
Better Batch for Deep Probabilistic Time Series Forecasting

Vincent Zhihao Zheng
McGill University

Seongjin Choi
University of Minnesota

Lijun Sun
McGill University

Abstract

Deep probabilistic time series forecasting has gained attention for its ability to provide nonlinear approximation and valuable uncertainty quantification for decision-making. However, existing models often oversimplify the problem by assuming a time-independent error process and overlooking serial correlation. To overcome this limitation, we propose an innovative training method that incorporates error autocorrelation to enhance probabilistic forecasting accuracy. Our method constructs a mini-batch as a collection of D consecutive time series segments for model training. It explicitly learns a time-varying covariance matrix over each mini-batch, encoding error correlation among adjacent time steps. The learned covariance matrix can be used to improve prediction accuracy and enhance uncertainty quantification. We evaluate our method on two different neural forecasting models and multiple public datasets. Experimental results confirm the effectiveness of the proposed approach in improving the performance of both models across a range of datasets, resulting in notable improvements in predictive accuracy.

1 INTRODUCTION

Time series forecasting stands as one of the primary tasks in the field of deep learning (DL) due to its broad range of applications (Benidis et al., 2022). Essentially, the problem of time series forecasting can be classified into deterministic forecasting and probabilistic forecasting. Deterministic forecasting provides point estimates for future time series values, while

probabilistic forecasting goes a step further by providing a distribution that quantifies the uncertainty associated with the predictions. As additional information on uncertainty assists users in making more informed decisions, probabilistic forecasting has become increasingly attractive and extensive efforts have been made to enhance uncertainty quantification. In time series analysis, errors can exhibit correlation for various reasons, such as the omission of essential covariates or model inadequacy. Autocorrelation (also known as serial correlation) and contemporaneous correlation are two common types of correlation in time series forecasting. Autocorrelation captures the temporal correlation present in errors, whereas contemporaneous correlation refers to the correlation among different time series at the same time.

This paper primarily investigates the issue of autocorrelation in errors. Modeling error autocorrelation is an important field in the statistical analysis of time series. A widely adopted method for representing autocorrelated errors is assuming the error series follows an autoregressive integrated moving average (ARIMA) process (Hyndman and Athanasopoulos, 2018). Similar issues may arise in learning nonlinear DL-based forecasting models. Previous studies have attempted to model error autocorrelation in deterministic DL models using the concept of dynamic regression (Sun et al., 2021; Zheng et al., 2023), assuming that the errors follow a first-order autoregressive process. However, since both neural networks and correlated errors can explain the data, these models may face challenges in balancing the two sources in the absence of an overall covariance structure. More importantly, these methods are not readily applicable to probabilistic models, where the model output typically consists of parameters of the predictive distribution rather than the estimated time series values.

In this paper, we propose a novel batch structure that allows us to explicitly model error autocorrelation. Each batch comprises multiple mini-batches, with each mini-batch grouping a fixed number of consecutive training instances. Our main idea draws inspiration from the generalized least squares (GLS) method used

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

in linear regression models with dependent errors. We extend the Gaussian likelihood of a univariate model to a multivariate Gaussian likelihood by incorporating a time-varying covariance matrix that encodes error autocorrelation within a mini-batch. The covariance matrix is decomposed into two components, a scale vector and a correlation matrix, that are both time-varying. In particular, we parameterize the correlation matrix using a weighted sum of several base kernel matrices, and the weights are dynamically generated from the output of the base probabilistic forecasting model. This enables us to improve the accuracy of estimated distribution parameters during prediction by using the learned dynamic covariance matrix to account for previously observed residuals. By explicitly modeling dynamic error covariance, our method enhances training flexibility, improves time series prediction accuracy, and provides high-quality uncertainty quantification. Our main contributions are as follows:

- We propose a novel method that enhances the training and prediction of univariate probabilistic time series models by learning a time-varying covariance matrix that captures the correlated errors within a mini-batch.
- We parameterize the dynamic correlation matrix with a weighted sum of several base kernel matrices. This ensures that the correlation matrix is a positive definite symmetric matrix with unit diagonals. This approach allows us to jointly learn the dynamic weights alongside the base model.
- We evaluate the effectiveness of the proposed approach on two base models with distinct architectures, DeepAR and Transformer, using multiple public datasets. Our method effectively captures the autocorrelation in errors and thus offers enhanced prediction quality. Importantly, these improvements are achieved through a statistical formulation without substantially increasing the number of parameters in the model.

2 PRELIMINARIES

2.1 Probabilistic Time Series Forecasting

Denote $\mathbf{z}_t = [z_{1,t}, \dots, z_{N,t}]^\top \in \mathbb{R}^N$ the time series variables at time step t , where N is the number of time series. Given the observed history $\{\mathbf{z}_t\}_{t=1}^T$, the task of probabilistic time series forecasting involves formulating the estimation of the joint conditional distribution $p(\mathbf{z}_{T+1:T+Q} | \mathbf{z}_{T-P+1:T}; \mathbf{x}_{T-P+1:T+Q})$. Here, $\mathbf{z}_{t_1:t_2} = [\mathbf{z}_{t_1}, \dots, \mathbf{z}_{t_2}]$, and \mathbf{x}_t represents known time-dependent or time-independent covariates, such as the

time of day or the time series identifier. Put differently, our focus is on predicting Q future values based on P historical values and covariates. The model can be further decomposed as

$$p(\mathbf{z}_{T+1:T+Q} | \mathbf{z}_{T-P+1:T}; \mathbf{x}_{T-P+1:T+Q}) = \prod_{t=T+1}^{T+Q} p(\mathbf{z}_t | \mathbf{z}_{t-P:t-1}; \mathbf{x}_{t-P:t}), \quad (1)$$

which is an autoregressive model that can be used for performing multi-step-ahead forecasting in a rolling manner. In this case, samples are drawn within the prediction range ($t \geq T+1$) and fed back for the subsequent time step until reaching the end of the prediction range. The conditioning is usually expressed as a state vector \mathbf{h}_t of a transition dynamics f_Θ that evolves over time $\mathbf{h}_t = f_\Theta(\mathbf{h}_{t-1}, \mathbf{z}_{t-1}, \mathbf{x}_t)$. Hence, Eq. (1) can be expressed in a simplified form as

$$p(\mathbf{z}_{T+1:T+Q} | \mathbf{z}_{T-P+1:T}; \mathbf{x}_{T-P+1:T+Q}) = \prod_{t=T+1}^{T+Q} p(\mathbf{z}_t | \mathbf{h}_t), \quad (2)$$

where \mathbf{h}_t is mapped to the parameters of a specific parametric distribution (e.g., Gaussian, Poisson). When $N = 1$, the problem is reduced to a univariate model:

$$p(z_{i,T+1:T+Q} | z_{i,T-P+1:T}; \mathbf{x}_{i,T-P+1:T+Q}) = \prod_{t=T+1}^{T+Q} p(z_{i,t} | \mathbf{h}_{i,t}), \quad (3)$$

where i is the identifier of a time series.

2.2 Error Autocorrelation

In the majority of probabilistic time series forecasting literature, the data under consideration is typically continuous, and the errors are assumed to follow an independent Gaussian distribution. Consequently, the time series variable associated with this framework is expected to follow a Gaussian distribution:

$$z_{i,t} | \mathbf{h}_{i,t} \sim \mathcal{N}(\mu(\mathbf{h}_{i,t}), \sigma^2(\mathbf{h}_{i,t})), \quad (4)$$

where $\mu(\cdot)$ and $\sigma(\cdot)$ map the state vector $\mathbf{h}_{i,t}$ to the mean and standard deviation of a Gaussian distribution. For instance, the DeepAR model (Salinas et al., 2020) adopts $\mu(\mathbf{h}_{i,t}) = \mathbf{w}_\mu^\top \mathbf{h}_{i,t} + b_\mu$ and $\sigma(\mathbf{h}_{i,t}) = \log(1 + \exp(\mathbf{w}_\sigma^\top \mathbf{h}_{i,t} + b_\sigma))$ and both parameters are time-varying. With Eq. (4), we can decompose the time series variable into $z_{i,t} = \mu_{i,t} + \eta_{i,t}$, where $\eta_{i,t} \sim \mathcal{N}(0, \sigma_{i,t}^2)$. Assuming the errors to be independent corresponds to $\text{Cov}(\eta_{i,t-\Delta}, \eta_{i,t}) = 0, \forall \Delta \neq 0$. In




Figure 1: Autocorrelation Function (ACF) of the One-step-ahead Prediction Residuals. The results depict the prediction outcomes generated by DeepAR for two time series in the `m4_hourly` dataset. The shaded area indicates the 95% confidence interval, highlighting regions where the correlation is statistically insignificant.

the following of this paper, we focus on this setting with Gaussian errors.

When there exists serial correlation in the error process, we will have $\boldsymbol{\eta}_{T+1:T+Q} = [\eta_{i,T+1}, \dots, \eta_{i,T+Q}]^\top$ follows a multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma}$ is the covariance matrix. Fig. 1 gives an example of residual autocorrelation functions (ACF) calculated using the prediction results of DeepAR on two time series from the `m4_hourly` dataset. The plot reveals a prevalent issue of lag-1 autocorrelation. Ignoring the systematic autocorrelation will undermine the performance of forecasting.

3 RELATED WORK

3.1 Probabilistic Time Series Forecasting

Probabilistic forecasting aims to offer the predictive distribution of the target variable rather than producing a single-point estimate, as seen in deterministic forecasting. Essentially, there are two approaches to achieve this: via the probability density function (PDF) or through the quantile function (Benidis et al., 2022). For example, MQ-RNN (Wen et al., 2017) employs sequence-to-sequence (Seq2Seq) recurrent neural networks (RNNs) to directly output specific quantiles of the predictive distribution.

In contrast, PDF-based models typically use a probabilistic model to describe the distribution of the target variables, with neural networks often employed to

generate the parameters of this probabilistic model. For example, DeepAR (Salinas et al., 2020) employs RNNs to model the transitions of hidden state. The hidden state at each time step is used to generate the parameters of a Gaussian distribution. Consequently, prediction samples for the target variables can be drawn from this distribution. GPVar (Salinas et al., 2019), a multivariate extension of DeepAR, utilizes a Gaussian copula to transform the original observations into Gaussian variables. Subsequently, a multivariate Gaussian distribution is assumed for these transformed variables. State space model (SSM) is also a popular choice for the probabilistic model. For instance, the deep SSM model proposed by Rangapuram et al. (2018) employs RNNs to generate parameters for the state space model, facilitating the generation of prediction samples. The normalizing kalman filter (de Bézenac et al., 2020) combines normalizing flows with the linear Gaussian state space model. This integration enables the modeling of nonlinear dynamics using RNNs and a more flexible probability density function for observations with normalizing flows. The deep factor model (Wang et al., 2019) employs a deterministic model and a probabilistic model separately to capture the global and local random effects of time series. The probabilistic model could be any classical probabilistic time series model such as a Gaussian noise process. Conversely, the global model, parameterized by neural networks, is dedicated to representing the deterministic patterns inherent in time series.

Various efforts have been undertaken to improve the quality of probabilistic forecasting. One avenue involves enhancing expressive conditioning for probabilistic models. For example, some approaches involve replacing RNNs with Transformer to model latent state dynamics, thus mitigating the Markovian assumption inherent in RNNs (Tang and Matteson, 2021). Another approach focuses on adopting more sophisticated distribution forms, such as normalizing flows (Rasul et al., 2020) and diffusion models (Rasul et al., 2021). For a recent and comprehensive review, readers are referred to Benidis et al. (2022).

3.2 Modeling Correlated Errors

Error correlation can be categorized into two main types: autocorrelation and contemporaneous correlation. In time series analysis, autocorrelation occurs when errors in a time series are correlated over different time points, while contemporaneous correlation refers to the correlation between errors at the same time step.

Autocorrelation has been extensively studied in classical time series models (Prado et al., 2021; Hynman and Athanasopoulos, 2018; Hamilton, 2020).

Statistical frameworks, including autoregressive (AR) and moving average (MA) models, have been well-developed to address autocorrelation. One important method is dynamic regression (Hyndman and Athanassopoulos, 2018), where errors are assumed to follow an ARIMA process. Recent DL-based models have also attempted to handle autocorrelation, such as reparameterizing the input and output of neural networks to model first-order error autocorrelation with an AR process (Sun et al., 2021). This method enhances the performance of DL-based time series models, and the parameters introduced to model autocorrelation can be jointly optimized with the base DL model. However, it is limited to one-step-ahead forecasting. This method was later extended to multivariate models for Seq2Seq traffic forecasting tasks by Zheng et al. (2023), assuming a matrix AR process for the matrix-valued errors.

Contemporaneous correlation modeling often appears in spatial regression tasks. For instance, Jia and Benson (2020) proposed using a multivariate Gaussian distribution to model label correlation in the node regression problem, where the predicted label at each node is often considered conditionally independent. This correlation can later be utilized to refine predictions for unknown nodes using information from known node labels. The introduction of GLS loss in Zhan and Datta (2023) captures the spatial correlation of errors in neural networks for geospatial data, bridging deep learning with Gaussian processes. A similar approach, involving the utilization of GLS loss with random forest, was proposed by Saha et al. (2023). Contemporaneous correlation is also modeled in time series forecasting to capture the interdependence between time series. The methods range from using a parametric multivariate Gaussian distribution (Salinas et al., 2019) to more expressive generative models such as normalizing flows (Rasul et al., 2020) and diffusion models (Rasul et al., 2021). In Choi et al. (2022), a dynamic mixture of matrix normal distributions was proposed to characterize spatiotemporally correlated errors in multivariate Seq2Seq traffic forecasting tasks.

To the best of our knowledge, our study presents an innovative training approach to addressing error autocorrelation in probabilistic time series forecasting. Our proposed method is closely related to those introduced in Sun et al. (2021), Zhan and Datta (2023), and Saha et al. (2023). While Zhan and Datta (2023), and Saha et al. (2023) primarily concentrate on modeling contemporaneous correlation, and Sun et al. (2021) focuses on modeling autocorrelation in deterministic forecasting. Our method concentrates on learning temporally correlated errors in the probabilistic time series forecasting context. We utilize a dynamic covari-

ance matrix to capture autocorrelation within a mini-batch, which is simultaneously learned alongside the base model. The introduction of the error covariance matrix not only provides a statistical framework for characterizing error autocorrelation but also enhances prediction accuracy.

4 OUR METHOD

Our approach is based on the formulation presented in Eq. (3), using an autoregressive model as the base model. Given the primary focus of this paper on univariate models, we will omit the subscript i for the remainder of this paper. In a general sense, an autoregressive probabilistic forecasting model comprises two key components: firstly, a transition model (e.g., an RNNs) to characterize state transitions \mathbf{h}_t , and secondly, a distribution head θ responsible for mapping \mathbf{h}_t to the parameters governing the desired distribution. Furthermore, the encoder-decoder framework is employed to facilitate multi-step forecasting, wherein an input sequence spanning P time steps is used to generate an output sequence spanning Q time steps. The likelihood is expressed as $p(z_t | \theta(\mathbf{h}_t))$ for an individual observation, and in the case of employing a Gaussian distribution, $\theta(\mathbf{h}_t)$ takes the form of (μ_t, σ_t) . In the training batch, the target time series variable can be decomposed as

$$z_t = \mu_t + \sigma_t \epsilon_t, \quad (5)$$

where ϵ_t is the normalized error term, which usually follows the assumption that $\epsilon_t = \frac{z_t - \mu_t}{\sigma_t} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. This assumption implies that ϵ_t are independent and identically distributed according to a standard normal distribution. Consequently, the parameters of the model can be optimized by maximizing its log-likelihood:

$$\mathcal{L} = \sum_{t=1}^T \log p(z_t | \theta(\mathbf{h}_t)) \propto \sum_{t=1}^T -\frac{1}{2} \epsilon_t^2 - \ln \sigma_t. \quad (6)$$

We adopt a unified univariate model trained across all time series, rather than training individual models for each time series. Moreover, if we assume the error process to be isotropic, the loss function is equivalent to the Mean Squared Error (MSE) commonly employed in training deterministic models (Sun et al., 2021). However, this assumption of independence ignores the potential serial correlation in ϵ_t .

4.1 Training with Mini-batch

We propose a novel training approach by constructing mini-batches instead of using individual training instances. For most existing deep probabilistic time series models including DeepAR, each training instance

consists of a time series segment with a length of $P+Q$, where P represents the conditioning range and Q denotes the prediction range. However, as mentioned, this simple approach cannot characterize the serial correlation of errors among consecutive time steps. To address this issue, we group D consecutive time series segments into a mini-batch, with each segment having a length of $P+1$ (i.e., $Q=1$). In other words, the new training instance (i.e., a mini-batch) becomes a collection of D time series segments with a prediction range $Q=1$. The composition of a mini-batch is illustrated in Fig. 2. An example of the collection of target variables in a mini-batch of size D (the time horizon we use for capturing serial correlation) is given by

$$\begin{aligned} z_t &= \mu_t + \sigma_t \epsilon_t, \\ z_{t+1} &= \mu_{t+1} + \sigma_{t+1} \epsilon_{t+1}, \\ &\dots \\ z_{t+D-1} &= \mu_{t+D-1} + \sigma_{t+D-1} \epsilon_{t+D-1}, \end{aligned} \quad (7)$$

where for time point t' , $\mu_{t'}$ and $\sigma_{t'}$ are the output of the model for each time series segment in the mini-batch, and $\epsilon_{t'}$ is the normalized error term. We use boldface symbols to denote the vectors of data and parameters in this mini-batch, e.g., $\mathbf{z}_t^{\text{bat}} = [z_t, z_{t+1}, \dots, z_{t+D-1}]^\top$ and the same notation applies to $\boldsymbol{\mu}_t^{\text{bat}}$ and $\boldsymbol{\sigma}_t^{\text{bat}}$.

Rather than assuming independence among the normalized errors, we consider modeling the joint distribution of the error vector in a mini-batch $\boldsymbol{\epsilon}_t$, denoted as $\boldsymbol{\epsilon}_t^{\text{bat}} = [\epsilon_t, \epsilon_{t+1}, \dots, \epsilon_{t+D-1}]^\top \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_t)$, where \mathbf{C}_t is a time-varying correlation matrix. To efficiently characterize the time-varying patterns, we parameterize \mathbf{C}_t as a dynamic weighted sum of several base kernel matrices: $\mathbf{C}_t = \sum_{m=1}^M w_{m,t} \mathbf{K}_m$, where $w_{m,t} \geq 0$ (with $\sum_m w_{m,t} = 1$) is the component weight. For each base component, we use a kernel matrix generated from a squared-exponential (SE) kernel function $\mathbf{K}_m^{ij} = \exp(-\frac{(i-j)^2}{l_m^2})$ with different lengthscales l_m (e.g., $l_m = 1, 2, 3, \dots$). An identity matrix is included in the additive structure to capture the independent noise process. Taken together, this parameterization ensures that \mathbf{C}_t is a positive definite symmetric matrix with unit diagonals, thus being a valid correlation matrix. A small neural network is attached to the original model to project the hidden state to the weights by setting the number of nodes in the final hidden layer to M (i.e., the number of components). A softmax layer is used as the output layer to ensure that these weights sum up to 1. The parameters of the small neural network can be learned jointly with the base model.

The utilization of a time-varying correlation matrix, as opposed to a static correlation matrix, offers the advantage of enabling the model to adapt dynami-

cally to the evolving structure of the error process. For example, the model can assign a higher weight to a kernel matrix generated by a kernel function with a lengthscales of $l=3$ when strong and long-range correlations are present in the current context, whereas it can favor the identity matrix when errors become white noise. This parameterization empowers the model to capture positive autocorrelation that diminishes over time lags. Alternatively, one could opt for a fully learnable positive definite symmetric Toeplitz matrix to parameterize the correlation matrix, which can also accommodate negative and complex correlations. With this formulation, the distribution of $\boldsymbol{\epsilon}_t$ becomes a multivariate Gaussian, which also leads to $\mathbf{z}_t^{\text{bat}} \sim \mathcal{N}(\boldsymbol{\mu}_t^{\text{bat}}, \boldsymbol{\Sigma}_t^{\text{bat}})$. The $D \times D$ covariance of the associated target variables can be decomposed as $\boldsymbol{\Sigma}_t^{\text{bat}} = \text{diag}(\boldsymbol{\sigma}_t^{\text{bat}}) \mathbf{C}_t \text{diag}(\boldsymbol{\sigma}_t^{\text{bat}})$. As both $\boldsymbol{\mu}_t^{\text{bat}}$ and $\boldsymbol{\sigma}_t^{\text{bat}}$ are default outputs of the base probabilistic model, the likelihood for a specific time series can be constructed for each mini-batch, and the overall likelihood is given by

$$\mathcal{L} = \sum_{t=1}^{T-D+1} \log p(\mathbf{z}_t^{\text{bat}} | \boldsymbol{\mu}_t^{\text{bat}}, \boldsymbol{\Sigma}_t^{\text{bat}}). \quad (8)$$

By allowing overlap, a total of $T-D+1$ mini-batches can be generated for each time series from the training data.

4.2 Multi-step Rolling Prediction

An autoregressive model performs forecasting in a rolling manner by drawing a sample of the target variable at each time step, feeding it to the next time step as input, and continuing this process until the desired prediction range is reached. Our method can provide extra calibration for this process using the proposed correlation matrix \mathbf{C}_t . Assume that we have observations till time step t and recall that the collection of normalized errors in a mini-batch jointly follows a multivariate Gaussian distribution. For the next time step ($t+1$) to be predicted, we have the conditional distribution of ϵ_{t+1} given the past $(D-1)$ errors using the conditional distribution properties of jointly Gaussian variables:

$$\begin{aligned} \epsilon_{t+1} | \epsilon_t, \epsilon_{t-1}, \dots, \epsilon_{t-D+2} \\ \sim \mathcal{N}(\mathbf{C}_* \mathbf{C}_{\text{obs}}^{-1} \boldsymbol{\epsilon}_{\text{obs}}, 1 - \mathbf{C}_* \mathbf{C}_{\text{obs}}^{-1} \mathbf{C}_*^\top), \end{aligned} \quad (9)$$

where $\boldsymbol{\epsilon}_{\text{obs}} = [\epsilon_{t-D+2}, \dots, \epsilon_{t-1}, \epsilon_t]^\top \in \mathbb{R}^{D-1}$ represents the set of observed $(D-1)$ residuals at forecasting step $t+1$. Here, \mathbf{C}_{obs} denotes the $(D-1) \times (D-1)$ partition of \mathbf{C}_{t+1} that captures the correlations within $\boldsymbol{\epsilon}_{\text{obs}}$, and \mathbf{C}_* denotes the $1 \times (D-1)$ partition of \mathbf{C}_{t+1} that captures the correlations between ϵ_{t+1} and $\boldsymbol{\epsilon}_{\text{obs}}$,




Figure 2: Example of a Mini-batch. Colored parts denote the time series segments that construct the mini-batch. Each mini-batch has a time span of $P + D$, and only one-step-ahead prediction is involved during training.

i.e., $\mathbf{C}_{t+1} = \begin{bmatrix} \mathbf{C}_{\text{obs}} & \mathbf{C}_*^\top \\ \mathbf{C}_* & 1 \end{bmatrix}$, where the weights for generating \mathbf{C}_{t+1} are obtained from the hidden state of the base model at time step $t+1$. Note that we remove the time index in \mathbf{C}_{obs} , \mathbf{C}_* and ϵ_{obs} for brevity. To obtain a sample of the target variable \tilde{z}_{t+1} , we can first draw a sample of the normalized error $\tilde{\epsilon}_{t+1}$ using Eq. (9). Since both μ_{t+1} and σ_{t+1} are deterministic outcomes from the base model, \tilde{z}_{t+1} can be derived by

$$\tilde{z}_{t+1} = \mu_{t+1} + \sigma_{t+1} \tilde{\epsilon}_{t+1}. \quad (10)$$

Based on Eq. (9) and Eq. (10), it can be seen that the final distribution for z_{t+1} becomes

$$z_{t+1} \mid \mathbf{h}_{t+1}, \epsilon_{\text{obs}} \sim \mathcal{N}(\bar{\mu}_{t+1}, \bar{\sigma}_{t+1}^2), \quad (11)$$

where

$$\begin{aligned} \bar{\mu}_{t+1} &= \mu_{t+1} + \sigma_{t+1} \mathbf{C}_* \mathbf{C}_{\text{obs}}^{-1} \epsilon_{\text{obs}}, \\ \bar{\sigma}_{t+1}^2 &= \sigma_{t+1}^2 \left(1 - \mathbf{C}_* \mathbf{C}_{\text{obs}}^{-1} \mathbf{C}_*^\top \right). \end{aligned} \quad (12)$$

Multi-step rolling prediction can be accomplished by treating the sample $\tilde{\epsilon}_{t+1}$ as a newly observed residual. Following this process for all subsequent time steps results in a trajectory of $\{\tilde{\epsilon}_{t+q}\}_{q=1}^Q$.

5 EXPERIMENTS

5.1 Datasets and Models

We apply the proposed framework to two base prediction models: DeepAR (Salinas et al., 2020) and an autoregressive decoder-only Transformer (i.e., the GPT model (Radford et al., 2018)). A Gaussian distribution head is employed to generate the distribution parameters for probabilistic forecasting based on the hidden state outputted by the model. It should be noted that our approach can be applied to other autoregressive univariate models without any loss of generality, as long as the final prediction follows a Gaussian distribution. We implemented these models using PyTorch Forecasting (Beitner, 2020). Both models utilize input data consisting of lagged time series values from

the preceding time step, accompanied by supplementary features including time of day, day of the week, and unique time series identifiers. We refer readers to the Supplementary Materials (SM) for comprehensive information on the experiment setup.

To assess our method’s effectiveness, we conducted extensive experiments on diverse real-world time series datasets sourced from GluonTS (Alexandrov et al., 2020). These datasets serve as important benchmarks in evaluating time series forecasting models. The prediction range (Q) and the number of rolling evaluations were acquired from each dataset’s configuration within GluonTS. Sequential splits into training, validation, and testing sets were performed for each dataset, with the validation set’s temporal length matching that of the testing sets. The temporal length of the testing set was computed using the prediction range and the number of rolling evaluations. For example, the testing set for `traffic` contains $24 + 7 - 1 = 30$ time steps, indicating 24-step predictions (Q) made in a rolling manner across 7 consecutive prediction start timestamps. The details of datasets are summarized in Table 1. Standardization of the data was carried out using the mean and standard deviation obtained from each time series within the training set. Predictions were subsequently rescaled to their original values for computing evaluation metrics.

5.2 Evaluation against Baseline

We evaluate the proposed approach by comparing it with models trained using Gaussian likelihood loss. To simplify the comparison and ensure fairness in terms of the data used during training, we set the autocorrelation range (D) to be identical to the prediction range (Q). This alignment ensures that each mini-batch in our method covers a time horizon of $P+D$, while in the conventional training method, each training instance spans a time horizon of $P+Q$. By setting $D=Q$, we guarantee that both methods involve the same amount of data per batch, given the same batch sizes. Further-

Table 1: Dataset Summary.

Dataset	Granularity	# of time series	# of time steps	Q	Rolling evaluation
m4_hourly	hourly	414	1,008	48	1
exchange_rate	workday	8	6,101	30	5
m1_quarterly	quarterly	203	48	8	1
pems03	5min	358	26,208	12	24
pems08	5min	170	17,856	12	24
solar	hourly	137	7,033	24	7
traffic	hourly	963	4,025	24	7
uber_daily	daily	262	181	7	1
uber_hourly	hourly	262	4,344	24	1

Table 2: CRPS_{sum} Accuracy Comparison.

	DeepAR			Transformer		
	w/o	w/	rel. impr.	w/o	w/	rel. impr.
m4_hourly	0.1529±0.0011	0.1421±0.0003	7.06%	0.1487±0.0011	0.1431±0.0009	3.77%
exchange_rate	0.0069±0.0001	0.0059±0.0000	14.49%	0.0081±0.0001	0.0074±0.0001	8.64%
m1_quarterly	0.3767±0.0006	0.3076±0.0018	18.34%	0.4449±0.0034	0.3302±0.0046	25.78%
pems03	0.0870±0.0000	0.0811±0.0000	6.78%	0.0907±0.0000	0.0845±0.0001	6.84%
pems08	0.0650±0.0001	0.0588±0.0000	9.54%	0.0622±0.0000	0.0591±0.0001	4.98%
solar	0.7627±0.0008	0.7063±0.0005	7.39%	0.6657±0.0010	0.5379±0.0010	19.20%
traffic	0.2765±0.0001	0.2387±0.0001	13.67%	0.2422±0.0002	0.2036±0.0001	15.94%
uber_daily	0.0897±0.0003	0.0890±0.0001	0.78%	0.0827±0.0003	0.0827±0.0002	0.00%
uber_hourly	0.1549±0.0007	0.1532±0.0005	1.10%	0.1503±0.0007	0.1462±0.0008	2.73%
		avg. rel. impr.	8.80%		avg. rel. impr.	9.76%

Note: The better results are in boldface (lower is better). All results are based on three runs of each model.

Table 3: Quantile Loss Accuracy Comparison using DeepAR.

	0.5-risk			0.9-risk		
	w/o	w/	rel. impr.	w/o	w/	rel. impr.
m4_hourly	0.1066±0.0006	0.0999±0.0006	6.29%	0.0579±0.0006	0.0497±0.0003	14.16%
exchange_rate	0.0043±0.0001	0.0042±0.0000	2.33%	0.0028±0.0000	0.0018±0.0000	35.71%
m1_quarterly	0.2011±0.0004	0.1671±0.0016	16.91%	0.3101±0.0013	0.2489±0.0030	19.74%
pems03	0.0601±0.0001	0.0563±0.0000	6.32%	0.0337±0.0000	0.0287±0.0000	14.84%
pems08	0.0456±0.0001	0.0410±0.0000	10.09%	0.0208±0.0000	0.0187±0.0000	10.10%
solar	0.5336±0.0009	0.4896±0.0005	8.25%	0.1827±0.0001	0.1699±0.0004	7.01%
traffic	0.1608±0.0000	0.1402±0.0001	12.81%	0.1167±0.0001	0.1051±0.0001	9.94%
uber_daily	0.0598±0.0006	0.0603±0.0004	-0.84%	0.0392±0.0005	0.0369±0.0002	5.87%
uber_hourly	0.1091±0.0005	0.1040±0.0003	4.67%	0.0541±0.0003	0.0574±0.0003	-6.10%
		avg. rel. impr.	7.42%		avg. rel. impr.	12.36%

Note: The better results are in boldface (lower is better). All results are based on three runs of each model.

more, we follow the default configuration in GluonTS by setting the context range equal to the prediction range, i.e., $P = Q$.

In our proposed method, we introduce a small number of additional parameters dedicated to projecting the hidden state into component weights $w_{m,t}$, which play a pivotal role in generating the dynamic correlation matrix C_t . In practice, the selection of base kernels should be data-specific, and one should perform residual analysis to determine the most appropriate structure. For simplicity, we use $M = 4$ base

kernels—three SE kernels with $l = 1, 2, 3$, respectively, and an identity matrix. The different lengthscales capture different decaying rates of autocorrelation. The time-varying component weights can help the model dynamically adjust to different correlation structures observed at different time points.

We use two different probabilistic scores, namely, the Continuous Ranked Probability Score (CRPS) (Gneiting and Raftery, 2007) and the quantile loss (ρ -risk) (Salinas et al., 2020), as the evaluation metrics for un-

certainty quantification. CRPS is defined as

$$\text{CRPS}(F, y) = \mathbb{E}_F |Y - y| - \frac{1}{2} \mathbb{E}_F |Y - Y'|, \quad (13)$$

where y is the observation, F is the cumulative distribution function (CDF) of the target variable z_t , Y and Y' are independent copies of a set of prediction samples associated with this distribution. To facilitate direct comparison between datasets, we use CRPS_{sum} calculated by first summing CRPS across the entire testing horizon of all time series and then normalizing the results by the sum of the corresponding observations. The quantile loss is defined as

$$L_\rho(Z, \hat{Z}^\rho) = 2(\hat{Z}^\rho - Z) \left((1 - \rho) \mathbb{I}_{\hat{Z}^\rho > Z} - \rho \mathbb{I}_{\hat{Z}^\rho \leq Z} \right), \quad (14)$$

where \mathbb{I} is a binary indicator function that equals 1 when the condition is met, \hat{Z}^ρ represents the predicted ρ -quantile, and Z represents the ground truth value. The quantile loss serves as a metric to assess the accuracy of a given quantile, denoted by ρ , from the predictive distribution. Following Salinas et al. (2020), we summarize the quantile losses for the testing set across all time series segments by evaluating a normalized summation of these losses: $(\sum_i L_\rho(Z_i, \hat{Z}_i^\rho)) / (\sum_i Z_i)$. In this paper, we evaluate the 0.5-risk and the 0.9-risk. Comprehensive details regarding the computation of the CRPS and ρ -risk can be found in the SM.

In Table 2, we present a comparative analysis of prediction performance using CRPS_{sum} . The variants combined with our proposed method (denoted as “w/”) are compared to their respective original implementations optimized with Gaussian likelihood loss (denoted as “w/o”). The results demonstrate the effectiveness of our approach in enhancing the performance of both models across a wide range of datasets, yielding notable improvements in predictive accuracy. Our method yields an average improvement of 8.80% for DeepAR and 9.76% for Transformer. It should be noted that our method exhibits versatile performance improvements that vary across different datasets. This variability can be attributed to a combination of factors, including the inherent characteristics of the data and the baseline performance of the original model on each specific dataset. Notably, when the original model already achieves exceptional performance on a particular dataset, our method could demonstrate minimal enhancement (e.g., `uber_daily` and `uber_hourly`). Moreover, the degree of alignment between our kernel assumption and the true error autocorrelation structure significantly impacts the performance of our method.

Furthermore, an additional comparison using quantile loss with DeepAR is shown in Table 3, present-

ing the outcomes for 0.5-risk and 0.9-risk. A noteworthy observation is the greater improvement in 0.9-risk (12.36%) compared to 0.5-risk (7.42%) with our method, implying that the model trained using our approach may provide more reliable prediction intervals for quantifying data uncertainty. Further analyses using other evaluation metrics such as MSE alongside naive baselines are also reported in the SM.

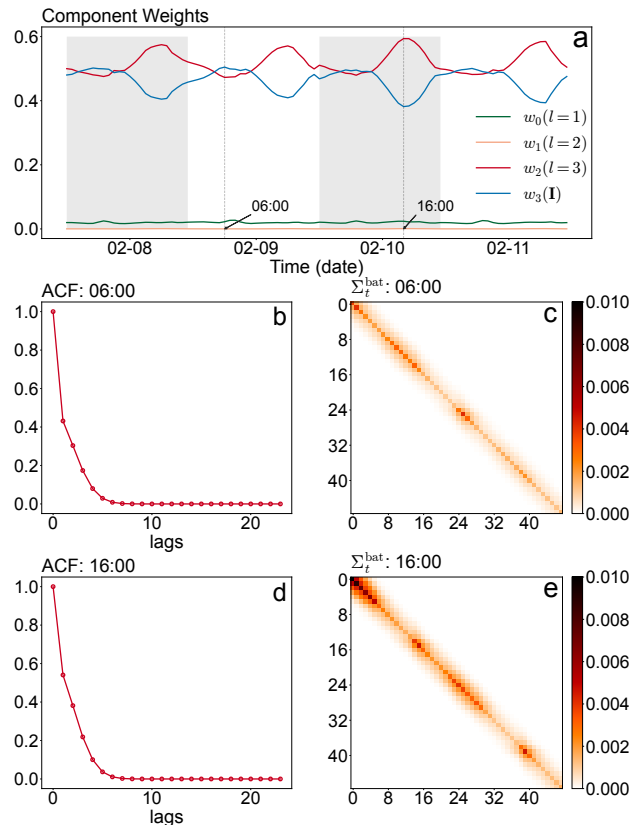


Figure 3: (a) Component weights for generating the correlation matrix of an example time series from the `m4_hourly` dataset. Parameters w_0, w_1, w_2 represent the component weights for kernel matrices associated with lengthscales $l = 1, 2, 3$, respectively, and w_3 is the component weight for the identity matrix. Shaded areas distinguish different days; (b) and (d): The auto-correlation function (ACF) indicated by the resulting error correlation matrix \mathbf{C}_t at 6:00 and 16:00. Given the rapid decay of the ACF, we only plot 24 lags to enhance clarity in our visualization; (c) and (e): The corresponding covariance matrix of the associated target variables $\Sigma_t^{\text{bat}} = \text{diag}(\sigma_t^{\text{bat}}) \mathbf{C}_t \text{diag}(\sigma_t^{\text{bat}})$ at 6:00 and 16:00, respectively.

5.3 Interpretation of Correlation

A key element of our method lies in its ability to capture error autocorrelation through the dynamic construction of a covariance matrix. This is achieved

by introducing a dynamic weighted sum of kernel matrices with different lengthscales. The choice of lengthscale significantly influences the structure of autocorrelation—a small lengthscale corresponds to short-range positive autocorrelation, while a large lengthscale can capture positive correlation spanning long lags.

In Fig. 3, we present the generated component weights and the resulting autocorrelation function (i.e., the first row in the learned correlation matrix \mathbf{C}_t) of an example time series from the `m4_hourly` dataset over a four-day duration. In particular, we observe that the component weights (w_0, w_1), corresponding to the kernel matrices with $l = 1$ and $l = 2$, consistently remain close to zero across the entire observation period. This suggests that the prevailing autocorrelation structure in this dataset is most effectively characterized by the kernel matrix associated with $l = 3$.

Furthermore, we observe the dynamic adjustment of correlation strengths facilitated by the identity matrix over time. Specifically, when w_3 (weight assigned to the identity matrix) increases, the error process tends to exhibit greater independence. In contrast, when the weight w_2 for the kernel matrix with $l = 3$ is larger, the error process becomes more correlated. Figs. 3 (b, d) reveal that the autocorrelation at 6:00 in the morning is less pronounced compared to that observed at 16:00. Additionally, Fig. 3 (a) demonstrates pronounced daily patterns in autocorrelation, particularly when errors exhibit an increased correlation around 16:00 each day. This underscores the crucial need for our methodology to dynamically adapt the covariance matrix, enabling the effective modeling of these temporal variations. Figs. 3 (c, e) depict the covariance matrix of the respective target variables within the autocorrelation horizon. The diagonal elements represent the variance of the target variables generated by the base model, while the off-diagonal elements depict the covariance of the target variables that are facilitated by our approach.

5.4 Computational Cost

The proposed method introduces additional computational cost due to the evaluation of the log-likelihood for the multivariate Gaussian distribution in Eq. (8). Specifically, this evaluation involves calculating the inverse and determinant of the covariance matrix Σ_t^{bat} . However, we can significantly expedite these computations by employing Cholesky decomposition, particularly when the dimensionality D is small. We provide a detailed account of the computational expenses in the SM. Remarkably, our method does not necessarily lead to an increase in training time per epoch or the number of epochs required for convergence. Consequently,

the total training time remains comparable to that of conventional training methods in certain scenarios.

6 CONCLUSION

This paper introduces an innovative training approach to model error autocorrelation in probabilistic time series forecasting. The method involves using mini-batches in training and learning a time-varying covariance matrix that captures the correlation among normalized errors within a mini-batch. Taken together with the standard deviation provided by the base model, we are able to model and predict a time-varying covariance matrix. We implement and evaluate the proposed method using DeepAR and Transformer on various public datasets, and our results confirm the effectiveness of the proposed solution in improving the quality of uncertainty quantification. The broader impact of our method can be observed in two aspects. First, since Gaussian errors are commonly assumed in probabilistic forecasting models, our method can be applied to enhance the training process of such models. Second, the learned autocorrelation can be leveraged to improve multi-step prediction by calibrating the distribution output at each forecasting step.

There are several directions for future research. First, the kernel-based covariance matrix may be too restrictive for capturing temporal processes. Investigating covariance structures with greater flexibility, such as parameterizing \mathbf{C}_t as a fully learnable positive definite symmetric Toeplitz matrix (e.g., AR(p) process has covariance with a Toeplitz structure) and directly factorizing the covariance matrix $\Sigma_t = \mathbf{U}_t \mathbf{U}_t^\top$ (e.g., Wishart process as in Wilson and Ghahramani (2011)) or the precision matrix $\Lambda_t = \Sigma_t^{-1} = \mathbf{V}_t \mathbf{V}_t^\top$ (e.g., using Cholesky factorization as in Fortuin et al. (2020)), could be promising avenues. Second, our method can be extended to multivariate models, in which the target output becomes a vector instead of a scalar. A possible solution is to use a matrix Gaussian distribution to replace the multivariate Gaussian distribution used in the current method. This would allow us to learn a full covariance matrix between different target series, thereby capturing any cross-correlations between them.

Acknowledgements

We express our gratitude to all reviewers for their insightful suggestions and comments. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). Vincent Zhihao Zheng acknowledges the support received from the FRQNT B2X Doctoral Scholarship Program.

References

- Alexandrov, A., Benidis, K., Bohlke-Schneider, M., Flunkert, V., Gasthaus, J., Januschowski, T., Maddix, D. C., Rangapuram, S., Salinas, D., Schulz, J., et al. (2020). Gluonts: Probabilistic and neural time series modeling in python. *The Journal of Machine Learning Research*, 21(1):4629–4634.
- Beitner, J. (2020). Pytorch forecasting. Feb 22, 2024.
- Benidis, K., Rangapuram, S. S., Flunkert, V., Wang, Y., Maddix, D., Turkmen, C., Gasthaus, J., Bohlke-Schneider, M., Salinas, D., Stella, L., et al. (2022). Deep learning for time series forecasting: Tutorial and literature survey. *ACM Computing Surveys*, 55(6):1–36.
- Choi, S., Saunier, N., Trepanier, M., and Sun, L. (2022). Spatiotemporal residual regularization with dynamic mixtures for traffic forecasting. *arXiv preprint arXiv:2212.06653*.
- de Bézenac, E., Rangapuram, S. S., Benidis, K., Bohlke-Schneider, M., Kurle, R., Stella, L., Hasson, H., Gallinari, P., and Januschowski, T. (2020). Normalizing kalman filters for multivariate time series analysis. *Advances in Neural Information Processing Systems*, 33:2995–3007.
- Fortuin, V., Baranchuk, D., Rätsch, G., and Mandt, S. (2020). Gp-vae: Deep probabilistic time series imputation. In *International Conference on Artificial Intelligence and Statistics*, pages 1651–1661. PMLR.
- Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.
- Hamilton, J. D. (2020). *Time series analysis*. Princeton University Press.
- Hyndman, R. J. and Athanasopoulos, G. (2018). *Forecasting: principles and practice*. OTexts.
- Jia, J. and Benson, A. R. (2020). Residual correlation in graph neural network regression. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 588–598.
- Prado, R., Ferreira, M. A., and West, M. (2021). *Time series: modeling, computation, and inference*. CRC Press.
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al. (2018). Improving language understanding by generative pre-training.
- Rangapuram, S. S., Seeger, M. W., Gasthaus, J., Stella, L., Wang, Y., and Januschowski, T. (2018). Deep state space models for time series forecasting. *Advances in Neural Information Processing Systems*, 31.
- Rasul, K., Seward, C., Schuster, I., and Vollgraf, R. (2021). Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*, pages 8857–8868. PMLR.
- Rasul, K., Sheikh, A.-S., Schuster, I., Bergmann, U., and Vollgraf, R. (2020). Multivariate probabilistic time series forecasting via conditioned normalizing flows. *arXiv preprint arXiv:2002.06103*.
- Saha, A., Basu, S., and Datta, A. (2023). Random forests for spatially dependent data. *Journal of the American Statistical Association*, 118(541):665–683.
- Salinas, D., Bohlke-Schneider, M., Callot, L., Medico, R., and Gasthaus, J. (2019). High-dimensional multivariate forecasting with low-rank gaussian copula processes. *Advances in Neural Information Processing Systems*, 32.
- Salinas, D., Flunkert, V., Gasthaus, J., and Januschowski, T. (2020). Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191.
- Sun, F.-K., Lang, C., and Boning, D. (2021). Adjusting for autocorrelated errors in neural networks for time series. *Advances in Neural Information Processing Systems*, 34:29806–29819.
- Tang, B. and Matteson, D. S. (2021). Probabilistic transformer for time series analysis. *Advances in Neural Information Processing Systems*, 34:23592–23608.
- Wang, Y., Smola, A., Maddix, D., Gasthaus, J., Foster, D., and Januschowski, T. (2019). Deep factors for forecasting. In *International Conference on Machine Learning*, pages 6607–6617. PMLR.
- Wen, R., Torkkola, K., Narayanaswamy, B., and Madeka, D. (2017). A multi-horizon quantile recurrent forecaster. *arXiv preprint arXiv:1711.11053*.
- Wilson, A. G. and Ghahramani, Z. (2011). Generalised wishart processes. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 736–744.
- Zhan, W. and Datta, A. (2023). Neural networks for geospatial data. *arXiv preprint arXiv:2304.09157*.
- Zheng, V. Z., Choi, S., and Sun, L. (2023). Enhancing deep traffic forecasting models with dynamic regression. *arXiv preprint arXiv:2301.06650*.

Checklist

1. For all models and algorithms presented, check if you include:

- (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes/No/Not Applicable]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes/No/Not Applicable]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes/No/Not Applicable]
2. For any theoretical claim, check if you include:
- (a) Statements of the full set of assumptions of all theoretical results. [Yes/No/Not Applicable]
 - (b) Complete proofs of all theoretical results. [Yes/No/Not Applicable]
 - (c) Clear explanations of any assumptions. [Yes/No/Not Applicable]
3. For all figures and tables that present empirical results, check if you include:
- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes/No/Not Applicable]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes/No/Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes/No/Not Applicable]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes/No/Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Yes/No/Not Applicable]
 - (b) The license information of the assets, if applicable. [Yes/No/Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes/No/Not Applicable]
 - (d) Information about consent from data providers/curators. [Yes/No/Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Yes/No/Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Yes/No/Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Yes/No/Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Yes/No/Not Applicable]

Better Batch for Deep Probabilistic Time Series Forecasting: Supplementary Materials

1 DATASET DETAILS

We conducted experiments on a diverse set of real-world datasets retrieved from GluonTS (Alexandrov et al., 2020), including:

- `m4_hourly` (Makridakis et al., 2020): This dataset consists of hourly time series data from various domains, covering microeconomics, macroeconomics, finance, industry, demographics, and other fields. The data originates from the M4-competition.
- `exchange_rate` (Lai et al., 2018): It provides daily exchange rate information for eight different countries spanning the period from 1990 to 2016.
- `m1_quarterly` (Makridakis et al., 1982): Quarterly time series data from seven different domains.
- `pems03` (Chen et al., 2001): Traffic flow records sourced from Caltrans District 3 and acquired through the Caltrans Performance Measurement System (PeMS). The records are aggregated at a 5-minute interval.
- `pems08` (Chen et al., 2001): Traffic flow records sourced from Caltrans District 8 and acquired through the Caltrans Performance Measurement System (PeMS). The records are aggregated at a 5-minute interval.
- `solar` (Lai et al., 2018): Hourly time series representing solar power production data in the state of Alabama for the year 2006.
- `traffic` (Caltrans, 2015): Hourly traffic occupancy rate recorded by sensors installed in the San Francisco freeway system between January 2008 and June 2008.
- `uber_daily` (Commission, 2015): Daily time series of Uber pickups in New York City from February to July 2015.
- `uber_hourly` (Commission, 2015): Hourly time series of Uber pickups in New York City from February to July 2015.

These datasets are widely used for benchmarking time series forecasting models. Each dataset follows its default configurations in GluonTS, including granularity and prediction range (Q). In our experiments, we set the context range (P) to match the prediction range, in line with the default setting in GluonTS. For simplicity, we configure the autocorrelation horizon (D) to also equal the prediction range (Q). Essentially, in this paper, we have $P = Q = D$.

2 CONTINUOUS RANKED PROBABILITY SCORE

Since the target time series variable is assumed to follow a normal distribution, the CRPS can be conveniently calculated as

$$\text{CRPS}(F, y) = y(2F(y) - 1) + 2f(y) - \frac{1}{\sqrt{\pi}}, \tag{1}$$

$$\text{CRPS}(F_{\mu, \sigma}, y) = \sigma \text{CRPS}\left(F, \frac{y - \mu}{\sigma}\right), \tag{2}$$

where y is the observation, for normal distribution, $f(x) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{x^2}{2})$, $F(x) = \int_{-\infty}^x f(t)dx$, and $F_{\mu, \sigma}(x) = F(\frac{x - \mu}{\sigma})$. μ and σ denote the mean and standard deviation evaluated by 100 samples drawn from the predictive distribution.

3 COMPARE WITH NAIVE BASELINES

We compare the proposed methods with two naive baselines: ARIMA (Box et al., 2015) and ETS exponential smoothing (Hyndman et al., 2008) using a more standard metric, namely mean squared error (MSE):

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n \left(Z_i - \hat{Z}_i \right)^2, \tag{3}$$

where \hat{Z}_i represents the sample mean for DeepAR and Transformer, computed from 100 samples drawn from the predictive distribution. In contrast, for ARIMA and ETS, the models provide deterministic forecasts directly. We conducted this experiment using the `auto.arima` function from the `pmdarima` package (Smith et al., 2017) and the `ETSModel` function from the `statsmodels` package (Seabold and Perktold, 2010). The `auto.arima` function can automatically determine the optimal order for an ARIMA model. The training data’s temporal length was set to be ten times that of the prediction range. The results are summarized in Table S1. Note that results for datasets with insufficient temporal length for training ARIMA models are not provided.

Table S1: MSE Accuracy Comparison.

	ARIMA	ETS	DeepAR		Transformer	
			w/o	w/	w/o	w/
<code>m4_hourly</code>	3416.2127	3718.8369	1482.3241	1358.1923	1295.8058	1357.3287
<code>exchange_rate</code>	0.0001	0.0001	0.0001	0.0001	0.0002	0.0001
<code>pems03</code>	1074.7741	1019.2916	618.2632	515.2057	647.8133	583.6474
<code>pems08</code>	1093.7592	793.8637	361.2450	325.1754	356.8268	320.8344
<code>solar</code>	1988.4526	1434.8300	1637.4979	1323.0383	1339.4904	949.7088
<code>traffic</code>	0.0013	0.0021	0.0011	0.0008	0.0009	0.0008
<code>uber_daily</code>	39688.8170	19413.8089	19366.8854	18370.4863	15591.5254	15718.5710
<code>uber_hourly</code>	371.6443	613.9445	64.9663	64.6354	62.1469	60.6224

Note: The better results between the “w/o” and “w/” variants are highlighted in boldface (lower values indicate better performance). All results are based on three runs of each model.

4 EXPERIMENT DETAILS

4.1 Model Architecture

In this paper, we employ DeepAR (Salinas et al., 2020) and an autoregressive decoder-only Transformer model, specifically, the GPT model (Radford et al., 2018), as our base prediction models. The architectures of these two models are visualized in Fig. S1 and Fig. S2. In the Transformer model, we adapt the output to the parameters of the predictive distribution, namely the mean and standard deviation. To achieve the autoregressive property, we apply a subsequent mask to the input sequence, ensuring that attention scores are computed exclusively based on the inputs preceding the current time step.

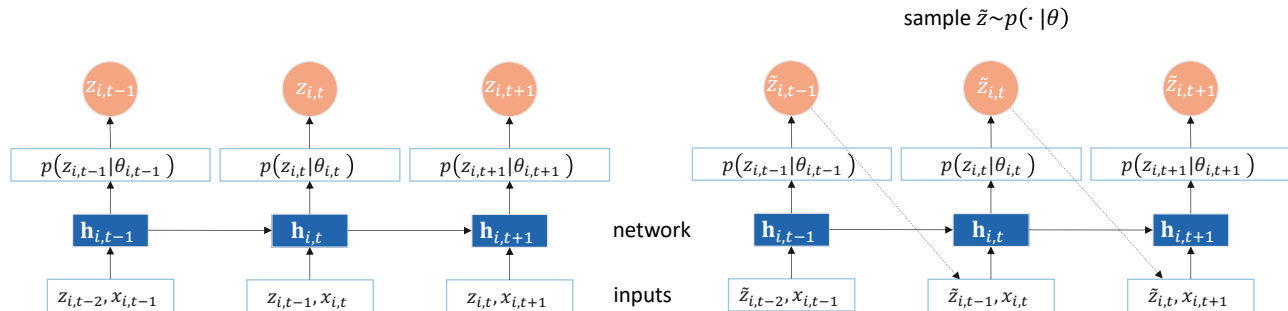


Figure S1: Summary of DeepAR (Salinas et al., 2020). Training (left): At each time step t , the inputs to the network are the target value from the previous time step $z_{i,t-1}$, the covariates $x_{i,t}$, and the hidden state from previous time step $\mathbf{h}_{i,t-1}$. The network output $\mathbf{h}_{i,t}$ is then used to compute the distribution parameters $\theta_{i,t} = (\mu_{i,t}, \sigma_{i,t})$. For prediction (right), a sample $\tilde{z} \sim p(\cdot | \theta_{i,t})$ is drawn and used as input for the next time step in the prediction range, resulting in the generation of a single sample trace. Repeating this prediction process generates multiple traces that collectively represent the joint predictive distribution. Within our framework, the component weights $\mathbf{w}_{i,t}$ are derived from the same hidden state $\mathbf{h}_{i,t}$ used for projecting the $\theta_{i,t}$.

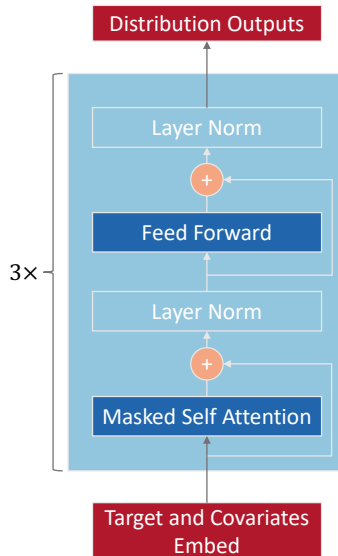


Figure S2: Summary of the Decoder-only Transformer (Radford et al., 2018). Similarly to DeepAR, Transformer uses ground truth values as inputs during training. However, during the prediction phase, we draw samples and utilize them as inputs for the subsequent time step within the prediction range.

4.2 Features

For datasets featuring hourly granularity, we incorporate two covariates: “hour of day” and “day of week”, as time-varying covariates. In the case of daily datasets, we exclusively employ “day of week” as the time-varying covariate. Furthermore, we utilize the identifier of each time series as a static covariate. These covariates are then concatenated with the RNN or Transformer input at each time step post-encoding by the embedding layer. Additionally, to normalize the target values of the time series, we scale them using the mean and standard deviation specific to each time series, obtained from the training dataset.

4.3 Hyperparameters

We employ batch sizes of 64, limiting each epoch to a maximum of 100 batches. The LSTM hyperparameters in DeepAR are adapted from Salinas et al. (2020): comprising three layers of LSTM, each containing 40 nodes, with a dropout rate of 0.1. Additionally, a single linear layer serves as the distribution head to output the distribution

parameters. Separate linear layers, along with ELU activation functions, are employed to produce component weights for the base kernel matrices. We configure the Transformer to have a comparable number of parameters with DeepAR. Specifically, we stack three Transformer decoding layers, each with a hidden size of 42, and utilize 2 attention heads ($H = 2$). The dropout rate for the Transformer is also set to 0.1.

4.4 Training Details

Every model underwent training for up to 100 epochs, employing a learning rate of 0.001 with the Adam optimizer. The training process was governed by an early stopping strategy with a patience of 10 epochs. Following the training phase, we restored the best-performing model, determined by validation loss. Details regarding training cost per epoch and the number of epochs until training termination are presented in Table S2.

Table S2: Training Cost Comparison.

	DeepAR				Transformer			
	w/o		w/		w/o		w/	
	time	epochs	time	epochs	time	epochs	time	epochs
m4_hourly	20	99	156	97	20	68	30	95
exchange_rate	17	57	92	33	25	48	28	39
m1_quarterly	15	10	32	10	28	17	18	18
pems03	76	65	50	50	84	51	46	45
pems08	49	25	39	79	68	34	35	99
solar	42	14	53	59	50	51	30	33
traffic	82	59	55	99	87	59	38	34
uber_daily	36	39	14	52	43	48	23	48
uber_hourly	34	46	33	74	38	67	20	98

Note: “w/o” implies the original implementation of models optimized with Gaussian likelihood loss, whereas “w/” implies the model combined with our method. “time” represents the training time per epoch in seconds.

4.5 Hardware Environment

Our experiments were conducted under a computer environment with one Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz and four NVIDIA Tesla V100 GPU.

References

- Alexandrov, A., Benidis, K., Bohlke-Schneider, M., Flunkert, V., Gasthaus, J., Januschowski, T., Maddix, D. C., Rangapuram, S., Salinas, D., Schulz, J., et al. (2020). Gluonts: Probabilistic and neural time series modeling in python. *The Journal of Machine Learning Research*, 21(1):4629–4634.
- Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time series analysis: forecasting and control*. John Wiley & Sons.
- Caltrans (2015). Caltrans performance measurement system. <https://pems.dot.ca.gov>.
- Chen, C., Petty, K., Skabardonis, A., Varaiya, P., and Jia, Z. (2001). Freeway performance measurement system: mining loop detector data. *Transportation Research Record*, 1748(1):96–102.
- Commission, N. T. L. (2015). Uber tlc foil response. <https://www.nyc.gov/site/tlc/index.page>.
- Hyndman, R., Koehler, A. B., Ord, J. K., and Snyder, R. D. (2008). *Forecasting with exponential smoothing: the state space approach*. Springer Science & Business Media.
- Lai, G., Chang, W.-C., Yang, Y., and Liu, H. (2018). Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 95–104.
- Makridakis, S., Andersen, A., Carbone, R., Fildes, R., Hibon, M., Lewandowski, R., Newton, J., Parzen, E., and Winkler, R. (1982). The accuracy of extrapolation (time series) methods: Results of a forecasting competition. *Journal of Forecasting*, 1(2):111–153.
- Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2020). The m4 competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, 36(1):54–74.
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al. (2018). Improving language understanding by generative pre-training.
- Salinas, D., Flunkert, V., Gasthaus, J., and Januschowski, T. (2020). Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191.
- Seabold, S. and Perktold, J. (2010). Statsmodels: Econometric and statistical modeling with python. In *Proceedings of the 9th Python in Science Conference*, volume 57, pages 10–25080. Austin, TX.
- Smith, T. G. et al. (2017). pmdarima: Arima estimators for Python. <https://www.alkaline-ml.com/pmdarima>.