

Stochastic Recurrent Neural Network for Multistep Time Series Forecasting

Zexuan Yin(✉)¹[0000-0002-1306-3858] and Paolo Barucca¹[0000-0003-4588-667X]

Department of Computer Science, University College London, WC1E 7JE, London,
United Kingdom

{zexuan.yin.20,p.barucca}@ucl.ac.uk

Abstract. Time series forecasting based on deep architectures has been gaining popularity in recent years due to their ability to model complex non-linear temporal dynamics. The recurrent neural network is one such model capable of handling variable-length input and output. In this paper, we leverage recent advances in deep generative models and the concept of state space models to propose a stochastic adaptation of the recurrent neural network for multistep-ahead time series forecasting, which is trained with stochastic gradient variational Bayes. To capture the stochasticity in time series temporal dynamics, we incorporate a latent random variable into the recurrent neural network to make its transition function stochastic. Our model preserves the architectural workings of a recurrent neural network for which all relevant information is encapsulated in its hidden states, and this flexibility allows our model to be easily integrated into any deep architecture for sequential modelling. We test our model on a wide range of datasets from finance to healthcare; results show that the stochastic recurrent neural network consistently outperforms its deterministic counterpart.

Keywords: state space models · deep generative models · variational inference.

1 Introduction

Time series forecasting is an important task in industry and academia, with applications in fields such as retail demand forecasting [1], finance [2]–[4], and traffic flow prediction [5]. Traditionally, time series forecasting was dominated by linear models such as the autoregressive integrated moving average model (ARIMA), which required prior knowledge about time series structures such as seasonality and trend. With an increasing abundance of data and computational power however, deep learning models have gained much research interest due to their ability to learn complex temporal relationships with a purely data-driven approach; thus requiring minimal human intervention and expertise in the subject matter. In this work, we combine deep learning with state space models (SSM) for sequential modelling. Our work follows recent trend that combines the powerful modelling capabilities of deep learning models with well understood theoretical frameworks such as SSMs.

Recurrent neural networks (RNN) are a popular class of neural networks for sequential modelling. There exists a great abundance of literature on time series modelling with RNNs across different domains [6]–[16]. However, vanilla RNNs have deterministic transition functions, which may limit their expressive power at modelling sequences with high variability and complexity [18]. There is recent evidence that the performance of RNNs on complex sequential data such as speech, music, and videos can be improved when uncertainty is incorporated in the modelling process [19]–[24]. This approach makes an RNN more expressive, as instead of outputting a single deterministic hidden state at every time step, it now considers many possible future paths before making a prediction. Inspired by this, we propose an RNN cell with stochastic hidden states for time series forecasting, which is achieved by inserting a latent random variable into the RNN update function. Our approach corresponds to a state space formulation of time series modelling where the RNN transition function defines the latent state equation, and another neural network defines the observation equation given the RNN hidden state. The main contributions of our paper are as follows:

1. we propose a novel deep stochastic recurrent architecture for multistep-ahead time series forecasting which leverages the ability of regular RNNs to model long-term dynamics and the stochastic framework of state space models.
2. we conduct experiments using publicly available datasets in the fields of finance, traffic flow prediction, air quality forecasting, and disease transmission. Results demonstrate that our stochastic RNN consistently outperforms its deterministic counterpart, and is capable of generating probabilistic forecasts

2 Related Works

2.1 Recurrent Neural Networks

The recurrent neural network (RNN) is a deep architecture specifically designed to handle sequential data, and has delivered state-of-the-art performance in areas such as natural language processing [25]. The structure of the RNN is such that at each time step t , the hidden state of the network - which learns a representation of the raw inputs - is updated using the external input for time t as well as network outputs from the previous step $t-1$. The weights of the network are shared across all time steps and the model is trained using back-propagation. When used to model long sequences of data, the RNN is subject to the vanishing/exploding gradient problem [26]. Variants of the RNN such as the LSTM [27] and the GRU [28] were proposed to address this issue. These variants use gated mechanisms to regulate the flow of information. The GRU is a simplification of the LSTM without a memory cell, which is more computationally efficient to train and offers comparable performance to the LSTM [29].

2.2 Stochastic Gradient Variational Bayes

The authors in [21] proposed the idea of combining an RNN with a variational auto-encoder (VAE) to leverage the RNN’s ability to capture time dependencies and the VAE’s role as a generative model. The proposed structure consists of an encoder that learns a mapping from data to a distribution over latent variables, and a decoder that maps latent representations to data. The model can be efficiently trained with Stochastic Gradient Variational Bayes (SGVB) [30] and enables efficient, large-scale unsupervised variational learning on sequential data. Consider input \mathbf{x} of arbitrary size, we wish to model the data distribution $p(\mathbf{x})$ given some unobserved latent variable \mathbf{z} (again, of arbitrary dimension). The aim is to maximise the marginal likelihood $p(\mathbf{x}) = \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z}) d\mathbf{z}$, which is often intractable when the likelihood $p(\mathbf{x}|\mathbf{z})$ is expressed by a neural network with non-linear layers. Instead we apply variational inference and maximise the evidence lower-bound (ELBO):

$$\begin{aligned} \log p(\mathbf{x}) &= \log \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z}) d\mathbf{z} = \log \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z}) \frac{q(\mathbf{z})}{q(\mathbf{z})} d\mathbf{z} \\ &\geq \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] - KL(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) = ELBO, \quad (1) \end{aligned}$$

where $q(\mathbf{z}|\mathbf{x})$ is the variational approximation to true the posterior distribution $p(\mathbf{z}|\mathbf{x})$ and KL is the Kullback-Leibler divergence. For the rest of this paper we refer to $p(\mathbf{x}|\mathbf{z})$ as the decoding distribution and $q(\mathbf{z}|\mathbf{x})$ as the encoding distribution. The relationship between the marginal likelihood $p(\mathbf{x})$ and the *ELBO* is given by

$$\begin{aligned} \log p(\mathbf{x}) &= \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] - KL(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \\ &\quad + KL(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z}|\mathbf{x})), \quad (2) \end{aligned}$$

where the third KL term specifies the tightness of the lower bound. The expectation $\mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})]$ can be interpreted as an expected negative reconstructed error, and $KL(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))$ serves as a regulariser.

2.3 State Space Models

State space models provide a unified framework for time series modelling; they refer to probabilistic graphical models that describe relationships between observations and the underlying latent variable [35]. Exact inference is feasible only for hidden Markov models (HMM) and linear Gaussian state space models (LGSS) and both are not suitable for long-term prediction [31]. SSMs can be viewed as a probabilistic extension of RNNs. Inside an RNN, the evolution of the hidden states \mathbf{h} is governed by a non-linear transition function $f: \mathbf{h}_{t+1} = f(\mathbf{h}_t, \mathbf{x}_{t+1})$ where \mathbf{x} is the input vector. For an SSM however, the hidden states are assumed to be random variables. It is therefore intuitive to combine the non-linear gated mechanisms of the RNN with the stochastic transitions of the SSM; this creates a sequential generative model that is more expressive than the RNN and better

capable of modelling long-term dynamics than the SSM. There are many recent works that draw connections between SSM and VAE using an RNN. The authors in [18] and [19] propose a sequential VAE with nonlinear state transitions in the latent space, in [32] the authors investigate various inference schemes for variational RNNs, in [22] the authors propose to stack a stochastic SSM layer on top of a deterministic RNN layer, in [23] the authors propose a latent transition scheme that is stochastic conditioned on some inferable parameters, the authors in [33] propose a deep Kalman filter with exogenous inputs, the authors in [34] propose a stochastic variant of the Bi-LSTM, and in [37] the authors use an RNN to parameterise a LGSS.

3 Stochastic Recurrent Neural Network

3.1 Problem Statement

For a multivariate dataset comprised of $N+1$ time series, the covariates $\mathbf{x}_{1:T+\tau} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{T+\tau}\} \in \mathbb{R}^{N \times (T+\tau)}$ and the target variable $y_{1:T} \in \mathbb{R}^{1 \times T}$. We refer to the period $\{T+1, T+2, \dots, T+\tau\}$ as the prediction period, where $\tau \in \mathbb{Z}^+$ is the number of prediction steps and we wish to model the conditional distribution

$$P(y_{T+1:T+\tau} | y_{1:T}, \mathbf{x}_{1:T+\tau}). \quad (3)$$

3.2 Stochastic GRU Cell

Here we introduce the update equations of our stochastic GRU, which forms the backbone of our temporal model:

$$\mathbf{u}_t = \sigma(\mathbf{W}_u \cdot \mathbf{x}_t + \mathbf{C}_u \cdot \mathbf{z}_t + \mathbf{M}_u \cdot \mathbf{h}_{t-1} + \mathbf{b}_u) \quad (4)$$

$$\mathbf{r}_t = \sigma(\mathbf{W}_r \cdot \mathbf{x}_t + \mathbf{C}_r \cdot \mathbf{z}_t + \mathbf{M}_r \cdot \mathbf{h}_{t-1} + \mathbf{b}_r) \quad (5)$$

$$\tilde{\mathbf{h}}_t = \tanh(\mathbf{W}_h \cdot \mathbf{x}_t + \mathbf{C}_h \cdot \mathbf{z}_t + \mathbf{r}_t \odot \mathbf{M}_h \cdot \mathbf{h}_{t-1} + \mathbf{b}_h) \quad (6)$$

$$\mathbf{h}_t = \mathbf{u}_t \odot \mathbf{h}_{t-1} + (1 - \mathbf{u}_t) \odot \tilde{\mathbf{h}}_t, \quad (7)$$

where σ is the sigmoid activation function, \mathbf{z}_t is a latent random variable which captures the stochasticity of the temporal process, \mathbf{u}_t and \mathbf{r}_t represent the update and reset gates, \mathbf{W} , \mathbf{C} and \mathbf{M} are weight matrices, \mathbf{b} is the bias matrix, \mathbf{h}_t is the GRU hidden state and \odot is the element-wise Hadamard product. Our stochastic adaptation can be seen as a generalisation of the regular GRU, i.e. when $\mathbf{C} = 0$, we have a regular GRU cell [28].

3.3 Generative Model

The role of the generative model is to establish probabilistic relationships between the target variable y_t , the intermediate variables of interest $(\mathbf{h}_t, \mathbf{z}_t)$, and the input \mathbf{x}_t . Our model uses neural networks to describe the non-linear transition and emission processes, and we preserve the architectural workings of an

RNN - relevant information is encoded within the hidden states that evolve with time, and the hidden states contain all necessary information required to estimate the target variable at each time step. A graphical representation of the generative model is shown in Fig 1a, the RNN transitions are now stochastic, facilitated by the random variable \mathbf{z}_t . The joint probability distribution of the generative model can be factorised as follows:

$$p_{\theta}(y_{2:T}, \mathbf{z}_{2:T}, \mathbf{h}_{2:T} | \mathbf{x}_{1:T}) = \prod_{t=2}^T p_{\theta_1}(y_t | \mathbf{h}_t) p_{\theta_2}(\mathbf{h}_t | \mathbf{h}_{t-1}, \mathbf{z}_t, \mathbf{x}_t) p_{\theta_3}(\mathbf{z}_t | \mathbf{h}_{t-1}) \quad (8)$$

where

$$p_{\theta_3}(\mathbf{z}_t | \mathbf{h}_{t-1}) = N(\boldsymbol{\mu}(\mathbf{h}_{t-1}), \boldsymbol{\sigma}^2(\mathbf{h}_{t-1})\mathbf{I}) \quad (9)$$

$$\mathbf{h}_t = GRU(\mathbf{h}_{t-1}, \mathbf{z}_t, \mathbf{x}_t) \quad (10)$$

$$y_t \sim p_{\theta_1}(y_t | \mathbf{h}_t) = N(\mu(\mathbf{h}_t), \sigma^2(\mathbf{h}_t)), \quad (11)$$

where GRU is the stochastic GRU update function given by (4)–(7). (9) defines the prior distribution of \mathbf{z}_t , which we assume to have an isotropic Gaussian prior (covariance matrix is diagonal) parameterised by a multi-layer perceptron (MLP). When conditioning on past time series for prediction, we use (9), (10) and the last available hidden state \mathbf{h}_{last} to calculate \mathbf{h}_1 for the next sequence, otherwise we initialise them to $\mathbf{0}$. We refer to the collection of parameters of the generative model as θ , i.e. $\theta = \{\theta_1, \theta_2, \theta_3\}$. We refer to (11) as our generative distribution, which is parameterised by an MLP.




Fig. 1: Proposed generative and inference models

3.4 Inference Model

We wish to maximise the marginal log-likelihood function $\log p_{\theta}(y_{2:T} | \mathbf{x}_{2:T})$, however the random variable \mathbf{z}_t of the non-linear SSM cannot be analytically integrated out. We instead maximise the variational lower bound (ELBO) with

respect to the generative model parameters θ and some inference model parameter which we call ϕ [36]. The variational approximation of the true posterior $p(\mathbf{z}_{2:T}, \mathbf{h}_{2:T} | \mathbf{y}_{1:T}, \mathbf{x}_{1:T})$ can be factorised as follows:

$$q_\phi(\mathbf{z}_{2:T}, \mathbf{h}_{2:T} | \mathbf{y}_{1:T}, \mathbf{x}_{1:T}) = \prod_{t=2}^T q_\phi(\mathbf{z}_t | \mathbf{y}_{1:T}) q_\phi(\mathbf{h}_t | \mathbf{h}_{t-1}, \mathbf{z}_t, \mathbf{x}_t) \quad (12)$$

and

$$q_\phi(\mathbf{h}_t | \mathbf{h}_{t-1}, \mathbf{z}_t, \mathbf{x}_t) = p_{\theta_2}(\mathbf{h}_t | \mathbf{h}_{t-1}, \mathbf{z}_t, \mathbf{x}_t), \quad (13)$$

where p_{θ_2} is the same as in (8); this is due to the fact that the GRU transition function is fully deterministic conditioned on knowing \mathbf{z}_t and hence p_{θ_2} is just a delta distribution centered at the GRU output value given by (4)–(7). The graphical model of the inference network is given in Fig 1b. Since the purpose of the inference model is to infer the filtering distribution $q_\phi(\mathbf{z}_t | \mathbf{y}_{1:t})$, and that an RNN hidden state contains a representation of current and past inputs, we use a second GRU model with hidden states \mathbf{g}_t as our inference model, which takes the observed target values y_t and previous hidden state \mathbf{g}_{t-1} as inputs and maps \mathbf{g}_t to the inferred value of \mathbf{z}_t :

$$\mathbf{g}_t = GRU(\mathbf{g}_{t-1}, \mathbf{y}_t) \quad (14)$$

$$\mathbf{z}_t \sim q_\phi(\mathbf{z}_t | \mathbf{y}_{1:t}) = N(\boldsymbol{\mu}(\mathbf{g}_t), \boldsymbol{\sigma}^2(\mathbf{g}_t)I). \quad (15)$$

3.5 Model Training

The objective function of our stochastic RNN is the ELBO $L(\theta, \phi)$ given by:

$$\begin{aligned} L(\theta, \phi) &= \int \int q_\phi \log \frac{p_\theta}{q_\phi} d\mathbf{z}_{2:T} d\mathbf{h}_{2:T} \\ &= \sum_{n=2}^T \mathbb{E}_{q_\phi} [\log p_\theta(y_t | \mathbf{h}_t)] - KL(q_\phi(\mathbf{z}_t | \mathbf{y}_{1:t}) || p_\theta(\mathbf{z}_t | \mathbf{h}_{t-1})), \end{aligned} \quad (16)$$

where p_θ and q_ϕ are the generative and inference distributions given by (8) and (12) respectively. During training, we use the posterior network (15) to infer the latent variable \mathbf{z}_t used for reconstruction. During testing we use the prior network (9) to forecast 1-step-ahead \mathbf{z}_t , which has been trained using the KL term in the ELBO function. We seek to optimise the ELBO with respect to decoder parameters θ and encoder parameters ϕ jointly, i.e. we wish to find:

$$(\theta^*, \phi^*) = \underset{\theta, \phi}{\operatorname{argmax}} L(\theta, \phi). \quad (17)$$

Since we do not back-propagate through a sampling operation, we apply the reparameterisation trick [30] to write

$$\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}, \quad (18)$$

where $\boldsymbol{\epsilon} \sim N(0, I)$ and we sample from $\boldsymbol{\epsilon}$ instead. The KL divergence term in (16) can be analytically computed since we assume the prior and posterior of \mathbf{z}_t to be normally distributed.

3.6 Model Prediction

Given the last available GRU hidden state \mathbf{h}_{last} , prediction window τ and covariates $\mathbf{x}_{T+1:T+\tau}$, we generate predicted target values in an autoregressive manner, assuming that at every time step the hidden state of the GRU \mathbf{h}_t contains all relevant information up to time t . The prediction algorithm of our stochastic GRU is given by Algorithm 1.

```

Input:  $\tau, \mathbf{h}_{last}, \mathbf{x}_{T+1:T+\tau}$ 
Output:  $y_{T+1:T+\tau}$ 
for  $t \leftarrow 1$  to  $\tau$  do
     $\mathbf{z}_t \sim p_{\theta_3}(\mathbf{z}_t | \mathbf{h}_{last})$ 
     $\mathbf{h}_t \leftarrow GRU(\mathbf{h}_{last}, \mathbf{z}_t, \mathbf{x}_t)$ 
     $y_t \sim p_{\theta_1}(y_t | \mathbf{h}_t)$ 
     $\mathbf{h}_{last} \leftarrow \mathbf{h}_t$ 
end

```

Algorithm 1: Prediction algorithm for stochastic GRU

4 Experiments

We highlight the model performance on 6 publicly available datasets:

1. Equity options trading price time series available from the Chicago Board Options Exchange (CBOE) datashop. This dataset describes the minute-level traded prices of an option throughout the day. We study 3 options with Microsoft and Amazon stocks as underlyings where $\mathbf{x}_t =$ underlying stock price and $y_t =$ traded option price
2. The Beijing PM2.5 multivariate dataset describes hourly PM2.5 (a type of air pollution) concentrations of the US Embassy in Beijing, and is freely available from the UCI Machine Learning Repository. The covariates we use are $\mathbf{x}_t =$ temperature, pressure, cumulated wind speed, Dew point, cumulated hours of rainfall and cumulated hours of snow, and $y_t =$ PM2.5 concentration. We use data from 01/11/2014 onwards
3. The Metro Interstate Traffic Volume dataset describes the hourly interstate 94 Westbound traffic volume for MN DoT ATR station 301, roughly midway between Minneapolis and ST Paul, MN. This dataset is available on the UCI Machine Learning Repository. The covariates we use in this experiment are $\mathbf{x}_t =$ temperature, mm of rainfall in the hour, mm of snow in the hour, and percentage of cloud cover, and $y_t =$ hourly traffic volume. We use data from 02/10/2012 9AM onwards
4. The Hungarian Chickenpox dataset describes weekly chickenpox cases (childhood disease) in different Hungarian counties. This dataset is also available on the UCI Machine Learning Repository. For this experiment, $y_t =$ number of chickenpox cases in the Hungarian capital city Budapest, $\mathbf{x}_t =$ number

of chickenpox cases in Pest, Bacs, Komárom and Heves, which are 4 nearby counties. We use data from 03/01/2005 onwards

We generate probabilistic forecasts using 500 Monte-Carlo simulations and we take the mean predictions as our point forecasts to compute the error metrics. We tested the number of simulations from 100 to 1000 and found that above 500, the differences in performance were small, and with fewer than 500 we could not obtain realistic confidence intervals for some time series. We provide graphical illustrations of the prediction results in Fig 2a–2f. We compare our model performance against an AR(1) model assuming the prediction is the same as the last observed value ($y_{T+\tau} = y_T$), a standard LSTM model and a standard GRU model. For the performance metric, we normalise the root-mean-squared-error (rmse) to enable comparison between time series:

$$nrmse = \frac{\sqrt{\frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{N}}}{\bar{y}}, \quad (19)$$

where $\bar{y} = mean(y)$, \hat{y}_i is the mean predicted value of y_i , and N is the prediction size. For replication purposes, in Table 1 we provide (in order): number of training, validation and conditioning steps, (non-overlapping) sequence lengths used for training, number of prediction steps, dimensions of \mathbf{z}_t , \mathbf{h}_t and \mathbf{g}_t , details about the MLPs corresponding to (9)(\mathbf{z}_t prior) and (15)(\mathbf{z}_t post) in the form of (n layers, n hidden units per layer), and lastly the size of the hidden states of the benchmark RNNs (LSTM and GRU). we use the ADAM optimiser with a learning rate of 0.001. In Table 2–5 we observe that the nrmse of the

Table 1: Model and training parameters

dataset	train	val	cond	seq length	pred	\mathbf{z}_t	\mathbf{h}_t	\mathbf{g}_t	\mathbf{z}_t prior	\mathbf{z}_t post	RNN hid
Options	300	30	10	10	30	50	64	64	(4,64)	(4,64)	64
PM2.5	1200	200	10	10	30	50	64	64	(4,64)	(4,64)	64
Traffic volume	1000	200	20	20	30	30	128	128	(4,128)	(4,128)	128
Hungarian chickenpox	300	150	10	10	30	50	128	128	(4,128)	(4,128)	128

stochastic GRU is lower than its deterministic counterpart for all datasets investigated and across all prediction steps. This shows that our proposed method can better capture both long and short-term dynamics of the time series. With respect to multistep time series forecasting, it is often difficult to accurately model the long-term dynamics. Our approach provides an additional degree of freedom facilitated by the latent random variable which needs to be inferred using the inference network; we believe this allows the stochastic GRU to better capture the stochasticity of the time series at every time step. In Fig 2e for example, we observe that our model captures well the long-term cyclicality of the traffic volume, and in Fig 2d where the time series is much more erratic, our model can still accurately predict the general shape of the time series in the prediction period.

Table 2: nrmse for 30 steps-ahead options price predictions

Option	Description	Ours	AR(1)	LSTM	GRU
MSFT call	strike 190, expiry 17/09/2021	0.0010	0.0109	0.0015	0.0015
MSFT put	strike 315, expiry 16/07/2021	0.0004	0.0049	0.0006	0.0007
AMZN put	strike 3345, expiry 22/01/2021	0.0032	0.0120	0.0038	0.0038

Table 3: nrmse for 30 steps-ahead PM2.5 concentration predictions

steps	5	10	15	20	25	30
Ours	0.1879	0.2474	0.4238	0.4588	0.6373	0.6523
AR(1)	0.3092	1.0957	0.7330	0.6846	1.0045	1.1289
LSTM	0.4797	0.6579	0.4728	0.4638	0.8324	0.8318
GRU	0.4846	0.5553	0.4789	0.4919	0.6872	0.6902

Table 4: nrmse for 30 steps-ahead traffic volume predictions

steps	5	10	15	20	25	30
Ours	0.4284	0.2444	0.2262	0.2508	0.2867	0.2605
AR(1)	1.2039	1.0541	1.0194	1.0283	1.1179	1.0910
LSTM	0.8649	0.5936	0.4416	0.4362	0.5591	0.5446
GRU	0.8425	0.5872	0.4457	0.4376	0.5510	0.5519


Table 5: nrmse for 30 steps-ahead Hungarian chickenpox predictions

steps	5	10	15	20	25	30
Ours	0.6585	0.6213	0.5795	0.5905	0.6548	0.5541
AR(1)	0.7366	0.7108	0.9126	0.9809	1.0494	1.0315
LSTM	0.7215	0.6687	0.9057	1.0717	0.8471	0.7757
GRU	0.6795	0.6379	0.6825	0.6196	0.6355	0.6739


Table 6: nrmse of MLP benchmark and our proposed model for 30 steps-ahead forecasts

	MSFT call	MSFT put	AMZN put	PM2.5	Metro	Chickenpox
Ours	0.0010	0.0004	0.0032	0.6523	0.2605	0.5541
MLP	0.0024	0.0005	0.0141	0.7058	0.6059	0.5746


To investigate the effectiveness of our temporal model, we compare our prediction errors against a model without a temporal component, which is constructed using a 3-layer MLP with 5 hidden nodes and ReLU activation functions. Since we are using covariates in the prediction period (3), we would like to verify that our model can outperform a simple regression-type benchmark which approximates a function of the form $y_t = f_\psi(\mathbf{x}_t)$; we use the MLP to parameterise the function f_ψ . We observe in Table 6 that our proposed model outperforms a regression-type benchmark for all the experiments, which shows the effectiveness of our temporal model. It is also worth noting that in our experiments we use the actual values of the future covariates. In a real forecasting setting, the future covariates themselves could be outputs of other mathematical models, or they could be estimated using expert judgement.




(a) MSFT call option, strike 190, expiry 17/09/2021




(b) MSFT put option, strike 315, expiry 16/07/2021




(c) AMZN put option, strike 3345, expiry 22/01/2021



(d) PM2.5 concentration forecasts up to 30 steps ahead



(e) Traffic volume forecasts up to 30 steps ahead



(f) Hungarian chickenpox cases forecasts up to 30 steps ahead

5 Conclusion

In this paper we have presented a stochastic adaptation of the Gated Recurrent Unit which is trained with stochastic gradient variational Bayes. Our model design preserves the architectural workings of an RNN, which encapsulates all relevant information into the hidden state, however our adaptation takes inspiration from the stochastic transition functions of state space models by injecting a latent random variable into the update functions of the GRU, which allows the GRU to be more expressive at modelling highly variable transition dynamics compared to a regular RNN with deterministic transition functions. We have tested the performance of our model on different publicly available datasets and results demonstrate the effectiveness of our design. Given that GRUs are now popular building blocks for much more complex deep architectures, we believe that our stochastic GRU could prove useful as an improved component which can be integrated into sophisticated deep learning models for sequential modelling.

Acknowledgments

We would like to thank Dr Fabio Caccioli (Dpt of Computer Science, UCL) for proofreading this manuscript and for his questions and feedback.

References

1. Bandara, K., Shi, P., Bergmeir, C., Hewamalage, H., Tran, Q., Seaman, B.: Sales demand forecast in E-commerce using a long short-term memory neural network methodology. In: ICONIP, Sydney, NSW, Australia, pp. 462–474 (2019)
2. McNally, S., Roche, J., Caton, S.: Predicting the price of Bitcoin using machine learning. In: PDP, Cambridge, UK, pp. 339–343 (2018)
3. Hu, Z., Zhao, Y., Khushi, M.: A Survey of Forex and stock price prediction using deep learning. *Appl.Syst.Innov.* vol. 4, no. 9 (2021). DOI: 10.3390/asi4010009.
4. Zhang, R., Yuan, Z., Shao, X.: A new combined CNN-RNN model for sector stock price analysis. In: COMPSAC, Tokyo, Japan (2018)
5. Lv, Y., Duan, Y., Kang, W.: Traffic flow prediction with big data: a deep learning approach. *IEEE T. INTELL. TRANSP.* vol. 16, no. 2, pp. 865–873 (2014)
6. Dolatabadi, A., Abdeltawab, H., Mohamed, Y.: Hybrid deep learning-based model for wind speed forecasting based on DWPT and bidirectional LSTM Network. *IEEE Access.* vol. 8, pp. 229219–229232 (2020). DOI: 10.1109/ACCESS.2020.3047077.
7. Alazab, M., Khan, S., Krishnan, S., Pham, Q., Reddy, M., Gadekallu, T.: A multi-directional LSTM model for predicting the stability of a smart grid. *IEEE Access.* vol. 8, pp. 85454–85463 (2020). DOI: 10.1109/ACCESS.2020.2991067.
8. Liu, T., Wu, T., Wang, M, Fu, M, Kang, J., Zhang, H.: Recurrent neural networks based on LSTM for predicting geomagnetic field. In: ICARES, Bali, Indonesia (2018)
9. Lai, G., Chang, W., Yang, Y., Liu, H.: Modelling long-and short-term temporal patterns with deep neural networks. In: SIGIR (2018)
10. Apaydin, H., Feizi, H., Sattari, M.T., Cloak, M.S., Shamshirband, S., Chau, K.W.: Comparative analysis of recurrent neural network architectures for reservoir inflow forecasting: *Water.* vol. 12 (2020). DOI: 10.3390/w12051500.
11. Di Persio, L., Honchar, O.: Analysis of recurrent neural networks for short-term energy load forecasting. In: AIP (2017)
12. Meng, X., Wang, R., Zhang, X., Wang, M., Ma, H., Wang, Z.: Hybrid neural network based on GRU with uncertain factors for forecasting ultra-short term wind power. In: IAI (2020)
13. Khaldi, R., El Afia, A., Chiheb, R.: Impact of multistep forecasting strategies on recurrent neural networks performance for short and long horizons. In: BDIoT (2019)
14. Mattos, C.L.C., Barreto, G.A.: A stochastic variational framework for recurrent gaussian process models. *Neural Networks.* vol. 112, pp. 54–72 (2019). DOI: 10.1016/j.neunet.2019.01.005.
15. Seleznev, A., Mukhin, D., Gavrilov, A., Loskutov, E., Feigin, A.: Bayesian framework for simulation of dynamical systems from multidimensional data using recurrent neural network. *Chaos.* vol. 29 (2019). DOI: 10.1063/1.5128372.
16. Alhussein, M., Aurangzeb, K., Haider, S.: Hybrid CNN-LSTM model for short-term individual household load forecasting. *IEEE Access.* vol. 8, pp. 180544–180557 (2020). DOI: 10.1109/ACCESS.2020.3028281.
17. Sezer, O., Gudelek, M., Ozbayoglu, A.: Financial time series forecasting with deep learning: A systematic literature review: 2005-2019. *Appl. Soft Comput.* vol. 90 (2020). DOI: 10.1016/j.asoc.2020.106181.
18. Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A., Bengio, Y.: A recurrent latent variable model for sequential data. In: NIPS, Montreal, Canada (2015)
19. Bayer, J., Osendorfer, C.: Learning stochastic recurrent networks. *arXiv preprint. arXiv: 1411.7610* (2014)

20. Goyal, A., Sordoni, A., Cote, M., Ke, N., Bengio, Y.: Z-forcing: training stochastic recurrent networks. In: NIPS, California, USA, pp. 6716–6726 (2017)
21. Fabius, O., Amersfoort, J.R.: Variational recurrent auto-encoders. arXiv preprint. arXiv: 1412.6581 (2017)
22. Fraccaro, M., Sønderby, S., Paquet, U., Winther, O.: Sequential neural models with stochastic layers. In: NIPS, Barcelona, Spain (2016)
23. Karl, M., Soelch, M., Bayer, J., Smagt, P.: Deep Variational Bayes Filters: Unsupervised learning of state space models from raw data. In: ICLR, Toulon, France (2017)
24. Franceschi, J., Delasalles, E., Chen, M., Lamprier, S., Gallinari, P.: Stochastic latent residual video prediction. In: ICML, pp. 3233–3246 (2020)
25. Young, T., Hazarika, D., Poria, S., Cambria, E.: Recent trends in deep learning based natural language processing. *IEEE Comput. Intell. Mag.* vol. 13, no. 3, pp. 55–75 (2018). DOI: 10.1109/MCI.2018.2840738.
26. Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training recurrent neural networks. In: ICML, Atlanta, GA, USA, pp. 1310–1318 (2013)
27. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* vol. 9, no. 8, pp. 1735–1780 (1997)
28. Cho, K., Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: encoder-decoder approaches. In: SSST, Doha, Qatar, pp. 103–111 (2014)
29. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modelling. arXiv preprint. arXiv: 1412.3555 (2014)
30. Kingma, D., Welling, M.: Auto-encoding variational Bayes. In: ICLR, Banff, Canada (2014)
31. Liitiainen, E., Lendasse, A.: Long-term prediction of time series using state-space models. In: ICANN, Athens, Greece (2006)
32. Krishnan, R., Shalit, U., Sontag, D.: Structured inference networks for nonlinear state space models. In: AAAI, California, USA, pp. 2101–2109 (2017)
33. Krishnan, R., Shalit, U., Sontag, D.: Deep Kalman filters. arXiv preprint. arXiv: 1511.05121 (2015)
34. Shabanian, S., Arpit, D., Trischler, A., Bengio, Y.: Variational bi-LSTMs. arXiv preprint. arXiv: 1711.05717 (2017)
35. Durbin, J., Koopman, S.: Time series analysis by state space methods. volume 38. Oxford University Press (2012)
36. Jordan, M., Ghahramani, Z., Jaakkola, T., Saul, L.: An introduction to variational methods for graphical models. *Machine Learning.* vol. 37, no. 2, pp. 183–233 (1999)
37. Rangapuram, S., Seeger, M., Gasthaus, J., Stella, L., Yang, Y., Janushowski, T.: Deep state space models for time series forecasting. In: NIPS, Montreal, Canada (2018)