

Hybrid variational autoencoder for time series forecasting

Borui Cai ^{a,*}, Shuiqiao Yang ^b, Longxiang Gao ^c, Yong Xiang ^a

^a School of Information Technology, Deakin University, Burwood, 3125, VIC, Australia

^b School of Computer Science and Engineering, University of New South Wales, Sydney, 2032, NSW, Australia

^c Qilu University of Technology (Shandong Academy of Sciences), Jinan, China

ARTICLE INFO

Keywords:

Time series forecasting
Variational autoencoder
Deep learning

ABSTRACT

Variational autoencoders (VAE) are powerful generative models that learn the latent representations of input data as random variables. Recent studies show that VAE can flexibly learn the complex temporal dynamics of time series and achieve more promising forecasting results than deterministic models. However, a major limitation of existing works is that they fail to jointly learn the local patterns (e.g., seasonality and trend) and temporal dynamics of time series for forecasting. Accordingly, we propose a novel hybrid variational autoencoder (HyVAE) to integrate the learning of local patterns and temporal dynamics by variational inference for time series forecasting. Experimental results on four real-world datasets show that the proposed HyVAE achieves better forecasting results than various counterpart methods, as well as two HyVAE variants that only learn the local patterns or temporal dynamics of time series, respectively.

1. Introduction

Time series forecasting aims at learning the generation process of time series and uses previously observed samples to predict future values [1]. Accurate forecasting is essential and can help with the success of many applications/businesses. For example, an electricity company can design effective energy policies in advance by predicting the future energy consumption [2]; a corporation can minimize its investment risk if the future stock prices are accurately predicted [3].

Time series forecasting has been studied in the literature for decades, but to date, it remains a challenging and active research problem due to the complexity of time series. Classical time series forecasting methods, including autoregressive models (AR), moving average models (MA), and autoregressive integrated moving average models (ARIMA) [4], predict future values by assuming they have linear relationships with observed values; however, this simplification normally leads to unsatisfactory results for complex real-world time series. With the booming of deep learning techniques, deep neural networks (DNN) are widely used to tackle time series forecasting problems. Unlike classical models, DNNs are flexible non-linear models that can capture the temporal information of time series for forecasting [5]. Convolutional neural networks (CNN) [6] and recurrent neural networks (RNN) [7] are two types of DNN widely adopted for time series forecasting. CNN captures salient local patterns of short time series

subsequences/segments (e.g., seasonality [8] and trend [9]), while RNN learns long-term or mid-term temporal dynamics/dependencies of the entire time series [10]. In fact, many works capture both types of temporal information by proposing hybrid DNN models and obtaining more accurate forecasting results [9,11]. For example, the researchers [12] adopt a hybrid neural network, which stacks CNN with RNN, for DNA sequence prediction. Specifically, CNN can capture short and recurring sequence motifs, which represent biological function units in a DNA sequence. RNN, i.e., long short-term memory (LSTM) [10], is stacked with the output of CNN to learn the spatial arrangement of these motifs.

However, these DNN-based models cannot capture temporal information from time series with high accuracy since they are sensitive to small perturbations on time series [13]. Recent works refer to variational autoencoder (VAE) [14], which is a type of deep generative model, to learn representations of time series as latent random variables and obtain improved results [15]. Compared with directly fitting the exact values of time series, the latent random variables learned by VAE represent the generation process of time series and thus can more accurately capture essential temporal information of time series [16]. Based on this, existing methods learn either local seasonal-trend patterns [17] or temporal dynamics [15]; but to date, there is no VAE model that can jointly capture both information for time series forecasting.

* Corresponding author.

E-mail addresses: b.cai@deakin.edu.au (B. Cai), shuiqiao.yang@unsw.edu.au (S. Yang), gaolx@sdas.org (L. Gao), yong.xiang@deakin.edu.au (Y. Xiang).

In this paper, we bridge this gap by proposing a novel hybrid variational autoencoder (HyVAE) method for time series forecasting. HyVAE follows the variational inference [14] to jointly learn local patterns and temporal dynamics of time series. To achieve this goal, HyVAE is designed based on two objectives: (1) capturing local patterns by encoding time series subsequences into latent representations; and (2) learning temporal dynamics through the temporal dependencies among latent representations of different time series subsequences. HyVAE integrates the two objectives following the variational inference. Extensive experiments conducted on four real-world time series datasets show that HyVAE can improve the time series forecasting accuracy over strong counterpart methods. The contributions of this paper are summarized as follows:

- We propose a novel hybrid variational autoencoder (HyVAE) for time series forecasting. HyVAE derives an objective following variational inference to integrate the learning of local patterns and temporal dynamics of time series, thereby improving the accuracy of forecasting.
- We conduct comprehensive experiments on four real-world datasets to demonstrate the effectiveness of the proposed HyVAE method, and the results show that HyVAE achieves better forecasting accuracy than strong counterpart methods.

The rest of this paper is organized as follows. The related works are reviewed in Section 2. The preliminary knowledge is introduced in Section 3. The proposed method is detailed in Section 4, and is evaluated in Section 5. The paper is summarized in Section 6.

2. Related work

In this section, we briefly review time series forecasting methods and VAE-related forecasting approaches.

2.1. Time series forecasting

Classical auto-regressive model (AR) predicts by the linear aggregation of past time series values and a stochastic term (e.g., white noise). ARIMA extends AR to non-stationary time series by incorporating moving average (MA) and differencing. Other statistical models, such as linear regression [18] and support vector regression [19], enhance the model capacity but still have limited expressiveness. DNNs are flexible non-linear models and have been widely used for time series forecasting in recent years. Specifically, RNNs memorize historical information with feedback loops and can conveniently learn the temporal dynamics of time series. Long short-memory network (LSTM) [10] is a typical RNN that alleviates gradient vanishing with forget gates, and that enables the learning of long-term temporal dynamics for time series. Other types of RNN, e.g., GRU [20], and Informer [1], which uses the attention mechanism [21], are also used to improve the effectiveness of different forecasting scenarios. In addition, CNNs [22] are further used to capture local patterns of time series (such as seasonality [8] and trends [9]). Many works stack CNN and RNN to learn both the local patterns and the temporal dynamics for challenging forecasting problems; for example, combining multi-layer one-dimensional CNNs with bi-directional LSTM for air quality forecasting [9] and DNA sequence forecasting [12]; integrating a Savitzky–Golay filter (to avoid noise) and a stacked TCN-LSTM for traffic forecasting [11]. In addition, transformers [23] and GNNs are also adopted for forecasting. Transformer [23] better captures the complicated period dynamics [24] or resolves over-stationarization [25] with the flexible multi-head attention module; however, recent work [26] also suggests that its position encoding incurs losses of temporal information. Meanwhile, GCN is applied on specific graph representations, e.g., the time-conditioned graph structures in Z-GCNETs [27] (by introducing time-aware zigzag persistence), for robust time series forecasting.

2.2. Variational autoencoder-based forecasting

Variational autoencoder (VAE) [14] is a powerful deep generative model that encodes the input data as latent random variables, instead of deterministic values. To enhance the flexibility of VAE (learns independent latent random variables), follow-up methods introduce extra dependencies among the latent random variables. For example, ladder variational autoencoder [28] specifies a top-down hierarchical dependency among the latent random variables, fully-connected variational autoencoder [29] includes all possible dependencies among variables, and graph variational autoencoder [30] automatically learns an acyclic dependency graph. Due to the high flexibility, it is introduced to time series forecasting [31]. To improve the performance of the vanilla VAE, VRNN [15] introduces an RNN as the backbone to capture the long-term temporal dynamics of time series. LaST [17] develops disentangled VAE to learn dissociated seasonality and trend patterns of time series for forecasting. The proposed HyVAE is different from existing methods as it integrates the learning of both local patterns and the temporal dynamic for time series forecasting.

3. Preliminaries

In this section, we first define the problem and then introduce the preliminary knowledge of VAE.

3.1. Notation and problem statement

A scalar is denoted as a lowercase character, a vector is denoted as a bold lowercase character, and a matrix is denoted as an uppercase character. A time series is denoted as $s = \{s_1, s_2, \dots, s_m, s_{m+1}, \dots, s_{m+n}\}$, the time series forecasting problem is defined as determining $\{s_{m+1}, \dots, s_{m+n}\}$ with known $\{s_1, s_2, \dots, s_m\}$, where n is the step of forecasting. For the convenience, we denote $y = \{s_{m+1}, \dots, s_{m+n}\}$, and the forecasting problem can be formulated as $\hat{y} = f(s_1, s_2, \dots, s_m)$, where \hat{y} is the predicted values for y . The error of forecasting is measured as follows:

$$Errr(y, \hat{y}) = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2, \quad (1)$$

Time series subsequence is denoted as $x_l = \{s_t, \dots, s_{t+l-1}\}$, where l is the length. Time series subsequence contains contextual information that expresses local patterns [9], and thus we use subsequences in the forecasting task. Following [32], we obtain a series of l length subsequences from time series, using a sliding window. Time series represented by subsequences is denoted as $\{x^1, \dots, x^T\}$, where $T = m-l+1$ is the number of its subsequences. Thus, the forecasting problem becomes $\hat{y} = f(x^{\leq T}) = f(x^1, \dots, x^T)$ (see Table 1).

3.2. Variational autoencoder

Variational autoencoder (VAE) [14] is an unsupervised generative learning model that learns the latent representation of the input data as random variables. Similar to the conventional autoencoder [33], VAE has an encoding process that encodes the input into latent representations, and a decoding process that reconstructs the original input with the learned representations. We show the process of VAE in Fig. 1.

VAE learns the generative model as $p(x, z) = p(x|z)p(z)$, where x is the input data and z is its latent representations. The prior of z , $p(z)$, is normally defined as a multivariate Gaussian distribution, i.e., $z \sim \mathcal{N}(0, I)$; we denote that as $p(z) = \mathcal{N}(z|0, I)$ for convenience. The posterior $p(z|x)$ normally can be an arbitrary non-linear non-Gaussian distribution and thus is intractable. To resolve that, VAE approximates the posterior with $q(z|x) = \mathcal{N}(z|\mu(x), \sigma(x))$, where mean and variance are determined by x . Then, VAE defines the learning problem as the maximum likelihood estimation of $\log p(x)$, which can be formulated as:

$$\log p(x) = KL(q(z|x)||p(z|x)) + \ell, \quad (2)$$

Table 1
Summary of notations.

Notation	Description
s	time series
y	ground truth future values, $\{s_{m+1}, \dots, s_{m+n}\}$
\hat{y}	predicted future values
x	time series subsequence, $\{s_1, \dots, s_{t+l-1}\}$
z	latent representations learnt with VAE
h	hidden states learned with RNN (i.e., GRU)
L	<i>ladder size</i> for subsequence encoding
$\mathcal{N}(\mu, \sigma)$	Gaussian distribution
$KL(q(x) p(x))$	KL divergence from $q(x)$ to $p(x)$




Fig. 1. The framework of VAE. VAE encodes the input (x) into the latent random variables (as Gaussian distributions). Then, z is sampled from the distribution of latent random variables to reconstruct the input (\hat{x}).

where the first term is the KL divergence between the approximated posterior and the true posterior. Specifically, the KL divergence of two distributions $q(x)$ and $p(x)$ measures their similarity and is defined as:

$$KL(q(x)||p(x)) = \int q(x) \frac{q(x)}{p(x)} = \mathbb{E}_{q(x)} \frac{q(x)}{p(x)}. \quad (3)$$

In Eq. (2), since $p(z|x)$ is intractable and KL divergence is non-negative, maximizing $\log p(x)$ is achieved by maximizing ℓ , which is the evidence lower bound (ELBO) of $\log p(x)$ defined as follows:

$$\ell = E_{q(z|x)} \log p(x|z) - KL(q(z|x)||p(z)), \quad (4)$$

The first term in ℓ maximizes the conditional probability of x given the latent representation z and can be seen as the reconstruction loss, while the second term minimizes the difference between the prior and the approximated posterior.

4. The proposed method

In this section, we first provide an overview of the proposed hybrid variational autoencoder (HyVAE) method and then elaborate on its details.

4.1. Overview of HyVAE

We propose a novel generative hybrid variational autoencoder (HyVAE) model for time series forecasting, inspired by existing hybrid deterministic deep neural models. HyVAE jointly learns the local patterns from time series subsequences and the temporal dynamics among time series subsequences. To achieve that, HyVAE is derived based on variational inference to integrate two processes: (1) the encoding of time series subsequences, which captures local patterns; and (2) the encoding of entire time series, which learns temporal dynamics among time series subsequences. In the following content, we separately detail the encoding of time series subsequences and the encoding of the entire time series, respectively, and then explain the integration of these two processes for time series forecasting.

4.2. Encoding of time series subsequence

As discussed in Section 1, many existing models have shown that learning the local patterns can effectively improve time series forecasting [9]. To capture the flexible local patterns, we encode time series subsequences as latent random variables, rather than deterministic values.

An intuitive choice of the encoder is the conventional VAE, which maps a time series subsequence (x^t) into latent random variables (z^t)




Fig. 2. The encoder (inference process) and the decoder (generative process) of subsequence encoding.

to learn the local patterns. In VAE, z^t are assumed to be independent variables as $p(z^t) = \mathcal{N}(z^t|\mu^t, \sigma^t)$, where σ^t is the diagonal covariance matrix. However, such a simplified assumption may lead to inaccurate local pattern learning since essential causal information is neglected. In particular, a sample in a time series is always highly affected by its previous samples (e.g., autoregressive). Therefore, we further enforce causal dependency among variables in z^t to capture causal information within local patterns.

We separate latent random variables in z^t as L ladders (groups) (L is the *ladder size*) [28], i.e., $\{z_1^t, \dots, z_L^t\}$, and the groups have a sequential causal dependency (from 1 to L). We illustrate the encoding and decoding process of subsequence encoding in Fig. 2, in which the top row is the encoding process and the bottom row is the decoding process. For the convenience of implementation, we adopt the same causal dependency among the latent random variables ($z_1^t \rightarrow \dots \rightarrow z_L^t$) in the encoding and decoding processes. Based on this, the prior distribution of z^t can be factorized as:

$$p(z^t) = p(z_L^t) \prod_{i=1}^{L-1} p(z_i^t|z_{i+1}^t), \quad (5)$$

$$p(z_i^t|z_{i+1}^t) = \mathcal{N}(z_i^t|\mu_i^t(z_{i+1}^t), \sigma_i^t(z_{i+1}^t)),$$

where $\{\mu(\star), \sigma(\star)\} = \varphi(\star)$ and we implement $\varphi(\star)$ as a multilayer perceptron (MLP). By changing the size of dependency (L , the *ladder size*), we can regulate how well causal information is preserved, and no causal information when $L = 1$ (i.e., all latent random variables are independent). Based on this, the generative model of subsequence encoding can further be factorized as follows:

$$p(x^t, z^t) = p(x^t|z_1^t) p(z_L^t) \prod_{i=1}^{L-1} p(z_i^t|z_{i+1}^t), \quad (6)$$

where $p(x^t|z_1^t) = \mathcal{N}(x^t|\mu^t(z_1^t), \sigma^t(z_1^t))$. This causal dependency ensures the latent random variables have sufficient flexibility to model the complex local patterns of subsequences. Since the posterior $p(z^t|x^t)$ is intractable, $q(z^t|x^t)$ is used as an approximation. Meanwhile, to avoid $\{z_1^t, \dots, z_L^t\}$ converging to arbitrary variables, they all depend on x^t in the inference model similar to [28] as follows:

$$q(z^t|x^t) = q(z_L^t|x^t) \prod_{i=1}^{L-1} q(z_i^t|z_{i+1}^t, x^t), \quad (7)$$

$$q(z_L^t|x^t) = \mathcal{N}(z_L^t|\mu_L^t(x^t), \sigma_L^t(x^t)),$$

$$q(z_i^t|z_{i+1}^t, x^t) = \mathcal{N}(z_i^t|\mu_i^t(z_{i+1}^t, x^t), \sigma_i^t(z_{i+1}^t, x^t)),$$

where $\{\mu(\star, \star), \sigma(\star, \star)\} = \varphi([\star; \star])$ and $[\cdot; \cdot]$ is the concatenation operation.

4.3. Encoding of entire time series

From the global perspective, we encode all time series subsequences $\{x^1, \dots, x^T\}$ as $\{z^1, \dots, z^T\}$ to learn the temporal dynamics of entire time series. Since time series subsequences are normally not independent across different time stamps, we first impose a temporal dependency for consecutive subsequences (e.g., $p(z^t, z^{t-1}) = p(z^t|z^{t-1})p(z^{t-1})$). In addition, we capture long-term temporal dependency with other




Fig. 3. The illustration of HyVAE; (a) shows the prior defined by Eq. (13); (b) is the recurrent updating of GRU hidden states in Eq. (9); (c) shows the inference operation in Eq. (14); and (d) represents the generation operation in Eq. (15).

subsequences by hidden states of a recurrent neural network, i.e., gated recurrent unit (GRU) [20]. Therefore, we have the following derivation: $p(z^t|z^{<t})$ can be derived as follows:

$$\begin{aligned} p(z^t|z^{<t}) &= p(z^t|z^{t-1}, h^{t-1}), \\ p(z^t|z^{t-1}, h^{t-1}) &= \mathcal{N}(z^t|\mu_i^t(z^{t-1}, h^{t-1}), \sigma_i^t(z^{t-1}, h^{t-1})), \end{aligned} \quad (8)$$

where h is the hidden state and is obtained by:

$$h^t = \text{GRU}(h^{t-1}, x^t). \quad (9)$$

GRU(*) is the calculation of hidden states in a GRU unit. GRU adopts gates and memory cells and alleviates the gradient vanishing problem while being easier to train than LSTM due to fewer gates used. The structure of GRU is formulated as follows:

$$\begin{aligned} r^t &= \sigma(W_r[h^{t-1}; x^t]), \\ \zeta_t &= \sigma(W_\zeta[h^{t-1}; x^t]), \\ \tilde{h}^t &= \tanh(W_{\tilde{h}}[r^t \circ h^{t-1}; x^t]), \\ h^t &= (1 - \zeta_t) \circ h^{t-1} + \zeta_t \circ \tilde{h}^t, \end{aligned} \quad (10)$$

where \circ is the element-wise product. Specifically, r^t and ζ_t are the reset gate vector and update gate vector, which decides how much past information needs to be forgotten/preserved, respectively. Meanwhile, \tilde{h}^t is the candidate activation vector that memorizes the past information, and h^t is obtained as the balanced sum of the short (h^{t-1}) memory and the long (\tilde{h}^t) memory.

For the generative process $p(x^{\leq T}, z^{\leq T})$, we explicitly simplify $p(x^t|z^{\leq t})$ as $p(x^t|z^t)$ to ensure the local pattern of x^t is mainly preserved in z^t ; this simplification also can largely reduce the complexity of the reconstruction/decoding process. Based on (8), the generation model can be factorized as follows:

$$p(x^{\leq T}, z^{\leq T}) = \prod_{i=1}^T p(x^i|z^i, x^{<i})p(z^i|x^{<i}, z^{i-1}), \quad (11)$$

$p(x^t|z^t, x^{<t})$ also can be denoted as $p(x^t|z^t, h^{t-1})$, due to the recursive nature of GRU, which requires h^{t-1} being obtained by the recursive calculation with $x^{<t}$.

Similarly, we derive the inference model as:

$$\begin{aligned} q(z^t|x^{\leq t}, z^{t-1}) &= \mathcal{N}(z^t|\mu_i^t(x^{\leq t}, z^{t-1}), \sigma_i^t(x^{\leq t}, z^{t-1})) \\ &= \mathcal{N}(z^t|\mu_i^t(h^{t-1}, x^t, z^{t-1}), \sigma_i^t(h^{t-1}, x^t, z^{t-1})). \end{aligned} \quad (12)$$

The above approximated posterior of z^t captures the long-term dynamics carried by $x^{<t}$ (h^{t-1}), the neighboring dependency with z^{t-1} , and the corresponding subsequence x^t .

4.4. Integration and joint learning

Based on the encoding of a subsequence and the encoding of the entire time series (represented as subsequences) discussed above, we now integrate them into a HyVAE model, which can jointly learn the local patterns and temporal dynamics for time series forecasting. The jointly learned latent random variables for both time series subsequences and the entire time series are denoted as $\{(z_L^1, \dots, z_1^1), \dots, (z_L^T, \dots, z_1^T)\}$, with respect to time series $\{x^1, \dots, x^T\}$, and the encoding process is

illustrated in Fig. 3. By combining the prior of subsequence encoding in Eq. (5) and the prior of entire time series encoding in Eq. (8), we obtain the prior of HyVAE, which is factorized as follows:

$$\begin{aligned} p(z^t|z^{t-1}, h^{t-1}) &= p(z_L^t|z_1^{t-1}, h^{t-1}) \prod_{i=1}^L p(z_i^t|z_{i+1}^{t-1}), \\ p(z_L^t|z_1^{t-1}, h^{t-1}) &= \mathcal{N}(z_L^t|\mu^t(z_1^{t-1}, h^{t-1}), \sigma^t(z_1^{t-1}, h^{t-1})), \\ p(z_i^t|z_{i+1}^{t-1}) &= \mathcal{N}(z_i^t|\mu_i^t(z_{i+1}^{t-1}), \sigma_i^t(z_{i+1}^{t-1})). \end{aligned} \quad (13)$$

As shown in Fig. 3(a), the prior of HyVAE integrates the long-term temporal dynamics by affecting the first latent random variable of each subsequence (e.g., z_L^t) with the hidden states (e.g., h^{t-1} , generated by GRU) of its precedent subsequence. Meanwhile, h^t is obtained by the recurrence process with GRU as shown in Fig. 3(b). We then obtain the inference model of HyVAE by integrating Eqs. (7) and (12) as follows (Fig. 3(c)):

$$\begin{aligned} q(z^t|x^{\leq t}, z^{t-1}) &= q(z_L^t|x^{\leq t}, z_1^{t-1}) \prod_{i=1}^{L-1} q(z_i^t|z_{i+1}^{t-1}, x^t), \\ q(z_L^t|x^{\leq t}, z_1^{t-1}) &= \mathcal{N}(z_L^t|\mu^t(z_1^{t-1}, x^{\leq t}), \sigma^t(z_1^{t-1}, x^{\leq t})), \\ q(z_i^t|z_{i+1}^{t-1}, x^t) &= \mathcal{N}(z_i^t|\mu_i^t(z_{i+1}^{t-1}, x^t), \sigma_i^t(z_{i+1}^{t-1}, x^t)). \end{aligned} \quad (14)$$

Similarly, for x^t , the above encoding process also includes the temporal dynamics carried by h^{t-1} during the encoding of z_L^t , while the rest latent random variables of z^t only learn from x^t . Then, as shown in Fig. 3(d), the generation model of HyVAE is obtained by combining Eqs. (6) and (11) as follows:

$$\begin{aligned} p(x^t|z^t, x^{<t}) &= p(x^t|z_1^t, h^{t-1}), \\ &= \mathcal{N}(x^t|\mu_i^t(z_1^t, h^{t-1}), \sigma_i^t(z_1^t, h^{t-1})). \end{aligned} \quad (15)$$

Following the derivative process of variational inference, with Eq. (13) to (15), HyVAE learns the latent representations by maximizing its ELBO defined as follows:

$$\begin{aligned} \ell_{enc} &= \sum_{t=1}^T \left\{ \mathbb{E}_{q(z_i^t|x^{\leq t}, z_1^{t-1})} \log p(x^t|h^{t-1}, z_1^t) \right. \\ &\quad - \sum_{i=1}^{L-1} KL(q(z_i^t|z_{i+1}^{t-1}, x^t) || p(z_i^t|z_{i+1}^{t-1})) \\ &\quad \left. - KL(q(z_L^t|x^{\leq t}, z_1^{t-1}) || p(z_L^t|x^{<t}, z_1^{t-1})) \right\}. \end{aligned} \quad (16)$$

The first term in ℓ_{enc} implies the reconstruction loss of HyVAE for each time series subsequence, i.e., between the input x^t and the \hat{x}^t reconstructed with (z_1^t, h^{t-1}) (see Fig. 3(d)). The second and third terms are regularization terms that enforce the encoded latent random variables to jointly capture the local patterns of individual subsequences and learn the temporal dynamics of the entire time series. The expectation of ℓ_{enc} is approximated by Monte Carlo estimation [34] and is estimated with the average of the ℓ_{enc} of each sample time series.

We use h^t and z^t for the final time series forecasting, i.e., $\hat{y} = \psi(h^t, z^t)$, where $\psi(\ast)$ is a single-layer fully-connected neural network. The forecasting loss is measured by Eq. (1):

$$\ell_{pred} = \text{Err}(y, \hat{y}). \quad (17)$$

Table 2
Statistics of the datasets.

Dataset	Train	Valid	Test	Description
Parking	2856	357	358	Car park occupancy
Stock	1081	135	136	NASDAQ stock index
Electricity	1120	140	140	Electricity load values
Sealevel	1120	140	140	Sea level pressure

Then, the overall loss minimizes the negative ELBO of HyVAE and the forecasting loss as follows:

$$\ell = -\ell_{enc} + \ell_{pred}. \quad (18)$$

In ℓ , ℓ_{enc} aims at learning representations that capture the latent distribution of time series, while ℓ_{pred} can be regarded as a regularization term that ensures the latent representations can provide insights for accurate forecasting. We perform ADMA [35] for the optimization and use the reparameterization trick [14] for the model training. For ℓ_{enc} , we adopt the warm-up scheme [28] during the implementation to avoid inactive latent random variables caused by the variational regularization.

5. Evaluation

In this section, we first introduce the real-world datasets used to evaluate the proposed method. Then, we explain the accuracy metrics for time series forecasting and briefly describe the counterpart methods. Finally, we analyze the results and compare HyVAE with counterpart methods regarding the effectiveness of time series forecasting. All the experiments are implemented with Python 3.7 and run on a Linux platform with a 2.6G CPU and 132G RAM.

5.1. Datasets

We select four datasets widely used for time series forecasting. Parking Birmingham dataset [36] is collected from car parks in Birmingham, which regularly records the total occupancy of all available parking spaces between October 4, 2016, and December 19, 2016. We down-sample the recording frequency to every 5 h and result in 3571 records. Another NASDAQ stock dataset [37] consists of stock prices of 104 corporations together with the overall NASDAQ100 index, which is collected from July 26, 2016, to December 22, 2016. We use the NASDAQ100 index for forecasting and down-sample the records every 30 min, which results in 1352 records. The other two datasets¹ record the electricity load values of Poland from the 1990s and monthly Darwin sea level pressures from 1882 to 1998, respectively; both datasets contain 1400 records. We preprocess each dataset with Min-Max normalization by:

$$s'_i = \frac{s_i - \min(s)}{\max(s) - \min(s)}. \quad (19)$$

Then, each dataset is split into a training set, a validation set and a test set by {80%, 10%, 10%}. The number of known time series values used for forecasting is fixed as 50 for all datasets. The statistics of the datasets are shown in Table 2.

5.2. Performance metric

We use three different metrics widely used for time series forecasting [9,15] in the evaluation, and they are mean square error (MSE), mean absolute error (MAE), and mean absolute percentage error (MAPE). MSE and MAE respectively measure the variance and

average of the residuals of the forecasting results to ground truth and are respectively defined as follows:

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2, MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|. \quad (20)$$

MAPE measures the proportion of forecasting deviation to the ground truth as follows:

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{\max(\epsilon, y_i)} \right|, \quad (21)$$

where ϵ is an arbitrarily small positive value to ensure the dividing is always legal.

5.3. Counterpart methods

We select three types of counterpart methods to compare with the proposed method, i.e., the classical statistical models, deterministic DNN-based methods, and VAE-based methods. The classical models include the widely used AR, ARIMA, and SVR. For deterministic DNN-based methods, we choose the LSTM and Informer and implement a stacked CNN and LSTM model (CNN+LSTM) following [9] for time series forecasting. For the VAE-based methods, other than the vanilla VAE, we adopt VRNN [15] and LaST [17]. We brief these methods as follows:

- **AR** forecasts with the weighted sum of past values. **ARIMA** incorporates moving average and differencing to AR for non-stationary time series.
- **SVR** [19] is based on the support vector machine (SVM) and the principle of structural risk minimization.
- **LSTM** [10] is an RNN model that can learn the long dynamics with its forget gates.
- **Informer** [23] uses multi-head attention with position encoding to learn the latent structure of time series for forecasting.
- **CNN+LSTM** [9] stacks CNN and LSTM for accurate air quality forecasting. CNN+LSTM includes three TCN layers and two bi-LSTM layers.
- **Vanilla VAE** [14] is the basic variational autoencoder that learns latent representations as independent Gaussian random variables.
- **VRNN** [15] extends VAE to be capable of learning temporal dynamics by introducing temporal dependency among the latent representations.
- **LaST** [17] adopts disentangled variational autoencoder to capture seasonality and trend, with auxiliary objectives to ensure dissociate representations.
- **GBT** [38] decouples the Transformer-based forecasting into an auto-regressive stage and a self-regression stage to overcome the severe over-fitting problem.
- **N-HITS** [39] incorporates hierarchical interpolation and multi-rate data sampling to address the volatility of predictions.

5.4. Experiment setup

In all the experiments, we use the validation sets to tune optimal parameters and use the test sets for forecasting accuracy measurement. For AR, we search the optimal number of lag (past time series values) from 1 to 10, and use the same strategy to search optimal p (the number of past observations) and q (the size of moving average window) for ARIMA, with the optimal differencing degree searched from 0 to 3. For SVR, we adopt the radial basis function (*RBF*) kernel for running, with its parameters C (regularization parameter) searched from {1, 10, 100, 1000} and γ (kernel coefficient) searched from {0.00005, 0.0005, 0.005, 0.05}.

For LSTM, Informer, CNN+LSTM, GBT, N-HITS, vanilla VAE, VRNN, LaST, and HyVAE, we search the optimal batch size from {32, 64, 128} and set the maximum iteration to be 100 epochs. The learning rate

¹ <https://research.cs.aalto.fi/aml/datasets.shtml>

Table 3
Time series forecasting results on the datasets, with the best displayed in bold.

Methods	Parking			Stock			Electricity			Sealevel		
	MSE ($\times 10^{-2}$)	MAE	MAPE	MSE ($\times 10^{-2}$)	MAE	MAPE	MSE ($\times 10^{-2}$)	MAE	MAPE	MSE ($\times 10^{-2}$)	MAE	MAPE
AR	1.043	0.085	0.474	1.876	0.114	0.146	1.093	0.086	0.336	1.481	0.095	0.187
ARIMA	0.637	0.066	0.289	0.880	0.080	0.100	0.743	0.056	0.249	1.613	0.107	0.204
SVR	1.077	0.082	0.288	0.606	0.075	0.092	0.563	0.057	0.254	1.003	0.079	0.161
LSTM	0.571	0.057	0.249	0.557	0.068	0.078	0.321	0.036	0.211	0.801	0.068	0.139
Informer	0.425	0.051	0.224	0.728	0.078	0.090	0.305	0.037	0.219	1.083	0.083	0.173
CNN+LSTM	0.397	0.046	0.200	0.254	0.043	0.049	0.149	0.023	0.210	0.667	0.062	0.127
GBT	0.400	0.046	0.200	0.420	0.056	0.061	0.117	0.017	0.178	0.760	0.072	0.138
N-HITS	0.401	0.469	0.197	0.127	0.027	0.031	0.336	0.042	0.229	0.687	0.064	0.129
Vanilla VAE	0.713	0.067	0.312	15.142	0.373	0.754	5.445	0.200	0.457	4.386	0.178	0.332
VRNN	0.454	0.056	0.226	0.190	0.039	0.042	0.199	0.029	0.195	0.665	0.066	0.125
LaST	0.366	0.043	0.191	0.119	0.026	0.029	0.116	0.018	0.164	0.674	0.064	0.128
HyVAE	0.133	0.028	0.144	0.087	0.021	0.023	0.097	0.015	0.143	0.623	0.060	0.123

Table 4
Multi-step forecasting results (MSE $\times 10^{-2}$) on the datasets, with the best displayed in bold.

Methods	Parking			Stock			Electricity			Sealevel		
	3	4	5	3	4	5	3	4	5	3	4	5
LSTM	0.712	0.696	0.743	0.669	0.778	1.388	0.340	0.396	0.386	0.938	1.052	1.122
Informer	0.662	0.638	0.649	0.933	1.163	1.297	0.307	0.409	0.450	1.105	1.177	1.056
CNN+LSTM	0.635	0.643	0.636	0.524	0.710	0.822	0.163	0.156	0.197	0.890	0.959	1.032
GBT	0.640	0.661	0.682	0.722	0.848	0.951	0.140	0.153	0.201	0.958	1.023	1.247
N-HITS	0.640	0.653	0.655	0.436	0.493	0.670	0.362	0.403	0.492	0.847	0.936	1.096
VRNN	0.614	0.700	0.760	0.557	0.653	0.895	0.232	0.437	0.458	0.870	1.257	1.322
LaST	0.442	0.533	0.574	0.425	0.463	0.617	0.136	0.138	0.189	0.819	0.892	1.021
HyVAE	0.446	0.466	0.502	0.367	0.365	0.431	0.123	0.137	0.151	0.787	0.858	0.931

is searched from $\{0.001, 0.01, 0.1\}$. The dimension of the LSTM/GRU hidden states and latent representations are searched from $\{8, 16, 32, 64, 128\}$, and the number of layers is no more than 3. For HyVAE, the *ladder size* and the *subsequence length* are searched from $\{2, 4, 6, 8, 10\}$ and $\{10, 20, 30, 40\}$, respectively. We run each method 50 times and report the average accuracy as the final results.

5.5. Main results

In this experiment, we compare the accuracy of HyVAE with counterpart methods, with respect to single-step forecasting and multi-step (3, 4, and 5 steps) forecasting, respectively, on the four datasets.

As shown in Table 3, HyVAE generally achieves the best performance among all methods on the four datasets. Notably, we observe on the Parking dataset, the MSE achieved by HyVAE (0.133×10^{-2}) is nearly three times smaller than that of the second-best performed LaST (0.366×10^{-2}). The least improvement over all the counterpart methods is shown in the Sealevel dataset, in which HyVAE reduces the MSE, MAE, and MAPE of VRNN (the second best) by 6.3%, 9.1% and 1.6%, respectively. When further considering the type of the counterpart methods, first, we see HyVAE achieves significant improvement over the classical AR, ARIMA, and SVR methods, by achieving nearly one magnitude smaller MSE on the Parking and Stock datasets. Second, compared with the deterministic DNN-based LSTM, Informer, CNN+LSTM, GBT, and N-HITS, HyVAE also shows significant improvement; especially on the Parking dataset, HyVAE achieves nearly two times smaller MSE, MAE, and MAPE than the best-performed deterministic neural model (CNN-LSTM). Although CNN+LSTM also considers both local patterns and temporal dynamics of time series and outperforms LSTM and Informer on all the datasets, HyVAE constantly being more effective and thus is better at capturing the complex structure of time series for forecasting. Third, we can see that HyVAE achieves more accurate forecasting results than other VAE-based methods that only learn part of the information of time series. That includes the vanilla VAE, which misses the temporal dynamics, VRNN, which only learns the temporal dynamics, and LaST for seasonality/trend patterns of time series. This

observation shows the effectiveness of HyVAE in learning both the local patterns and temporal dynamics for time series forecasting.

For multi-step forecasting, in Table 4, we show the MSE of LSTM, Informer, CNN+LSTM, GBT, N-HITS, VRNN, LaST, and HyVAE; AR, ARIMA, SVR, and vanilla VAE are excluded due to low performance. The results show that HyVAE achieves more accurate forecasting results than compared counterpart methods. Although generally, the forecasting accuracy decreases with larger forecasting steps, except for the Parking dataset, the forecasting accuracy of HyVAE decreases much slower than the compared counterpart methods since it captures more informative patterns of time series. For example, from 3-step forecasting to 5-step forecasting in the Electricity dataset, the MSE of HyVAE only decreases by 0.028, while LSTM, Informer, CNN+LSTM, GBT, N-HITS, VRNN, and LaST decrease by 0.046, 0.143, 0.034, 0.061, 0.130, 0.226 and 0.063, respectively. Meanwhile, CNN+LSTM and HyVAE in most cases produce more accurate forecasting results than other deterministic DNN-based methods and VAE-based methods, respectively, and that again supports the effectiveness of learning both local patterns and temporal dynamics for time series forecasting.

In addition, we compare the variational inference-based methods (VRNN, LaST, and HyVAE) on their performance on probabilistic forecasting, and the results are measured by continuous ranked probability score (CRPS). With F denoting the cumulative distribution function of the forecasts distribution, CRPS is defined as $CRPS(F, x) = \int_{-\infty}^{\infty} (F(y) - \mathbb{1}(y - x)^2) dy$, where $\mathbb{1}$ is the Heaviside step function. CRPS measures the similarity of the forecasted distribution with the true prediction and is minimized when they are identical. The results in Table 5 show that on all four datasets, HyVAE achieves more accurate probabilistic forecasting performance than VRNN, which only learns the global temporal dynamics, and LaST, which aims to capture specific temporal patterns (trend and seasonality).

5.6. Ablation analysis

We conduct an ablation analysis to further understand the effectiveness of learning both the local patterns and the temporal dynamics in HyVAE. To do that, we implement two variants of HyVAE by removing

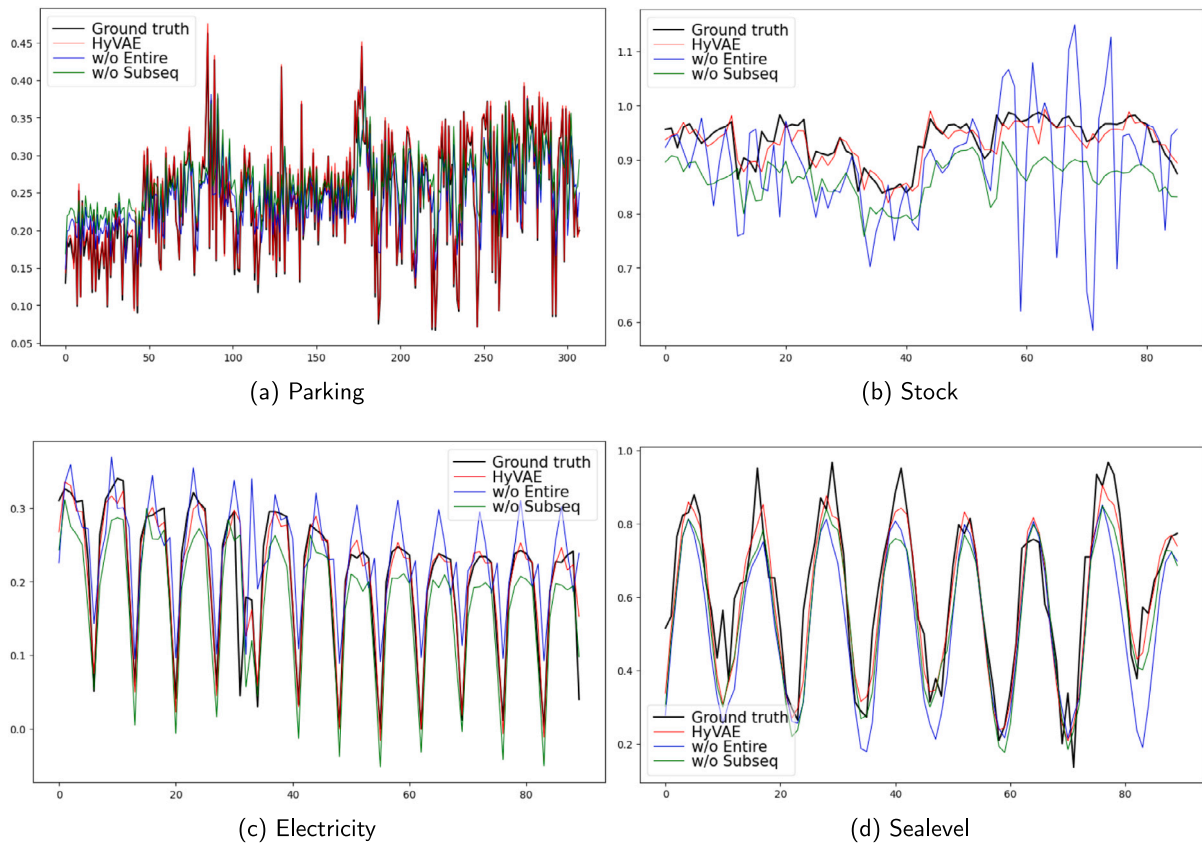


Fig. 4. Forecasting results of HyVAE, w/o Entire and w/o Subseq on the datasets.

Table 5

Probabilistic forecasting results (measured by CRPS) on the datasets, with the best displayed in bold.

Methods	Parking	Stock	Electricity	Sealevel
VRNN	1.220	0.231	0.129	0.574
LaST	0.838	0.172	0.060	0.389
HyVAE	0.649	0.138	0.052	0.306

Table 6

Ablation analysis of HyVAE ($MSE \times 10^{-2}$), with the best displayed in bold.

Methods	Parking	Stock	Electricity	Sealevel
w/o Entire	0.410	0.980	0.503	1.742
w/o Subseq	0.389	0.513	0.160	1.088
HyVAE	0.133	0.087	0.116	0.623

the learning of one type of information, respectively; that is, w/o Subseq that excludes the learning of local patterns from subsequences, and w/o Entire that does not learn the temporal dynamics of the entire time series. The parameters of w/o Subseq and w/o Entire are tuned with the validation set the same as HyVAE, and we show the results of time series forecasting measured by MSE in Table 6.

In Table 6, HyVAE that learns both information achieves higher forecasting accuracy than the two variants. Specifically, the largest improvement of HyVAE towards the variants is shown in the Stock dataset (0.087), i.e., around six times smaller than that of w/o Subseq (0.513, second best). The smallest improvement appears in the Sealevel dataset, but the MSE of HyVAE is still around two times smaller than that of the second-best performed w/o Subseq (0.623 to 1.088). Meanwhile, it is interesting to observe that w/o Entire, which misses the temporal dynamics, constantly performs worse than w/o Subseq, which misses local patterns, on the four datasets.

We further show the forecasting results of HyVAE and the variants against the ground truth in Fig. 4. For w/o Entire, since no temporal dynamics is learned, it cannot properly capture the global trend of time series; especially for the Stock dataset (Fig. 4(b)), it misunderstands the steady curves between the 55 and 75 timestamps as sharp spikes. As for the results of the Electricity dataset and the Sealevel dataset shown in Fig. 4(c–d), w/o Entire only emphasizes recurring local patterns but misses their differences at different timestamps. Meanwhile, w/o Subseq can better express the temporal dynamics than w/o Entire, as clearly shown in Fig. 4(b); however, it fails to properly capture local details. By combining the strengths of w/o Entire and w/o Subseq, HyVAE achieves the best forecasting results that are quite close to the ground truth.

Based on the above analysis, we summarize that the improved forecasting accuracy of HyVAE is due to its joint learning of local patterns and global dynamics of time series. Compared with only learning the global dynamics (w/o Subseq), HyVAE can more accurately capture the local details, e.g., the flat peaks in Fig. 4(c), thus significantly reducing the forecasting error in detail-rich regions on time series. While compared with only capturing the local patterns (w/o Entire), HyVAE improves the forecasting of global trends and dynamics, leading to an overall reduction of forecasting errors on most time series samples.

5.7. Parameter analysis

In this experiment, we analyze the impact of three parameters of HyVAE, i.e., the *ladder size*, the *subsequence length*, and the *embedding size*, on its performance. Specifically, the *ladder size* determines the causal information during subsequence encoding, and we vary the *ladder size* from 0 to 10, where 0 means HyVAE learns no causal information of subsequences (see Fig. 2). The *subsequence length* balances the local patterns and the temporal dynamics, i.e., HyVAE is degraded

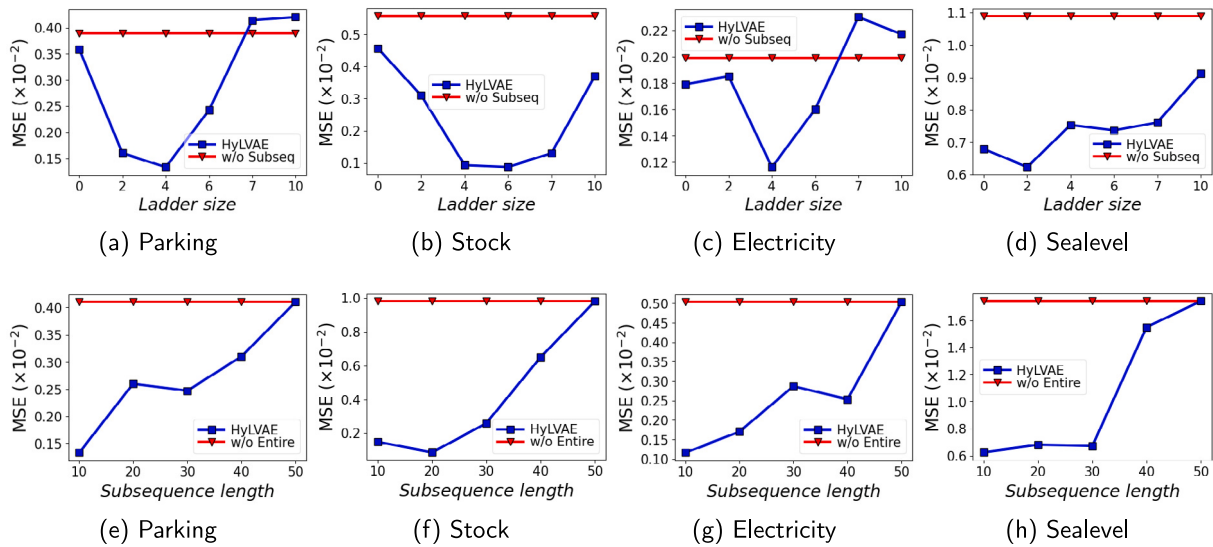


Fig. 5. Parameter analysis of HyVAE with respect to *ladder size* and *subsequence length*.




Fig. 6. Parameter analysis of HyVAE with respect to *embedding size*.

to w/o Subseq or w/o Entire if *subsequence length* is 0 (no subsequence) or the maximum (50, the subsequence becomes the entire time series), respectively. The results measured by MSE are shown in Fig. 5. For the *ladder size* shown in Fig. 5(a–d), the forecasting accuracy significantly decreases when *ladder size* is too small or too large. The optimal ladder sizes are relatively small (2 for the Sealevel dataset, 4 for the Parking and Stock datasets, and 6 for the Electricity dataset). Meanwhile, we see that when the *ladder size* equals 0, HyVAE still outperforms w/o Subseq on all the datasets. The results of *subsequence length* are shown in Fig. 5(c–d), in which we see that HyVAE prefers short subsequences to obtain optimal forecasting results, i.e., *subsequence length* is 10 (Parking, Electricity, and Sealevel datasets) or 20 (Stock dataset). The reason is that if the subsequences are too long, temporal dynamics can hardly be preserved. Not surprisingly, HyVAE that learns temporal dynamics with different *subsequence length* achieves better forecasting accuracy than w/o Entire, which does not learn temporal dynamics at all.

We then run HyVAE with varying *embedding size* {8, 16, 32, 64, 128}, which determines the dimension of latent representation and the dimension of hidden states in neural networks, and the results are shown in Fig. 6. On all the datasets, accuracy measured by MSE, MAE, and MAPE has similar trends. First, the forecasting accuracy is low with small *embedding size*, mainly because the small size of latent random variables cannot properly capture the complex non-linear processes of time series. When the *embedding size* becomes too large (128), the accuracy decreases due to over-fitting. The results show that HyVAE can obtain optimal forecasting results with relatively small *embedding size*.

6. Conclusion

This paper proposes a novel hybrid variational autoencoder (HyVAE) model for time series forecasting. HyVAE integrates the learning of local patterns and temporal dynamics into a variational autoencoder. Through comprehensive evaluation on four real-world datasets, we show that HyVAE achieves better time series forecasting accuracy than various counterpart methods, including a deterministic DNN-based method (CNN+LSTM) that also learns both information of time series. Moreover, the ablation analyses demonstrate that by jointly learning local patterns and temporal dynamics, HyVAE outperforms its two variants which only learn local patterns and temporal dynamics from time series, respectively.

CRedit authorship contribution statement

Borui Cai: Conceptualization, Methodology, Writing – original draft, Writing – review & editing. **Shuiqiao Yang:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing. **Longxiang Gao:** Conceptualization, Writing – original draft, Writing – review & editing. **Yong Xiang:** Conceptualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

This work was supported in part by the Australian Research Council under Grant LP190100594.

References

- [1] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, W. Zhang, Informer: Beyond efficient transformer for long sequence time-series forecasting, in: AAAI Conference on Artificial Intelligence, Vol. 35, No. 12, 2021, pp. 11106–11115.
- [2] K. Wang, C. Xu, Y. Zhang, S. Guo, A.Y. Zomaya, Robust big data analytics for electricity price forecasting in the smart grid, *IEEE Trans. Big Data* 5 (1) (2017) 34–45.
- [3] A.W. Li, G.S. Bastos, Stock market forecasting using deep learning and technical analysis: A systematic review, *IEEE Access* 8 (2020) 185232–185242.
- [4] P.R. Junior, F.L.R. Salomon, E. de Oliveira Pamplona, et al., ARIMA: An applied time series forecasting model for the Bovespa stock index, *Appl. Math.* 5 (21) (2014) 3383.
- [5] B. Lim, S. Zohren, Time-series forecasting with deep learning: a survey, *Phil. Trans. R. Soc. A* 379 (2194) (2021) 20200209.
- [6] Z.D. Akşehir, E. Kiliç, How to handle data imbalance and feature selection problems in CNN-based stock price forecasting, *IEEE Access* 10 (2022) 31297–31305.
- [7] R. He, Y. Liu, Y. Xiao, X. Lu, S. Zhang, Deep spatio-temporal 3D densenet with multiscale ConvLSTM-Resnet network for citywide traffic flow forecasting, *Knowl.-Based Syst.* 250 (2022) 109054.
- [8] S. Liu, H. Ji, M.C. Wang, Nonpooling convolutional neural network forecasting for seasonal time series with trends, *IEEE Trans. Neural Netw. Learn. Syst.* 31 (8) (2019) 2879–2888.
- [9] S. Du, T. Li, Y. Yang, S.-J. Horng, Deep air quality forecasting using hybrid deep learning framework, *IEEE Trans. Knowl. Data Eng.* 33 (6) (2019) 2412–2424.
- [10] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (8) (1997) 1735–1780.
- [11] J. Bi, X. Zhang, H. Yuan, J. Zhang, M. Zhou, A hybrid prediction method for realistic network traffic with temporal convolutional network and LSTM, *IEEE Trans. Autom. Sci. Eng.* 19 (3) (2021) 1869–1879.
- [12] D. Quang, X. Xie, DanQ: a hybrid convolutional and recurrent deep neural network for quantifying the function of DNA sequences, *Nucleic Acids Res.* 44 (11) (2016) e107.
- [13] J.E. Van Engelen, H.H. Hoos, A survey on semi-supervised learning, *Mach. Learn.* 109 (2) (2020) 373–440.
- [14] D.P. Kingma, M. Welling, Auto-encoding variational Bayes, in: *International Conference on Learning Representations*, 2014.
- [15] U. Ullah, Z. Xu, H. Wang, S. Menzel, B. Sendhoff, T. Bäck, Exploring clinical time series forecasting with meta-features in variational recurrent models, in: *International Joint Conference on Neural Networks*, IEEE, 2020, pp. 1–9.
- [16] W. Chen, L. Tian, B. Chen, L. Dai, Z. Duan, M. Zhou, Deep variational graph convolutional recurrent network for multivariate time series anomaly detection, in: *International Conference on Machine Learning*, PMLR, 2022, pp. 3621–3633.
- [17] Z. Wang, X. Xu, W. Zhang, G. Trajcevski, T. Zhong, F. Zhou, Learning latent seasonal-trend representations for time series forecasting, in: *Advances in Neural Information Processing Systems*, 2022.
- [18] J.F. de Oliveira, E.G. Silva, P.S. de Mattos Neto, A hybrid system based on dynamic selection for time series forecasting, *IEEE Trans. Neural Netw. Learn. Syst.* 33 (8) (2021) 3251–3263.
- [19] U. Thissen, R. Van Brakel, A. De Weijer, W. Melssen, L. Buydens, Using support vector machines for time series prediction, *Chemometr. Intell. Lab. Syst.* 69 (1–2) (2003) 35–49.
- [20] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using RNN encoder–decoder for statistical machine translation, in: *Conference on Empirical Methods in Natural Language Processing*, 2014, pp. 1724–1734.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: *Advances in Neural Information Processing Systems*, 2017.
- [22] R. Sen, H.-F. Yu, I.S. Dhillon, Think globally, act locally: A deep neural network approach to high-dimensional time series forecasting, in: *Neural Information Processing Systems*, 2019.
- [23] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, W. Zhang, Informer: Beyond efficient transformer for long sequence time-series forecasting, in: AAAI Conference on Artificial Intelligence, 2021.
- [24] W. Chen, W. Wang, B. Peng, Q. Wen, T. Zhou, L. Sun, Learning to rotate: Quaternion transformer for complicated periodical time series forecasting, in: *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022, pp. 146–156.
- [25] Y. Liu, H. Wu, J. Wang, M. Long, Non-stationary transformers: Exploring the stationarity in time series forecasting, in: *Advances in Neural Information Processing Systems*, Vol. 35, 2022, pp. 9881–9893.
- [26] A. Zeng, M. Chen, L. Zhang, Q. Xu, Are transformers effective for time series forecasting? in: AAAI Conference on Artificial Intelligence, Vol. 37, No. 9, 2023, pp. 11121–11128.
- [27] Y. Chen, I. Segovia, Y.R. Gel, Z-GCNets: Time zigzags at graph convolutional networks for time series forecasting, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 1684–1694.
- [28] C.K. Sønderby, T. Raiko, L. Maaløe, S.K. Sønderby, O. Winther, Ladder variational autoencoders, in: *Advances in Neural Information Processing Systems*, Vol. 29, 2016, pp. 3745–3753.
- [29] D.P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, M. Welling, Improved variational inference with inverse autoregressive flow, in: *Advances in Neural Information Processing Systems*, 2016.
- [30] J. He, Y. Gong, J. Marino, G. Mori, A. Lehrmann, Variational autoencoders with jointly optimized latent dependency structure, in: *International Conference on Learning Representations*, 2018.
- [31] A. Zeroual, F. Harrou, A. Dairi, Y. Sun, Deep learning methods for forecasting COVID-19 time-series data: A comparative study, *Chaos Solitons Fractals* 140 (2020) 110121.
- [32] D. Hallac, S. Vare, S. Boyd, J. Leskovec, Toeplitz inverse covariance-based clustering of multivariate time series data, in: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, pp. 215–223.
- [33] G.E. Hinton, R.R. Salakhutdinov, Reducing the dimensionality of data with neural networks, *Science* 313 (5786) (2006) 504–507.
- [34] L. Li, J. Yan, X. Yang, Y. Jin, Learning interpretable deep state space model for probabilistic time series forecasting, in: *International Joint Conference on Artificial Intelligence*, 2019.
- [35] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: *International Conference on Learning Representations*, 2015.
- [36] D.H. Stolfi, E. Alba, X. Yao, Predicting car park occupancy rates in smart cities, in: *International Conference on Smart Cities*, Springer, 2017, pp. 107–117.
- [37] Y. Qin, D. Song, H. Cheng, W. Cheng, G. Jiang, G.W. Cottrell, A dual-stage attention-based recurrent neural network for time series prediction, in: *International Joint Conference on Artificial Intelligence*, 2017, pp. 2627–2633.
- [38] L. Shen, Y. Wei, Y. Wang, GBT: Two-stage transformer framework for non-stationary time series forecasting, *Neural Netw.* 165 (2023) 953–970.
- [39] C. Challu, K.G. Olivares, B.N. Oreshkin, F.G. Ramirez, M.M. Canseco, A. Dubrawski, NHITS: Neural hierarchical interpolation for time series forecasting, in: AAAI Conference on Artificial Intelligence, Vol. 37, No. 6, 2023, pp. 6989–6997.