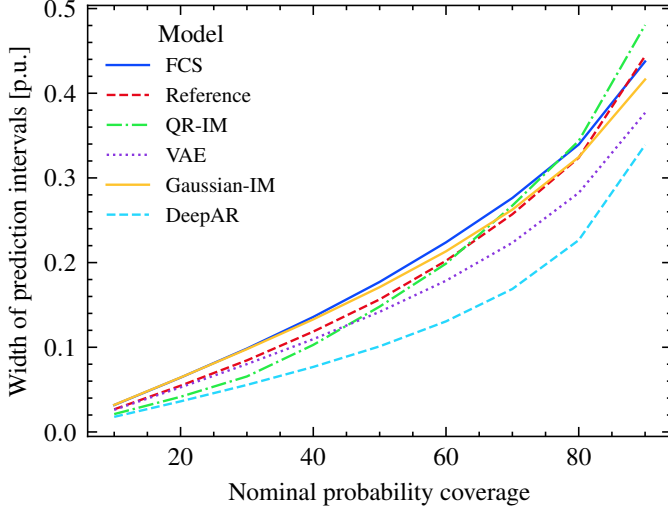


Fig. 6. Sharpness diagrams of 1-step forecasts for all models.

displays the CRPS for 1-step forecasts, varying with different values of  $K$ . Empirically, larger values of  $K$  correspond to smaller CRPS values for forecasts. As demonstrated in [25], a larger  $K$  results in a more stringent bound, approximating the likelihood more closely as  $K$  reaches a sufficient magnitude. A more stringent evidence lower bound leads to more accurate estimation of the proposed model’s parameters. On the other hand, greater values of  $K$  introduce heightened computational load during both model estimation and forecasting phases.

TABLE IV  
CRPS VALUES OF 1-STEP FORECASTS ISSUED BY THE PROPOSED MODEL BASED ON DIFFERENT  $K$  RATES (IN TERMS OF NORMALIZED CAPACITY).

$K$	1	5	10	20	50
CRPS	7.8	7.5	7.4	7.4	7.3

### E. Effectiveness of Posterior Approximation

In contrast to establishing the posterior distribution within a known distribution class, we employ a normalizing flow model to approximate it within this study. This approach provides enhanced flexibility in modeling the posterior distribution. In fact, the lower bound in formula (15) can be also written as

$$\sum_t [-KL[q(\mathbf{u}_t|g(\mathbf{z}_t); \phi)||p(\mathbf{z}_t)] + \mathbb{E}_q[\log p(\mathbf{z}_t^o|\mathbf{u}_t; \theta)]],$$

where  $KL(\cdot)$  represents the Kullback–Leibler divergence. Thus, the posterior distribution achieves optimality when  $q(\mathbf{u}_t|g(\mathbf{z}_t); \phi)$  aligns with the true posterior distribution. For

comparative purposes, we select  $q(\mathbf{u}_t|g(\mathbf{z}_t); \phi)$  to be a multivariate Gaussian distribution in Case 1. Consequently, the CRPS for 1-step forecasts amounts to 7.5, exceeding that of the proposed model.

### F. Complexity Analysis

As disclosed in [15], the training process of FCS involves iterative estimation for every conditional distribution, causing its complexity to exhibit super-linear scaling with the feature dimension. Representing the feature dimension as  $h$ , FCS constructs  $h$  regression models, each corresponding to a distinct conditional distribution. Once the regression model is set, minor changes to the training time occur with feature modifications. Accordingly, the training time for a regression model is denoted as  $T$ . The total training time for FCS in a iteration will be

$$M * h * T,$$

which is linear in the number of scenarios  $M$ . If we assume the training time for a sub neural network in the proposed model is  $T$ , then the overall training time becomes contingent on the quantity of neural networks  $L$ . As the proposed model samples  $K$  times from the latent distribution, the total training time is

$$K * L * T.$$

In essence, the training time of the proposed model remains unaffected by the number of scenarios and is less impacted by the feature dimension. To clarify, we display the training times for each model under case 1 in Table V. The training time for the proposed model is within acceptable limits. Moreover, while the feature dimension in case 2 is three times that of case 1, the training time for the proposed model exhibits a marginal increase compared to case 1. Conversely, the training of FCS in case 2 would demand over two hours. Additionally, during the forecasting stage, the FCS model outlined in [15] retains reliance on iterative imputation, with potentially impractical computation times. In contrast, the proposed approach draws upon ancestral sampling and importance resampling, enabling the generation of scenarios without resorting to iterations.

TABLE V  
TRAINING TIME FOR EACH MODEL WITHIN CASE 1 (MINUTES).

	QR-IM	Gaussian-IM	DeepAR	FCS	Proposed
Time	1	32	67	41	7

## VII. CONCLUSION

In comparison to the straightforward “impute, then predict” strategy, the “universal imputation” strategy derives Bayesian estimate for parameters by maximizing the likelihood of observations, assuming data are missing at random. Hence, when we can effectively estimate distribution parameters in the presence of missing values using advanced techniques, the forecasting quality surpasses that of models based on the “impute, then predict” approach. In this paper, we present an efficient approach based on the generative model, which exhibits significantly greater computational efficiency compared to the fully conditional specification. It is easy to use at the operational forecasting stage, and allowed to generate several probabilistic scenarios at once.

However, our current investigation is limited to analyzing the distribution within time windows of a fixed length. This focus, however, results in the omission of certain information pertaining to the sequential structure inherent in time series data. Additional endeavors might be necessary to exploit the sequential structure of time series data for updating distributions across consecutive time windows. Furthermore, it’s important to note that this approach is designed to handle cases where missing data occurs at random. However, there remains a necessity to formulate forecasting methodologies tailored to situations where data is missing not at random.

## REFERENCES

- [1] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, and H. Zareipour, “Energy forecasting: A review and outlook,” *IEEE Open Access Journal of Power and Energy*, 2020.
- [2] P. Pinson, C. Chevallier, and G. N. Kariniotakis, “Trading wind generation from short-term probabilistic forecasts of wind power,” *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1148–1156, 2007.
- [3] P. Pinson, “Very-short-term probabilistic forecasting of wind power with generalized logit–normal distributions,” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 61, no. 4, pp. 555–576, 2012.
- [4] H. Wen, J. Ma, J. Gu, L. Yuan, and Z. Jin, “Sparse variational gaussian process based day-ahead probabilistic wind power forecasting,” *IEEE Transactions on Sustainable Energy*, vol. 13, no. 2, pp. 957–970, 2022.
- [5] R. Koenker and K. F. Hallock, “Quantile regression,” *Journal of economic perspectives*, vol. 15, no. 4, pp. 143–156, 2001.
- [6] M. Landry, T. P. Erlinger, D. Patschke, and C. Varrichio, “Probabilistic gradient boosting machines for gefcom2014 wind forecasting,” *International Journal of Forecasting*, vol. 32, no. 3, pp. 1061 – 1066, 2016.
- [7] H. Wen, J. Gu, J. Ma, L. Yuan, and Z. Jin, “Probabilistic load forecasting via neural basis expansion model based prediction intervals,” *IEEE Transactions on Smart Grid*, vol. 12, no. 4, pp. 3648–3660, 2021.
- [8] P. Pinson and G. Kariniotakis, “Conditional prediction intervals of wind power generation,” *IEEE Transactions on Power Systems*, vol. 25, no. 4, pp. 1845–1856, 2010.
- [9] H. Wen, P. Pinson, J. Ma, J. Gu, and Z. Jin, “Continuous and distribution-free probabilistic wind power forecasting: A conditional normalizing flow approach,” *IEEE Transactions on Sustainable Energy*, vol. 13, no. 4, pp. 2250–2263, 2022.
- [10] R. Tawn, J. Browell, and I. Dinwoodie, “Missing data in wind farm time series: Properties and effect on forecasts,” *Electric Power Systems Research*, vol. 189, p. 106640, 2020.
- [11] R. H. Jones, “Maximum likelihood fitting of arma models to time series with missing observations,” *Technometrics*, vol. 22, no. 3, pp. 389–395, 1980.
- [12] A. C. Harvey and R. G. Pierse, “Estimating missing observations in economic time series,” *Journal of the American statistical Association*, vol. 79, no. 385, pp. 125–131, 1984.
- [13] R. Kohn and C. F. Ansley, “Estimation, prediction, and interpolation for arima models with missing data,” *Journal of the American statistical Association*, vol. 81, no. 395, pp. 751–761, 1986.
- [14] A. Stratigakos, P. Andrianesis, A. Michiorri, and G. Kariniotakis, “Towards resilient energy forecasting: A robust optimization approach,” *IEEE Transactions on Smart Grid*, 2023.
- [15] H. Wen, P. Pinson, J. Gu, and Z. Jin, “Wind energy forecasting with missing values within a fully conditional specification framework,” *International Journal of Forecasting*, 2023.
- [16] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” 2013.
- [17] P.-A. Mattei and J. Frellsen, “Miwae: Deep generative modelling and imputation of incomplete data sets,” in *International conference on machine learning*. PMLR, 2019, pp. 4413–4423.
- [18] D. Rezende and S. Mohamed, “Variational inference with normalizing flows,” in *International conference on machine learning*. PMLR, 2015, pp. 1530–1538.
- [19] R. J. Little and D. B. Rubin, *Statistical analysis with missing data*. John Wiley & Sons, 2019, vol. 793.
- [20] D. B. Rubin, “Inference and missing data,” *Biometrika*, vol. 63, no. 3, pp. 581–592, 1976.
- [21] C. Draxl, A. Clifton, B.-M. Hodge, and J. McCaa, “The wind integration national dataset (wind) toolkit,” *Applied Energy*, vol. 151, pp. 355–366, 2015.
- [22] D. J. Stekhoven and P. Bühlmann, “Missforest—non-parametric missing value imputation for mixed-type data,” *Bioinformatics*, vol. 28, no. 1, pp. 112–118, 2012.
- [23] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, “Deepar: Probabilistic forecasting with autoregressive recurrent networks,” *International Journal of Forecasting*, vol. 36, no. 3, pp. 1181–1191, 2020.
- [24] T. Gneiting and M. Katzfuss, “Probabilistic forecasting,” *Annual Review of Statistics and Its Application*, vol. 1, pp. 125–151, 2014.
- [25] Y. Burda, R. Grosse, and R. Salakhutdinov, “Importance weighted autoencoders,” in *International Conference on Learning Representations*, 2016.