

Generative Probabilistic Time Series Forecasting and Applications in Grid Operations

Xinyi Wang

Electrical and Computer Engineering
Cornell University
Ithaca, USA
xw555@cornell.edu

Lang Tong

Electrical and Computer Engineering
Cornell University
Ithaca, USA
lt35@cornell.edu

Qing Zhao

Electrical and Computer Engineering
Cornell University
Ithaca, USA
qz16@cornell.edu

Abstract—Generative probabilistic forecasting produces future time series samples according to the conditional probability distribution given past time series observations. Such techniques are essential in risk-based decision-making and planning under uncertainty with broad applications in grid operations, including electricity price forecasting, risk-based economic dispatch, and stochastic optimizations. Inspired by Wiener and Kallianpur’s innovation representation, we propose a weak innovation autoencoder architecture and a learning algorithm to extract independent and identically distributed innovation sequences from nonparametric stationary time series. We show that the weak innovation sequence is Bayesian sufficient, which makes the proposed weak innovation autoencoder a canonical architecture for generative probabilistic forecasting. The proposed technique is applied to forecasting highly volatile real-time electricity prices, demonstrating superior performance across multiple forecasting measures over leading probabilistic and point forecasting techniques.

Index Terms—Probabilistic time series forecasting, innovation representation, autoencoder, generative adversarial networks

I. INTRODUCTION

We develop a *generative probabilistic forecasting (GPF) technique* for unknown and nonparametric time series models. Whereas standard probabilistic forecasting aims to estimate the conditional probability distribution of the time series at a future time, GPF obtains a generative model capable of producing arbitrarily many Monte Carlo samples of future time series realizations according to the conditional probability distribution of the time series given past observations. As a nonparametric probabilistic forecasting technique, GPF is essential for decision-making under uncertainty where data-driven risk-sensitive optimization requires conditional samples of future randomness. The Monte Carlo samples generated from GPF can be used to produce any form of point forecast.

Nonparametric probabilistic time series forecasting is challenging because the joint probability distributions that characterize temporal dependencies are infinite-dimensional and unknown. Most approaches rely on parametric and semiparametric models, restricting the infinite-dimensional inference to a finite-dimensional parameter space. Classic examples include the forecasting based on autoregressive moving average, GARCH, and Gaussian process models [1]–[3].

This work was supported in part by the National Science Foundation under Award 2218110.

This work pursues a new path to nonparametric probabilistic forecasting, following the generative AI principle and inspired by the classic notion of *innovation representation* pioneered by Norbert Wiener and Gopinath Kallianpur in 1958 [4]. In the parlance of modern machine learning, the Wiener-Kallianpur innovation representation is an autoencoder, where the *causal encoder* transforms a stationary time series (X_t) to an *innovation process* (V_t) defined by the independent and identical (i.i.d.) uniformly distributed sequence, followed by a *causal decoder* that reproduces almost surely the original time series. The most striking feature of the Wiener-Kallianpur innovation representation is that V_t at time t is statistically independent of the past $\mathcal{X}_t := \{X_{t-1}, X_{t-2}, \dots\}$, implying that V_t contains only new information at time t . The autoencoder’s capability to reproduce (X_t) makes the latent process sufficient statistics in any decision-making. The innovation being i.i.d. uniform means that the autoencoder captures the complete model temporal dependencies of the time series. This particular feature of the latent process plays a critical role in the proposed GPF methodology.

The main challenge of applying Wiener-Kallianpur innovation representation for inference and decision-making is twofold. First, obtaining a causal encoder to extract the innovation process requires knowing the marginal and joint distributions of the time series, which is rarely possible without imposing some parametric structure. Furthermore, even when the probability distribution is known, there is no known computationally tractable way to construct the causal encoder to obtain an innovation process. Second, the Wiener-Kallianpur innovation representation may not exist for a broad class of random processes, including some of the important cases of finite-state Markov chains [5]. These conceptual and computational barriers prevent employing innovation representation to a broad class of inference and decision problems except for Gaussian and additive Gaussian models [6].

A. Summary of Contributions

This paper makes methodological and practical advances in nonparametric probabilistic time series forecasting in three aspects. First, we propose a Weak Innovation AutoEncoder (WIAE) and a deep learning algorithm based on Rosenblatt’s weak innovation representation [5], a relaxation of Wiener-

Kallianpur's strong innovation representation, significantly extending the class of time series for which the canonical uniform i.i.d. innovation representation exists. We show further that the latent process of WIAE forms Bayesian sufficient statistics for generative probabilistic forecasting, resulting in no loss of optimality when the weak innovation representation is used.

Second, we propose Generative Probabilistic Forecasting with Weak Innovation (GPF-WI), a WIAE-based deep learning technique that produces future time series samples according to the conditional probability distribution on past realizations. Through WIAE, GPF-WI transforms the problem of generating Monte Carlo samples of the future time series with unknown probability distributions to using past innovation observations and Monte Carlo samples from the i.i.d. uniform distribution. See Sec. III-A.

Third, we apply GPF-WI to probabilistic forecasting of highly volatile wholesale electricity prices using real datasets from several major independent system operators. The frequent price spikes due to increasing renewable integration and significant weather events make conventional techniques ineffective. The empirical results demonstrate marked improvement over existing solutions, scoring near the top across multiple datasets and under different performance metrics. See Sec. V.

B. A Contextual Review of Literature

Nonparametric probabilistic time series forecasting has a long history. See [3] for a review up to the late 1990's. Machine learning techniques gained prominence in the last decade, which were reviewed and compared by [7]. Here we restrict our discussion to those machine learning techniques sharing architectural or methodological similarities with our approach. In particular, we focus on autoencoder-based forecasting techniques and the generative methods that rely on Monte-Carlo sampling to produce realizations of forecasts.

Most autoencoder-based GPF techniques rely on state-of-the-art autoencoder architectures for latent representation learning, followed by producing forecasts by Monte Carlo sampling of the latent representation. These methods generally adopt likelihood-based generative models such as Variational Autoencoder (VAE) [8], [9], normalizing flows [10] and denoising diffusion probabilistic models [11]. For instance, the authors in [8] developed a technique that models the conditional distribution of future time series as a Gaussian latent process of an autoencoder trained with historical data. A normalizing flow model conditioned on Recurrent Neural Network (RNN) or transformer model is adopted to learn the conditional probability distribution [10].

Both VAE and diffusion learning rely on statistically independent training samples with log-likelihood function. When deriving the variational lower bound for log-likelihood and implementing it for training, those methods assume that the samples in the training set are independent, which simplifies the minimization of an upper bound of negative log-likelihood. For time series segments, the samples aren't independent. Hence, minimizing the joint log-likelihood of multiple samples is

intractable, making these methods not suitable for time series. In contrast, when learning weak innovation representation through WIAE, as considered in our approach, time series training samples are allowed to have unknown temporal dependencies.

Novel learning architectures capable of modeling complex temporal dependencies have been proposed in recent years. In [12]–[14], the authors adopt RNN to develop a variety of parametric probabilistic forecasting techniques. The authors in [15], [16] adopted a dilated convolutional network to learn the conditional distribution.

The success of transformer-based deep learning in natural language processing inspired its applications to time series forecasting [17]–[19], some demonstrating promising performance. These point forecasting techniques have not been tested on the more challenging time series with high volatility, such as electricity prices. Although these techniques do not produce probabilistic forecasts, their strong performance in point forecasting makes them worth comparing with.

Clarification of notations used is in order. Random variables are highlighted in capital letters and their realizations in lower cases. When needed for clarity, vectors are in boldface. We use (X_t) for a time series indexed by t , extending from $-\infty$ to ∞ . For a partial time series realization up to time t is denoted by $\mathbf{x}_t := (x_t, x_{t-1}, \dots)$ and $\mathbf{x}_{t_o:t} := (x_t, x_{t-1}, \dots, x_{t_o})$.

II. INNOVATION REPRESENTATION

Innovation representation can be viewed as a causal autoencoder transformation (G, H) , as illustrated by Fig. 1. A key characteristics separating innovation autoencoder from other autoencoder architectures is that the latent process (V_t) is an i.i.d. uniform sequence extracted through the causal encoder G , and it can be used to produce an estimate (\hat{X}_t) of (X_t) by a causal decoder H . Specifically,

$$V_t = G(X_t, X_{t-1}, \dots), \quad (1.1)$$

$$\hat{X}_t = H(V_t, V_{t-1}, \dots), \quad (1.2)$$

$$V_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{U}[0, 1]. \quad (1.3)$$

The temporal independence of (V_t) implies that V_t is independent of the past $\mathbf{X}_t := (X_{t-1}, X_{t-2}, \dots)$. Hence, V_t represents the new information not contained in \mathbf{X}_t .

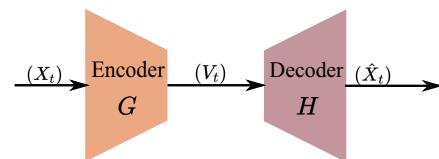


Fig. 1: An autoencoder interpretation of innovation representation.

The criterion for reconstruction distinguishes different types of innovation representations. The Wiener-Kallianpur innovation representation originally proposed by Wiener [4] requires almost sure equality between (X_t) and (\hat{X}_t) , hence referred to as *strong innovations representation*. That is, the transformation

pair (G, H) is causally invertible ensuring that $(X_t) \stackrel{a.s.}{=} (X_t)$. The causal invertibility makes the Wiener-Kallianpur innovation sequence (V_t) a sufficient statistic. Therefore, the decision-making problem based on strong innovations is lossless. Unfortunately, the Wiener-Kallianpur innovation can only be analytically computed for stationary Gaussian time series [20] and additive Gaussian models [6]. For those time series, the prediction error of the nonlinear minimum mean squared error predictor is its innovation sequence. Methods of extracting innovations from nonparametric time series were mostly unknown until recently when we obtained the first machine learning approach that learns the autoencoder of the Wiener-Kallianpur innovation representation [21].

Rosenblatt was the first to state that the Wiener-Kallianpur innovation does not exist for a broad class of random processes [5]. He thus suggested a weaker version of the Wiener-Kallianpur innovation representation by replacing the almost sure equality with distribution-wise equality:

$$(\hat{X}_t) \stackrel{d}{=} (X_t). \quad (2)$$

Herein, we refer to the innovation representation with matching input-output in distribution as weak innovation representation and (G, H) as the weak innovation autoencoder.

Although the existence of a WIAE remains an open problem, Rosenblatt has shown that his relaxation broadens the class of random processes for which weak innovation exists. With (V_t) not necessarily a sufficient statistic, decision-making based on (V_t) is lossy in general. An open question is whether forecasting decisions based on a weak innovation representation perform optimally. To this end, we provide a definitive answer for the probabilistic forecasting problem in Sec. III. Currently, there are no known techniques to extract weak innovation sequences from time series. We propose the first such technique based generative adversary network (GAN) learning in Sec. IV.

III. GENERATIVE PROBABILISTIC FORECASTING

Consider a stationary time series (X_t) . The probabilistic forecasting of X_{t+T} given current and past realizations $\mathbf{X}_t = \mathbf{x}_t$ is to estimate the conditional distribution $F_{t+T|t}(x|\mathbf{x}_t)$:

$$F_{t+T|t}(x|\mathbf{x}_t) := \Pr[x_{t+T} \leq x | \mathbf{X}_t = \mathbf{x}_t]. \quad (3)$$

In contrast, GPF produces Monte Carlo samples of X_{t+T} with the conditional distribution $F_{t|t+T}$. This is nontrivial because of the unknown and complex temporal dependencies of (X_t) .

A. GPF with Weak Innovation

This section introduces GPF-WI for forecasting. Generating samples directly from $F_{t+T|T}$ is difficult. As shown in Fig. 2, GPF-WI generates conditional samples $(V_{t+T}, \dots, V_{t+1})$ from the i.i.d. uniform distribution, and uses WIAE decoder to produce future samples. The key advantage of using innovations rather than the original time series is that future innovations are statistically independent of the past observations $\mathbf{V}_t = \mathbf{v}_t$.

As illustrated by Fig. 2, the weak innovations up to time t is the output of the weak innovation encoder G . Although future innovation $\{V_{t+1}, \dots, V_{t+T}\}$ are not yet realized, the fact that

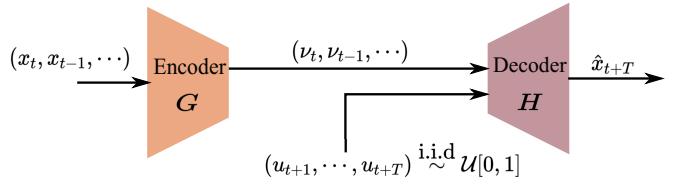


Fig. 2: Probabilistic Forecasting via Weak Innovations.

they are i.i.d. uniform allows us to sample them from the uniform marginal. These Monte Carlo samples $(u_{t+1}, \dots, u_{t+T})$, along with the encoder outputs (v_t, v_{t-1}, \dots) , form the input of the decoder H that produces an estimate of x_{t+T} . Different sets of Monte Carlo samples produce different estimates of x_{t+T} , generated from its conditional probability distribution.

While GPF-WI is easy to implement once the WIAE has been learned, one should question whether replacing \mathbf{X}_t by \mathbf{V}_t is justified, since weak innovations may not be sufficient statistics. We address this question next.

B. Weak Innovation are Bayesian Sufficient

The objective of probabilistic forecasting is to estimate the conditional probability distribution, which is deterministic once past observation is given. In contrast, GPF aims to obtain a generative model that produces estimates of a set of realizations with conditional $F_{t+T|t}$. In this aspect, GPF is a Bayesian estimation problem, where the notion of sufficient statistics takes a different form. In particular, $T(X)$ is Bayesian sufficient for estimating a realization of random variable Y if the posterior distribution given X is the same given $T(X)$ [22].

The theorem below establishes that the weak innovation sequence $\mathbf{v}_t = G(\mathbf{x}_t)$, under a mild assumption on WIAE, is Bayesian sufficient for estimating realizations of X_{t+T} . For simplicity, we define the following notation for the conditional distribution of $X_{t+T}|\mathbf{V}_t$:

$$F_{t+T|t}^{(\nu)}(x|\mathbf{v}_t) := \Pr[V_{t+T} \leq x | \mathbf{V}_t = \mathbf{v}_t].$$

Theorem 1 (Bayesian Sufficiency): Let (X_t) be a stationary time series for which the weak innovation exists. Let (V_t) be the weak innovation sequence of (X_t) , and assume that the causal decoder H is injective. Then, for almost all \mathbf{x}_t and x (with respect to Lebesgue measure),

$$F_{t+T|t}(x|\mathbf{x}_t) = F_{t+T|t}^{(\nu)}(x|\mathbf{v}_t),$$

where $F_{t+T|t}(x|\mathbf{x}_t)$ is defined in (3), $\mathbf{v}_t = G(\mathbf{x}_t)$, and

$$F_{t+T|t}^{\nu}(x|\mathbf{v}_t) := \Pr(X_{t+T} \leq x | \mathbf{v}_t).$$

Proof: See the appendix of [23].

IV. LEARNING INNOVATION REPRESENTATION

This section focuses on the learning of WIAE that generates the weak innovation sequence as its latent process. Shown in Fig. 3 is the schematic, where both the encoder G_θ and decoder H_η are convolutional neural networks parameterized by neural network coefficients θ and η , respectively.

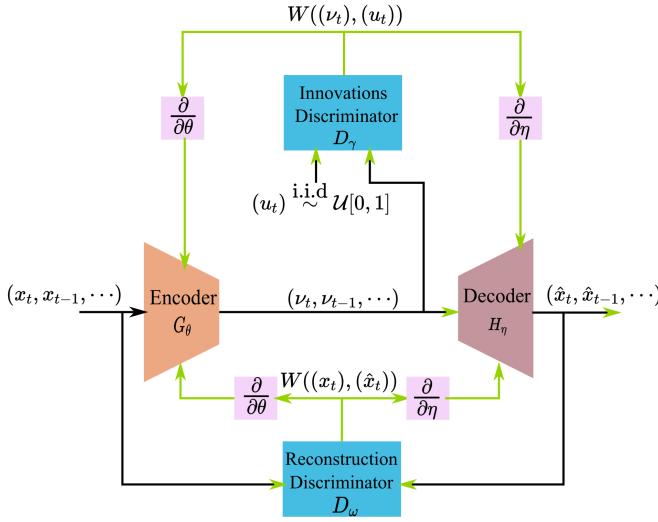


Fig. 3: Training WIAE with discriminators.

The WIAE is trained to enforce the weak innovation conditions for its latent process using gradients from two discriminators: the *innovation discriminator* D_γ to force (ν_t) to be i.i.d. uniform and the *reconstruction discriminator* D_ω to match the input-output distributions. We present the objective and an algorithm for WIAE learning in the following paragraphs.

We now define the objective and an algorithm for WIAE learning to generate weak innovation sequence (V_t) as its latent process. Both discriminators D_γ and D_ω in Fig. 3 compare two random processes to produce gradient updates: D_γ provides gradient signals for encoder G_θ to enforce the i.i.d. uniform condition on its output (ν_t) , and D_ω provides gradient signals for the encoder-decoder pair (G_θ, H_η) to produce the input-output matching distributions for WIAE. It is then natural to define the training objective function by a weighted sum of training losses observed by the two discriminators.

We use Wasserstein distance $W(f_1, f_2)$ to measure the distance between two distributions f_1 and f_2 . In particular, we adopt the learning paradigm of Wasserstein GAN (WGAN) [24] to train deep neural network discriminators D_γ and D_ω . By the Kantorovich-Rubinstein duality, Wasserstein distance between two random variables $X \sim f_1$ and $Y \sim f_2$ can be written as the maximum expected difference between $D(X)$ and $D(Y)$ under a 1-Lipschitz function D :

$$W(f_1, f_2) = \max_{D: 1\text{-Lipschitz}} (\mathbb{E}_{X \sim f_1}[D(X)] - \mathbb{E}_{Y \sim f_2}[D(Y)]). \quad (4)$$

WGAN employs a ω -parameterized neural network for the 1-Lipschitz function D_ω , with ω optimized according to (4).

Specializing to WIAE learning in Fig. 3, the two discriminators in WIAE are deep neural networks D_γ and D_ω minimized jointly under the following risk measure:

$$\begin{aligned} L((X_t), \theta, \eta) := \max_{\gamma, \omega} & (\mathbb{E}[D_\gamma(U_t)] - \mathbb{E}[D_\gamma(V_t)]) \\ & + \lambda(\mathbb{E}[D_\omega(X_t)] - \mathbb{E}[D_\omega(\hat{X}_t)]). \end{aligned} \quad (5)$$

λ is a real number that scales the two Wasserstein distances. The two parts of the loss function regularize the innovation according to (1.3) and (2). Minimizing Eq. (5) with respect to θ and η is thus equivalent to enforcing (V_t) being i.i.d uniform, and (\hat{X}_t) having the same distribution as (X_t) .

V. NUMERICAL RESULTS

This section presents an application of GPF-WI to real-world time series forecasting problems. We conducted numerical experiments on a benchmark M4 dataset and two real-time wholesale electricity price datasets for their high volatility. The detailed information about electricity price datasets can be found in Sec. V-A. Among the M4 datasets, we chose the electricity consumption one (hereafter referred to as Electricity) consisting of hourly load measurements from 370 clients.

We compared our methods with 6 other state-of-the-art time series forecasting techniques: DeepAR [12], Nonparametric Time Series Forecaster (NPTS) [25], Pyraformer [19], TLAE [8], Wavenet [15], and SNARX [1]. Both probabilistic forecasting methods and point-forecasting methods were included. DeepAR is an auto-regressive RNN time series model that estimates the parameters of parametric distributions. NPTS is a probabilistic forecasting technique that resembles naive forecasters. It randomly samples a past time index following a categorical probability distribution over time indices. TLAE is a nonparametric probabilistic forecasting technique that utilizes autoencoder architecture to learn the underlying Gaussian latent process, and uses it as the estimator. Wavenet is a parametric probabilistic forecasting technique that is based on dilated causal convolutions. Pyraformer is a point estimation technique that adopts multi-resolution attention modules, and is trained by minimizing mean squared error. SNARX is a semiparametric AR model that utilizes kernel density function to estimate the distribution of noise, which has superior empirical performance on electricity pricing datasets according to [2].

To comprehensively compare the strengths and weaknesses of different forecasting techniques, we adopted multiple evaluation metrics: normalized mean square error (NMSE), normalized mean absolute error (NMAE), symmetric mean absolute percentage error (sMAPE), and continuous ranked probability score (CRPS). For probabilistic forecasting methods, we used sample mean as the point estimator when calculating square errors, and sample median for absolute errors. In total, 1000 trajectories were sampled for probabilistic methods. The commonly used mean absolute percentage error was not adopted, since in electricity price datasets the absolute value of the prices can be very close to 0, which nullifies the effectiveness of the metrics. Since electricity price datasets exhibit high variability, we excluded the outliers, which are defined as samples that are three standard deviations away from the sample mean, when calculating all metrics, for all methods.

A. Electricity Prices and Datasets

Wholesale electricity prices in the U.S. are computed from an optimization over the power generation and demand levels

TABLE I: Numerical Results. The numbers in the parentheses after each dataset name indicates the prediction step. The number in parentheses after each value of the metrics indicates the ranking of the algorithm.

Metrics	Method	ISONE (75 min)	ISONE (2 hour)	NYISO (75 min)	NYISO (2 hour)	Electricity (5 hour)	Electricity (10 hour)
NMSE	GPF-WI	0.0857 (1)	0.0868 (1)	0.0876 (1)	0.0887 (1)	0.4682 (3)	0.5165 (3)
	DeepAR	0.1064 (2)	0.2852 (6)	0.0952 (2)	0.1632 (3)	0.5290 (4)	0.6272 (4)
	NPTS	0.1216 (3)	0.1250 (2)	0.1840 (5)	0.2031 (5)	1.3862 (6)	1.3867 (6)
	Pyraformer	0.1247 (6)	0.1779 (4)	0.4031 (6)	0.4229 (6)	0.1422 (1)	0.1735 (1)
	TLAE	0.1219 (4)	0.1837 (5)	0.1116 (3)	0.1492 (2)	0.1881 (2)	0.2304 (2)
	Wavenet	0.1245 (5)	0.1643 (3)	0.1128 (4)	0.1809 (4)	1.3367 (5)	1.3575 (5)
NMAE	SNARX	0.8632 (7)	0.9658 (7)	0.9586 (7)	1.8063 (7)	1.7193 (7)	1.8372 (7)
	GPF-WI	0.2327 (1)	0.2330 (1)	0.2112 (3)	0.2104 (1)	0.2296 (1)	0.3487 (1)
	DeepAR	0.2969 (5)	0.4458 (6)	0.1799 (2)	0.2451 (4)	0.3360 (3)	0.3559 (2)
	NPTS	0.2896 (4)	0.2958 (2)	0.2260 (4)	0.2138 (2)	0.4625 (5)	0.4627 (5)
	Pyraformer	0.2570 (3)	0.3518 (4)	0.4892 (6)	0.5234 (6)	0.3396 (4)	0.3794 (4)
	TLAE	0.2565 (2)	0.3536 (5)	0.2258 (5)	0.2590 (5)	0.3041 (2)	0.3597 (3)
sMAPE	Wavenet	0.3154 (6)	0.3310 (3)	0.1770 (1)	0.2330 (3)	0.6215 (6)	0.6216 (6)
	SNARX	0.8317 (7)	0.9611 (7)	0.9299 (7)	0.9967 (7)	0.7243 (7)	0.7624 (7)
	GPF-WI	0.2421 (2)	0.2419 (1)	0.2106 (4)	0.2091 (1)	0.3266 (1)	0.4034 (1)
	DeepAR	0.2372 (1)	0.7076 (5)	0.1709 (1)	0.2236 (2)	0.4917 (3)	0.5004 (3)
	NPTS	0.8601 (7)	0.8854 (6)	0.7395 (6)	0.7506 (6)	0.6051 (4)	0.6053 (4)
	Pyraformer	0.2785 (4)	0.2896 (2)	0.3672 (5)	0.4095 (5)	0.6385 (5)	0.6546 (5)
CRPS	TLAE	0.2440 (3)	0.3137 (3)	0.1974 (2)	0.2549 (4)	0.3974 (2)	0.4791 (2)
	Wavenet	0.3057 (5)	0.3375 (4)	0.2090 (3)	0.2475 (3)	1.9877 (7)	1.9858 (7)
	SNARX	0.5035 (6)	1.3589 (7)	1.2661 (7)	1.8736 (7)	1.0783 (6)	1.1125 (6)

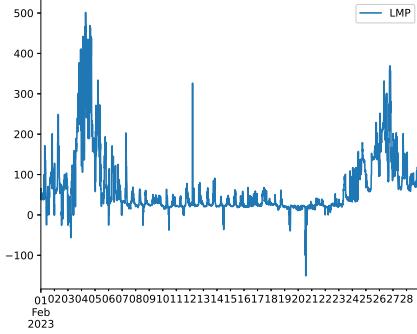


Fig. 4: Trajectory of ISONE real-time electricity price.

based on offers from producers and bids from power generators, subject to generation limits and power flow constraints. The price of electricity at a particular location is the Lagrange multiplier associated with the power balance constraint at that location, which is very sensitive to the line congestion pattern of the power grid. When a line changes from uncongested to congested due to even a tiny change in power flow, the price of electricity incurs a sudden change. Such discontinuities make electricity prices highly volatile. See Fig. 4 for the price trajectory of a typical day. The wholesale market price model can be found in [26], [27], where the authors also considered a very different price forecasting problem where the system operator has access to the network operating conditions and offers and bids from market participants. In contrast, the price forecasting problem studied here is time series forecasting without assuming how the time series is generated. We chose two publicly available real-time electricity price datasets from two independent system operators (ISO) to demonstrate the ability of forecasting time series with higher sampling frequency

and larger variability. We named the datasets with the names of the ISOs: New York ISO (NYISO) and ISO New England (ISONE). NYISO and ISONE consist of 5-minute real time electricity prices. We took one month (February of 2023) of electricity prices for NYISO and ISONE.

We conducted 75-min and 2-hour ahead prediction for 5-min real-time electricity price (ISONE & NYISO), and 5-hour and 10-hour ahead prediction for hourly sampled electricity consumption. The WIAE is trained with Adam optimizer.

B. Results

The simulation result is shown in Table. I. It can be seen that GPF-WI had better performance for most of the metrics. Pyraformer, as a point estimation technique optimized by MSE, performed worse under metrics using absolute errors. We also observed that the auto-regressive probabilistic forecasting methods (DeepAR, NPTS, Wavenet, SNARX) have the tendency to be affected by past observation, which leads to better performance when the time series is smooth with few fluctuations. Due to the fact that electricity prices being a Lagrangian multiplier that doesn't exhibit continuity as other measurements taken from a physical system, the auto-regressive methods predicted shifted peaks, which contributed to large errors under more volatile datasets. On the other hand, the generative probabilistic forecasting methods (GPF-WI, TLAE), where the prediction is conditioned on latent processes, suffered less from the volatility problem. For TLAE, the latent process being a correlated Gaussian process posed great difficulty to the Monte Carlo sampling of the latent process, essential to probabilistic forecasting. This was exhibited by the instability of probabilistic forecasts, which varies significantly with between trajectories.

VI. CONCLUSION

This paper presents a novel approach to generative probabilistic forecasting (GPF) of nonparametric time series with unknown probability structures, reviving the classic idea of innovation representation of Wiener-Kallianpur and Rosenblatt with modern generative machine learning techniques. Our main theoretical contribution lies in establishing that the weak innovation process of a time series forms Bayesian sufficient statistics for probabilistic forecasting. These results imply that WIAE is a canonical architecture for stochastic decision-making such as probabilistic forecasting. We have also made a methodological advance in developing a generative forecasting algorithm that produces Monte Carlo samples of future time series realizations. The application of the proposed technique to real world datasets demonstrates that innovation-based forecasting offers superior performance across multiple performance measures over leading benchmarks.

The vanilla implementation of GPF presented here can be extended in multiple directions; chief among them is the multivariate GPF. While the underlying principle of innovation-based forecasting applies to multivariate time series, the condition under which weak innovation representation exists may become more tenuous and the training of WIAE more challenging. More sophisticated neural network implementations such as RNN and LSTM may also improve the forecasting performance.

REFERENCES

- [1] R. Weron and A. Misiorek, "Forecasting spot electricity prices: A comparison of parametric and semiparametric time series models," *International Journal of Forecasting*, vol. 24, no. 4, pp. 744–763, 2008. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169207008000952>
- [2] R. Weron, "Electricity price forecasting: A review of the state-of-the-art with a look into the future," *International Journal of Forecasting*, vol. 30, no. 4, pp. 1030–1081, Oct. 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169207014001083>
- [3] W. Härdle, H. Lütkepohl, and R. Chen, "A review of nonparametric time series analysis," *International Statistical Review / Revue Internationale de Statistique*, vol. 65, no. 1, pp. 49–72, 1997, publisher: [Wiley, International Statistical Institute (ISI)]. [Online]. Available: <https://www.jstor.org/stable/1403432>
- [4] N. Wiener, *Nonlinear Problems in Random Theory*. Cambridge, MA: Technology Press of Massachusetts Institute of Technology, 1958.
- [5] M. Rosenblatt, "Stationary processes as shifts of functions of independent random variables," *Journal of Mathematics and Mechanics*, vol. 8, no. 5, pp. 665–681, 1959.
- [6] T. Kailath, "An innovations approach to least-squares estimation—Part I: Linear filtering in additive white noise," *IEEE Transactions on Automatic Control*, vol. 13, no. 6, pp. 646–655, 1968.
- [7] A. Gautam and V. Singh, "Parametric versus non-parametric time series forecasting methods: A review," *Journal of Engineering Science and Technology Review*, vol. 13, pp. 165–171, 2020.
- [8] N. Nguyen and B. Quanz, "Temporal latent auto-encoder: a method for probabilistic multivariate time series forecasting," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 10, pp. 9117–9125, May 2021, number: 10. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/17101>
- [9] L. Li, J. Zhang, J. Yan, Y. Jin, Y. Zhang, Y. Duan, and G. Tian, "Synergetic learning of heterogeneous temporal sequences for multi-horizon probabilistic forecasting," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 10, pp. 8420–8428, May 2021, number: 10. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/17023>
- [10] K. Rasul, A.-S. Sheikh, I. Schuster, U. M. Bergmann, and R. Vollgraf, "Multivariate probabilistic time series forecasting via conditioned normalizing flows," Feb. 2022. [Online]. Available: <https://openreview.net/forum?id=WiGQBFuVRv>
- [11] Y. Li, X. Lu, Y. Wang, and D. Dou, "Generative time series forecasting with diffusion, denoise, and disentanglement," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 23009–23022. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/91a85f3fb8f570e6be52b333b5ab017a-Paper-Conference.pdf
- [12] D. Salinas, V. Flunkert, and J. Gasthaus, "DeepAR: probabilistic forecasting with autoregressive recurrent networks," Feb. 2019, arXiv:1704.04110 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1704.04110>
- [13] Y. Wang, A. Smola, D. C. Maddix, J. Gasthaus, D. Foster, and T. Januschowski, "Deep factors for forecasting," May 2019, arXiv:1905.12417 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1905.12417>
- [14] H. Du, S. Du, and W. Li, "Probabilistic time series forecasting with deep non-linear state space models," *CAAI Transactions on Intelligence Technology*, vol. n/a, no. n/a, 2022, _eprint: <https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/cit2.12085>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1049/cit2.12085>
- [15] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "WaveNet: a generative model for raw audio," Sep. 2016, arXiv:1609.03499 [cs]. [Online]. Available: <http://arxiv.org/abs/1609.03499>
- [16] A. Borovykh, S. Bohte, and C. W. Oosterlee, "Conditional time series forecasting with convolutional neural networks," Sep. 2018, arXiv:1703.04691 [stat]. [Online]. Available: <http://arxiv.org/abs/1703.04691>
- [17] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: beyond efficient transformer for long sequence time-series forecasting," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, pp. 11106–11115, May 2021, number: 12. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/17325>
- [18] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin, "FEDformer: frequency enhanced decomposed transformer for long-term series forecasting," Jun. 2022, arXiv:2201.12740 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/2201.12740>
- [19] S. Liu, H. Yu, C. Liao, J. Li, W. Lin, A. X. Liu, and S. Dustdar, "Pyraformer: low-complexity pyramidal attention for long-range time series modeling and forecasting," Mar. 2022. [Online]. Available: <https://openreview.net/forum?id=0EXmFzUn5I>
- [20] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Trans. ASME J. of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [21] X. Wang and L. Tong, "Innovations autoencoder and its application in one-class anomalous sequence detection," *Journal of Machine Learning Research*, vol. 23, no. 49, pp. 1–27, 2022. [Online]. Available: <http://jmlr.org/papers/v23/21-0735.html>
- [22] P. J. Bickel and K. A. Doksum, *Mathematical statistics: basic ideas and selected topics*. (2nd ed.). Upper Saddle River, N.J.: Pearson Prentice Hall, 2007, vol. 1.
- [23] X. Wang, M. Lee, Q. Zhao, and L. Tong, "Non-parametric probabilistic time series forecasting via innovations representation," *arXiv preprint arXiv:2306.03782*, 2023.
- [24] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," Jan. 2017, arXiv:1701.07875.
- [25] A. Alexandrov, K. Benidis, M. Bohlke-Schneider, V. Flunkert, J. Gasthaus, T. Januschowski, D. C. Maddix, S. Rangapuram, D. Salinas, J. Schulz, L. Stella, A. C. Türkmen, and Y. Wang, "GluonTS: probabilistic time series models in Python," Jun. 2019, arXiv:1906.05264 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1906.05264>
- [26] Y. Ji, R. J. Thomas, and L. Tong, "Probabilistic forecasting of real-time lmp and network congestion," *IEEE Transactions on Power Systems*, vol. 32, no. 2, pp. 831–841, 2017.
- [27] Y. Ji, L. Tong, and W. Deng, "Probabilistic forecasting of power system and market operations," in *Advanced data analytics for power systems*. Cambridge : Cambridge University Press, 2021.