

1 **Exploring Bayesian deep learning for weather**
2 **forecasting with the Lorenz 84 system**

3 **Yang Liu^{1,2}, Jisk Attema¹, Wilco Hazeleger³**

4 ¹Netherlands eScience Center, Amsterdam, the Netherlands

5 ²Meteorology and Air Quality Group, Wageningen University, Wageningen, the Netherlands

6 ³Faculty of Geosciences, Utrecht University, Utrecht, the Netherlands

7 **Key Points:**

- 8 • Bayesian deep neural networks are able to represent uncertainties relevant to weather
9 forecasting.
10 • A trained Bayesian deep neural network can preserve the physical consistency of the
11 Lorenz 84 system.
12 • The forecast quality of the trained Bayesian deep neural network deteriorates with
13 forecast lead time and it is state-dependent.

Corresponding author: Yang Liu, y.liu@esciencecenter.nl

Abstract

[The need for uncertainty quantification placed by weather forecasting makes Bayesian deep learning (BDL) a suited candidate for data-driven weather forecasting. In this study, we use Bayesian Long-Short Term Memory neural networks (BayesLSTMs) to forecast output from the Lorenz 84 system with seasonal forcing. The latter represents the dynamics of large scale eddies (Rossby waves) on a westerly jet. We show that forecasts with the BayesLSTM can stay close to the attractor of the Lorenz model and conclude that they represent the nonlinear relations between each component in this simplified atmospheric circulation system. The forecasts are evaluated against persistence and a Vector Autoregressive Model (VAR). We demonstrate that the BayesLSTMs can produce reliable probabilistic forecasts and address uncertainties relevant to weather forecasting. Our study indicates that BDL is an easy and fast solution for probabilistic weather forecast and is promising to enhance weather forecasting capabilities at short to medium-range timescales.]

Plain Language Summary

[Recent developments in artificial intelligence (AI) have brought many techniques to climate science. Among these techniques, deep neural networks (DNN) serve as good candidates to improve and speed up weather forecasts. However, these DNN always have fixed structure and therefore can not satisfy the need of weather forecast for uncertainty estimation. To solve the problem, we introduce Bayesian deep learning (BDL), which is probabilistic and enables uncertainty quantification. In this study, we explore the BDL with a simplified chaotic system, the Lorenz 84 model with seasonal forcing. We test and use BDL to forecast the Lorenz 84 system and evaluate its probabilistic forecast skill against the persistence and a baseline statistical model. Our study indicates that the BDL is able to account for the uncertainty required by weather forecasting and it represents the nonlinear relations between each component in this simplified atmospheric circulation system. It is a promising tool for preliminary and quick probabilistic forecasts and therefore can enhance weather forecasting capabilities.]

1 Introduction

Deep neural networks (DNNs) are capable of representing intricate features of data and have been proven to be useful for many scientific disciplines (e.g., LeCun et al., 2015), including weather forecasting and climate science (Reichstein et al., 2019). It has been demonstrated by recent studies that typical DNN are able to mimic and predict the behavior of chaotic systems (e.g. Hochreiter & Schmidhuber, 1997; Chattopadhyay et al., 2019) and therefore they are potentially applicable to weather forecasting. However, mostly deterministic DNNs are considered and these are prone to overfitting and this can result in over-confident forecasts (Shridhar et al., 2019).

Due to the chaotic nature of the atmospheric dynamics and uncertainties in both initial conditions and models representing the atmosphere, weather forecasts are of probabilistic nature. In general, uncertainty estimation is achieved via an ensemble approach within trustworthy Numerical Weather Forecast systems (NWP) (Gneiting et al., 2007; Leutbecher & Palmer, 2008). However, this strategy is computationally expensive for NWP-based weather forecasts. Concerning the deep learning approaches, in order to meet the requirement for uncertainty quantification, many attempts have been made to adapt deterministic DNN to weather forecasting (e.g., Scher & Messori, 2018). These efforts mainly involve generating a DNN-based ensemble through perturbing either the training data or the structure of DNN (e.g., Zaier et al., 2010; H.-z. Wang et al., 2017). However, in practice, this technique is computationally expensive due to multiple training cycles that are needed and it is often difficult to manually select proper perturbations which can approximate the error growth of a real dynamical system. Fortunately, recent developments in deep learning have led

63 to a branch of DNN to cope with overfitting and address uncertainties, which is known as
64 Bayesian deep learning (BDL).

65 Unlike feed-forward DNN, BDL is constructed by replacing fixed weights with distri-
66 butions and therefore are designed to represent uncertainties (Blundell et al., 2015). With
67 a well-defined likelihood function, BDL is able to capture both the aleatoric and epistemic
68 uncertainty (Kendall & Gal, 2017; Shridhar et al., 2018, 2019). They can avoid making
69 over-confident decisions and incorporate regularization naturally by implementing the vari-
70 ational approaches (Shridhar et al., 2019). Together with the simplicity of implementing
71 BDL on an already defined deep neural network, these make BDL an attractive approach
72 for representing atmospheric dynamics and the practice of weather forecasting (Vandal et
73 al., 2018).

74 An operational numerical weather forecast system is very complex. Here, we want to
75 understand the characteristics of BDL within a simplified dynamical system that represents
76 the essence of midlatitude atmospheric dynamics and explore the types of uncertainties
77 addressed by BDL. In particular we examine how BDL can replicate the phase and amplitude
78 of midlatitude Rossby waves on a jet as represented in a Lorenz 84 model (Lorenz, 1984;
79 H. Wang et al., 2014). The predictive nature and time scale of propagation and development
80 of Rossby waves form the basis of short to medium-range weather forecasting. We will assess
81 whether BDL can represent the predictability of this simplified atmospheric circulation
82 system. We notice that the concept of BDL in the perspective of weather forecasting is
83 quite similar to the implementation of the Bayesian theorem in data assimilation (e.g., Ghil
84 & Malanotte-Rizzoli, 1991; Navon, 2009; Bannister, 2017).

85 Long-Short Term Memory neural networks (LSTMs) have a network structure and
86 characteristics that are found to be suitable to represent fluids in environmental studies
87 (Liu et al., 2020). In this study, we explore BDL by turning LSTMs into Bayesian LSTMs
88 (BayesLSTMs). We will use the BayesLSTMs to forecast the Lorenz 84 model and assess the
89 forecast quality in the spatial and temporal space at different lead times. The probabilistic
90 forecasts produced by the BayesLSTM will be evaluated against those with persistence of
91 initial conditions and a baseline statistical model. An emphasis is placed on the uncertainties
92 represented by the BayesLSTM and its capacity in preserving the physical consistency in a
93 simplified atmospheric circulation system.

94 The paper is organized as follows: we elaborate on the concept of BDL and Lorenz 84
95 model with seasonal forcing in the section Methodology. An analysis of uncertainty estima-
96 tion with BDL, and the procedure of sampling the BayesLSTM and generating ensemble
97 forecasts are also provided in this section. The probabilistic forecasts of the Lorenz 84 sys-
98 tem using the BayesLSTM are elucidated and analyzed in the section Results. This section
99 also includes forecasts with persistence and a baseline statistical model for comparison and
100 evaluation. Finally, in the section Conclusion and Discussion, we summarize this study and
101 provide our perspective for future work.

102 **2 Methodology**

103 In this section, we briefly introduce the Lorenz 84 model with seasonal forcing and
104 elaborate upon the concept of BDL as well as how an LSTM network is transformed into
105 a BayesLSTM. Based on the characteristics of BDL, the procedure of producing ensemble
106 forecasts and a description of uncertainty estimation with BayesLSTM is presented in this
107 section.

108 **2.1 Lorenz 84 Model with Seasonal Forcing**

109 The Lorenz 84 system represents the general circulation of the atmosphere in a low
110 dimensional space and therefore it is useful as a baseline model for exploring BayesLSTMs

111 in weather forecasting (Lorenz, 1984). To incorporate more realistic features into the simple
 112 Rossby wave evolution system, we add a seasonal forcing to the classical Lorenz 84 model.
 113 The dynamical system is formulated as follows:

$$\begin{aligned}
 \frac{dX}{dT} &= -Y^2 - Z^2 - aX + aF(1 + \epsilon \cos(\omega T)) \\
 \frac{dY}{dT} &= XY - bXZ - Y + G(1 + \epsilon \sin(\omega T)) \\
 \frac{dZ}{dT} &= bXY + XZ - Z
 \end{aligned} \tag{1}$$

114 where X represents the intensity of the westerly wind circulating around the globe, Y and
 115 Z represent the cosine and sine phases of a chain of superimposed large-scale eddies, T is
 116 the time, a and b indicate mechanical and thermal damping, F and G the symmetric and
 117 asymmetric thermal forcing, ϵ the intensity of seasonal forcing, and ω the angular frequency
 118 of seasonality (Freire et al., 2008). In this study, we mainly focus on the sensitivity of the
 119 forecast quality to variations in the initial condition X and model parameter a .

120 To obtain a chaotic system that is suitable for the assessment of the BayesLSTM
 121 forecast, we chose the model parameters to be $a = 0.25$, $b = 4.0$, $F = 8.0$, $G = 1.0$, $\epsilon = 0.4$,
 122 and the initial conditions as $X, Y, Z = 1.0$. One unit of time in the Lorenz model corresponds
 123 to 5 days. The damping time of the wave is about 5 days (Lorenz, 1984). We sample the
 124 system with a temporal resolution equal to 1/30 unit time, which is 4 hours. The period of
 125 seasonal forcing is taken as 73 unit time steps and then the period of the entire system is
 126 equivalent to 356 days. With this configuration, the trajectories and the time series of each
 127 variable are shown in Figure 1. Unless specifically noted, the time step and lead time steps
 128 in this paper are based on the sampling interval, which is 4 hours.

129 2.2 BayesLSTM and Bayes by Backprop

130 Our aim is to investigate whether the BayesLSTM can represent the Lorenz 84 model
 131 described above. We can add Bayesian inference to an existing neural network by replacing
 132 fixed weights with distributions (e.g. see Figure 1 in Blundell et al., 2015). Given the
 133 structure of an LSTM network (Hochreiter & Schmidhuber, 1997), the Bayesian form of an
 134 LSTM network can be represented by equation 2:

$$\begin{aligned}
 i_t &= \sigma(W_{xi}^s \circ x_t + W_{hi}^s \circ h_{t-1} + W_{ci} \circ c_{t-1} + b_i) \\
 f_t &= \sigma(W_{xf}^s \circ x_t + W_{hf}^s \circ h_{t-1} + W_{cf} \circ c_{t-1} + b_f) \\
 c_t &= f_t \circ c_{t-1} + i_t \circ \tanh(W_{xc}^s \circ x_t + W_{hc}^s \circ h_{t-1} + b_c) \\
 o_t &= \sigma(W_{xo}^s \circ x_t + W_{ho}^s \circ h_{t-1} + W_{co} \circ c_t + b_o) \\
 h_t &= o_t \circ \tanh(c_t)
 \end{aligned} \tag{2}$$

135 with i_t the input gate, f_t the forget gate, c_t the cell state, o_t the output gate, h_t the hidden
 136 state, W^s the weight distribution, x_t the input, b the bias, \circ the element-wise product, σ
 137 the sigmoid function and \tanh the hyperbolic tangent function. The subscripts describe
 138 the corresponding weight matrix to different gates and states. W_{xi}^s indicates the weight
 139 matrix of input values related to the input gate, while W_{hf}^s represents the weight matrix
 140 of hidden states corresponding to the forget gate. The subscript t indicates the time step.
 141 The structure of a BayesLSTM is illustrated in Figure 1c.

142 We need to search for the weight distribution W^s , thus the posterior $p(w|D)$ where
 143 w denotes the weight and $D = (x_j, y_j)_j$ indicates the training set. As the true poste-
 144 rior probability distribution is intractable (because of the marginal likelihood), we use a

145 variational inference scheme, namely the Bayes by Backprop approach, to approximate it
 146 (Blundell et al., 2015; Shridhar et al., 2018, 2019). The reason for choosing this method
 147 is elaborated upon in detail in the supporting material. A simple variational distribution
 148 $q(w|\theta)$ (where θ is the variational posterior parameter), such as a Gaussian distribution, or
 149 a lognormal distribution is often chosen (Blundell et al., 2015; Shridhar et al., 2018; Van-
 150 dal et al., 2018). Here we approximate the posterior $p(w|D)$ with a Gaussian distribution
 151 $q(w|\theta)$, which consists of two trainable parameters $\mu \in \mathbb{R}^d$ and $\sigma \in \mathbb{R}^d$. As a result, θ in the
 152 assumed variational distribution $q(w|\theta)$ can be denoted by $\mathcal{N}(\theta|\mu, \sigma^2)$.

153 The gap between the chosen variational distribution and the exact posterior distri-
 154 bution is reduced using the Kullback-Leibler (KL) divergence between $p(w|D)$ and $q(w|\theta)$
 155 (Graves, 2011; Blundell et al., 2015). KL divergence measures the similarity between two
 156 distributions and in this we define the optimal parameters θ^* as:

$$\begin{aligned} \theta^* &= \arg \min_{\theta} [q(w|\theta) || p(w|D)] \\ &= \arg \min_{\theta} KL[q(w|\theta) || p(w)] - \mathbb{E}_{q(w|\theta)} [\log p(D|w)] + \log p(D) \end{aligned} \quad (3)$$

157 where KL indicates the full KL divergence operation and \mathbb{E} represents the expectation.
 158 This equation includes a data dependent part $\mathbb{E}_{q(w|\theta)} [\log p(D|w)]$ and a prior dependent
 159 part $KL[q(w|\theta) || p(w)]$ (Neal & Hinton, 1998; Blundell et al., 2015; Shridhar et al., 2019).
 160 We sample the weight w from $q(w|\theta)$ and the cost function that we optimize is:

$$\mathcal{F}(D, \theta) = \sum_{n=1}^N \log q(w^{(n)}|\theta) - \log p(w^{(n)}) - \log p(D|w^{(n)}) \quad (4)$$

161 where $w^{(n)}$ denotes the n th Monte Carlo sampling from the variational posterior $q(w^{(n)}|\theta)$.

162 Together with the local reparameterization method (explained in the supplementary
 163 material), which translates the global uncertainty in the weights into a form of local uncer-
 164 tainty (Kingma et al., 2015; Shridhar et al., 2019), our BayesLSTMs are ready for training
 165 and back-propagation. We constructed the networks using the Pytorch library, and our code
 166 is published on Github (<https://github.com/geek-yang/DLACs>).

167 2.3 Ensemble Forecasting with BDL and Numerical Configurations

168 The ensemble method is generally used for uncertainty assessment in weather forecast-
 169 ing (Gneiting et al., 2005; Buizza et al., 2008; Leutbecher et al., 2017). In numerical weather
 170 prediction systems (NWP), uncertainties in the initial conditions and model parameters are
 171 projected by ensemble forecasts with perturbations in the initial conditions and model for-
 172 mulations (Palmer, 2002; Milinski et al., 2020). It has been explained in many previous
 173 studies that BDL is able to address the uncertainties in initial conditions and model pa-
 174 rameters (e.g., Kendall & Gal, 2017). More details about the uncertainty estimation with
 175 BDL are provided in the supplementary material. This characteristic is fundamental for
 176 any probabilistic forecast and therefore makes BDL a candidate for weather forecasting.
 177 However, they are treated differently than in operational NWP approaches.

178 The ability of BayesLSTM to characterize uncertainty is reflected in its forecasting
 179 procedure. During a prediction process, the whole time series preceding the forecast date
 180 ($t < t_0$) will be fed to the model to initialize the memory and position the state of the
 181 network. Therefore the model itself is constrained by the past and this is similar to an
 182 NWP-based forecast. When producing a forecast that takes uncertainties into account for
 183 a next time step, the BayesLSTM will first sample the weight distributions multiple times
 184 to build an ensemble and then use the sampled weight matrix to generate the predictions

185 for the target time step ($t = t_1$). The ensemble forecast can be extended to more time steps
186 ahead ($t = t_n > t_1$) by continuing with each individual LSTM.

187 In order to evaluate ensemble forecasts with the BayesLSTM, several scores are cal-
188 culated, including continuous ranked probability score (CRPS), root mean square error
189 (RMSE) and Euclidean distance (EuD). The mathematical expressions of these scores can
190 be found in the supplementary material.

191 For all the experiments in this paper, we generate sequences including 1500 time steps
192 (250 days) with the Lorenz 84 model. The training set contains 1300 time steps (about 216
193 days) and the validation set consists of 200 time steps (about 33 days). The optimization is
194 based on the minimization of training loss, which consists of likelihood cost (data-dependent)
195 and complexity cost (prior dependent) (Shridhar et al., 2019). A scaling factor between these
196 two sources of loss should be tuned, since it accounts for the trade-off between the width
197 of ensemble spread in terms of uncertainty estimation and saturation of forecasts around
198 the variance displayed in the observations. Note that the scaling factor is related to the
199 normalization of the distributions and cannot be calculated exactly. The training time
200 is about 20 hours on a single GPU (Nvidia Tesla K40m). The hyperparameters like the
201 learning rate, number of epochs and number of layers, were tested and determined in terms
202 of the EuD error. It shows that a combination of a learning rate equal to 0.01, a single
203 BayesLSTM layer and 3000 epochs is sufficient to achieve satisfying results. The training
204 loss is shown in Figure S1. More details about the numerical configurations are shown in
205 the supplementary material.

206 2.4 Vector Autoregressive Model

207 The VAR model is used as a baseline method to assess the probabilistic forecast skill
208 of BayesLSTM. As a variant of the autoregressive model (AR), the VAR model generalizes
209 univariate AR by allowing for multivariate time series and therefore can capture the relation
210 between multiple variables. The VAR model and many variants belonging to the VAR family
211 haven shown skill in many weather forecast applications (e.g., Gneiting et al., 2006; L. Wang
212 et al., 2016, and many others). To expand its forecast capacity from the deterministic
213 domain to the probabilistic domain, we replaced the Gaussian noise term (ϵ_t) with Gaussian
214 distributed variations based on the variance of input time series from the chosen lag step to
215 the current step. The optimal number of the lag to be included in the model is determined
216 based on the auto-correlation of each variable of the Lorenz 84 model output (shown in
217 Figure S2 in the supplementary material), and tests of forecast quality in terms of the
218 CRPS score. In our case, the VAR model with a lag equal to 3 provides the best probabilistic
219 forecast. Mathematically, our modified VAR model can be expressed as:

$$\begin{aligned}
 X_t &= \alpha_1 + \sum_{l=1}^{Lag} (\beta_{11,l}X_{t-l} + \beta_{12,l}Y_{t-l} + \beta_{13,l}Z_{t-l}) + \epsilon_{1,t} \\
 Y_t &= \alpha_2 + \sum_{l=1}^{Lag} (\beta_{21,l}X_{t-l} + \beta_{22,l}Y_{t-l} + \beta_{23,l}Z_{t-l}) + \epsilon_{2,t} \\
 Z_t &= \alpha_3 + \sum_{l=1}^{Lag} (\beta_{31,l}X_{t-l} + \beta_{32,l}Y_{t-l} + \beta_{33,l}Z_{t-l}) + \epsilon_{3,t}
 \end{aligned} \tag{5}$$

with

$$\begin{aligned}
 \epsilon_{1,t} &= \mathcal{N}[0, \sigma(X_{t-1}, X_{t-2}, \dots, X_{t-l})^2] \\
 \epsilon_{2,t} &= \mathcal{N}[0, \sigma(Y_{t-1}, Y_{t-2}, \dots, Y_{t-l})^2] \\
 \epsilon_{3,t} &= \mathcal{N}[0, \sigma(Z_{t-1}, Z_{t-2}, \dots, Z_{t-l})^2]
 \end{aligned}$$

220 Where α and β are trainable parameters in the model, ϵ_t the Gaussian distributed variations,
 221 and X_{t-l} , Y_{t-l} and Z_{t-l} the Lorenz model output at time lag l . The parameters were
 222 updated by fitting the model to the time series of Lorenz model output using maximum
 223 likelihood.

224 3 Results

225 We evaluate the capacity of BDL in representing the dynamics of Rossby wave propa-
 226 gation on a westerly jet by investigating the forecasts in the spatial and temporal domains.
 227 Based on the selected scoring metrics, we further assess the forecast quality of BayesLSTM
 228 against the forecasts with persistence and a VAR model.

229 3.1 Representing the Evolution of Lorenz 84 Model

230 A retrospective forecast of the Lorenz 84 system with the BayesLSTM is shown in
 231 Figure 2. The forecasts start every time step (4 hours) and each has been extended to a
 232 lead time of 3 days. Given the time series of the BayesLSTM forecasts in Figure 2a, it can
 233 be observed that in general the forecasts are close to the time series of the Lorenz 84 model
 234 output, which is considered to be the "truth". Although the forecast quality drops down
 235 with the increase of lead time as expected, the BayesLSTM shows good skill in replicating
 236 the variations of the Lorenz 84 model, especially for the state-transitions of the Lorenz 84
 237 system and the sinusoidal patterns of the eddy components, like the forecast of X around
 238 valid date 14 and the forecast of Y around valid date 16. This indicates that the BayesLSTM
 239 learns to predict the state of the Lorenz system. Considering the typical predicting process
 240 of an LSTM network, in which the whole time series of the Lorenz 84 system preceding the
 241 forecast time should be fed to the system, it implies that our BayesLSTM is well constrained
 242 by the Lorenz 84 model output. Given the fact that the learning process of a deep neural
 243 network is characterized by the relationship between input fields, it further indicates that
 244 the non-linear relations between the variables in this Lorenz 84 system, the westerly X and
 245 the large scale eddies Y and Z , were addressed by the BayesLSTM.

246 In addition, we plot the forecast trajectory in Figure 2d and compare it with the Lorenz
 247 model output to further evaluate the performance of BayesLSTM. It can be noticed that the
 248 forecast trajectory is close to the attractor and the "behavior" of the forecast trajectory as a
 249 function of lead time resembles the evolution of the Lorenz 84 model. The result is consistent
 250 with the assessment based on the time series of each component as shown in Figure 2a. As a
 251 follow-up check, we investigate the physical consistency of BayesLSTM forecasts via the log
 252 power spectrum density of forecast time series, which is shown in Figure 2c. Only the high
 253 frequency components of X (with the frequency between 0.9 and 1.5) differ from the Lorenz
 254 model output. In general, the power spectrum density of the BayesLSTM forecasts is similar
 255 to that of the Lorenz 84 model. This indicates that the phases of the waves simulated by
 256 BayesLSTM do not differ much from the Rossby waves in the Lorenz 84 model. Considering
 257 the time step (4 hours) and the damping time of the Lorenz system (5 days), such similarity
 258 over the whole frequency space reflects that the BayesLSTM can account for the dynamics
 259 of this Rossby wave system across different time scales, which potentially benefits from its
 260 ability of multiple-level information abstraction. Together with the similar amplitudes of
 261 waves displayed in Figure 2a, it implies that the BayesLSTM manages to learn the Rossby
 262 wave propagation. The interaction between the jet and eddy components in this simplified
 263 atmospheric circulation system and the forecasts are physically realistic.

264 In order to evaluate the probabilistic forecast skill of the BayesLSTM, we generated
 265 a 20-member ensemble by sampling the BayesLSTM network and the time series of these
 266 retrospective forecasts up to 3 lead days are shown in Figure 2b. The blue shades serve
 267 to approximate the error growth of the Lorenz 84 system, which are selected as the range
 268 between the current Lorenz model series persisting for 3 lead and lag days. Note that this
 269 selection is made based on the auto-correlation in Figure S2 and it aims to assist the evalu-

270 ation of the probabilistic forecasts, specifically for the uncertainty estimation. It is observed
 271 that the forecast members are located around the Lorenz model output and the spread is
 272 comparable to the error growth of this Rossby wave system. This indicates that the spread of
 273 the BayesLSTM ensemble is neither over-dispersive nor under-dispersive. The probabilistic
 274 forecasts therefore address uncertainties in a reasonable way. Collectively, the development
 275 of these forecasts as a function of lead time in 2b are similar to the single forecast in 2a.
 276 This means almost all the ensemble members capture the properties of the propagating
 277 waves and the jet strength while allowing for the occurring of uncertainty. Consequently,
 278 the probabilistic forecasts generated by sampling the BayesLSTM are physically plausible.

279 Nevertheless, the BayesLSTM forecasts may lose skill at certain valid time. For in-
 280 stance, in Figure 2a and b between valid date 0 to 6, forecasts of X drift away unrealistically.
 281 This might result from the state-dependency of the BayesLSTM, or in general the state-
 282 dependency of any deep learning approaches. For a numerical model, it is common to have
 283 state-dependency, for example, the prediction of NAO/blocking events in medium-range
 284 forecasts (e.g., Parker et al., 2018). This may also apply to the deep learning approaches if
 285 the training data fails to provide adequate information for forecasting at some points.

286 3.2 Evaluate the BayesLSTM ensemble forecasts

287 A reliability assessment of probabilistic forecasts with the BayesLSTM ensemble was
 288 performed using the chosen metrics. The BayesLSTM ensemble consists of 20 members
 289 and they are evaluated against a deterministic forecast with persistence and a probabilistic
 290 forecast with the VAR model, which is also a 20-member ensemble. The results are shown
 291 in Figure 3. Regarding the CRPS score, in general the BayesLSTM ensemble forecast is
 292 better than the VAR ensemble forecast considering all the variables for almost all lead days.
 293 Only around day 1 for predictand X , the VAR ensemble forecast shows slightly better skill.
 294 The error growth of the BayesLSTM ensemble forecast is much slower than that of the VAR
 295 ensemble forecast. Given the definition of CRPS score, which provides a quadratic measure
 296 of discrepancy between the forecast cumulative density function (CDF) and the empirical
 297 CDF of the scalar observation (Gneiting et al., 2005), this indicates that the forecast CDF
 298 with the BayesLSTM centered around the Lorenz model output, while the forecast CDF
 299 with the VAR is relatively over-dispersive.

300 Regarding the RMSE shown in Figure 3b, forecasts with persistence are better than
 301 that with BayesLSTM and VAR ensemble concerning only X . This is consistent with
 302 the high auto-correlation of the zonal wind X shown in Figure S2. While for the eddy
 303 components Y and Z , the BayesLSTM provides much better forecasts within 3 lead days,
 304 with the averaged RMSE error smaller than the standard deviation of the full time series
 305 of the Lorenz 84 model output. Considering the nonlinear relation between the westerly X
 306 and large scales eddies Y and Z , this means that the BayesLSTM is able to preserve the
 307 physical consistency between the zonal wind and the propagation of large scale eddies in
 308 this atmospheric circulation system, and therefore produces better probabilistic forecasts.
 309 It is evident by analyzing the time series in Figure 2a, that the variations of Y and Z are
 310 well represented by the BayesLSTM forecasts up to a lead time of 3 days.

311 More information about forecast quality in terms of the trajectories, which intrinsically
 312 embody the properties of Rossby waves and jet strength, is reflected by the EuD in Figure 3c.
 313 Starting from the first forecast time step (4 hours), the BayesLSTM shows better forecast
 314 skill concerning the EuD. Although the EuD error grows with the increase of lead time for
 315 all the forecast methods, the BayesLSTM forecasts are better than the others for the whole
 316 inspected lead time range. Note that within 2 lead days, the EuD error of BayesLSTM is
 317 smaller than the standard deviation of the Lorenz 84 model output, which is about 0.6.
 318 Since the EuD of BayesLSTM ensemble forecast shown in Figure 3c is the average of 20
 319 members with forecasts starting every time step, this implies that these ensemble members
 320 are able to replicate the patterns of the attractor and the spread of the ensemble is properly

distributed around the target Lorenz 84 model trajectory. It further demonstrates that probabilistic forecast with BayesLSTM can address uncertainties adequately.

4 Discussion

We demonstrate the capability of BayesLSTM in probabilistic weather forecasting. Intuitively, by perturbing the Lorenz 84 model, it seems possible to compare the BayesLSTM forecasts to the perturbed Lorenz model output and check if the BayesLSTMs are able to address uncertainties in the initial conditions and model formulation, respectively. However, there is no objective way to determine the amplitude of the perturbation which can appropriately approximate the error growth in the Lorenz system that is analogue to a realistic dynamical system of Rossby waves on a jet. So this experiment is not feasible at the moment.

In addition, we extended the ensemble forecasts to more than 60 lead days and noticed that after 20 days, the forecast errors increase dramatically with the increase of lead time (not shown). From this point, it seems that the BayesLSTM is useful for medium-range forecasts and it is not suitable for seasonal forecast and climate change predictions. The outcome of this study is insufficient to prove that, either the Bayesian deep neural networks can mathematically represent the differential equations which depict the Lorenz 84 system with seasonal forcing (note that due to the features of deep learning and the nature of deep neural networks, there is no direct mapping between weight matrix in a trained BayesLSTM and Lorenz model parameters), or BayesLSTMs only abstract and store the physical linkages in a latent space and use them to produce memory-based forecasts at relatively short time scales. This can be explored in the future.

Although not the main topic of this paper, we note that the formulations of BDL are very similar to data assimilation, specifically the Bayesian data assimilation, which is extensively used in weather forecasting to combine the knowledge from observations and models, and deal with the uncertainty in the initial conditions (Evensen, 1994; P. L. Houtekamer & Mitchell, 1998; P. Houtekamer & Zhang, 2016). Based on the Bayes' theorem, it incorporates model knowledge into the prior and corresponding observations as likelihood, and treats the observation involving uncertainty estimation as posterior. Given the large dimensional systems, in reality approximate solutions are always made based on different methods, like variational methods, Kalman-based methods and particle filters (Navon, 2009).

Given the fact that forecasts with BayesLSTM stay close to the Lorenz 84 attractor, the BayesLSTM may be also chaotic. This question can be answered by the chaotic system diagnostics, for instance, with the Lyapunov spectrum (Broer et al., 2002; Freire et al., 2008). However, this is beyond our scope now but worth the effort in the future. So far, it can be concluded that BayesLSTM is a useful candidate for weather forecasts, at relatively small lead times up to several days. For a long term climate forecast, the BayesLSTM may not be a good choice in terms of the error accumulation and the lack of skill in physical model representation. Also, for simulating and forecasting changes in the climate system boundary condition uncertainty will need to be taken into account. This can be further tested by studies using observational data and climate model ensembles in the future.

5 Conclusion

In this study, we explored the potential of BDL for weather forecasting using the modified Lorenz 84 model as a model for the atmosphere. The probabilistic character of the BDL is addressed and assessed using the chaotic nature of the Lorenz 84 system with seasonal forcing as 'truth'. Specifically, we chose BayesLSTM as an example of BDL to forecast the Lorenz 84 model and evaluate its forecast skill. It was observed that the retrospective forecasts are similar to those of the Lorenz model output in the spatial and temporal domain. The forecast trajectories are close to the attractor. This indicates that

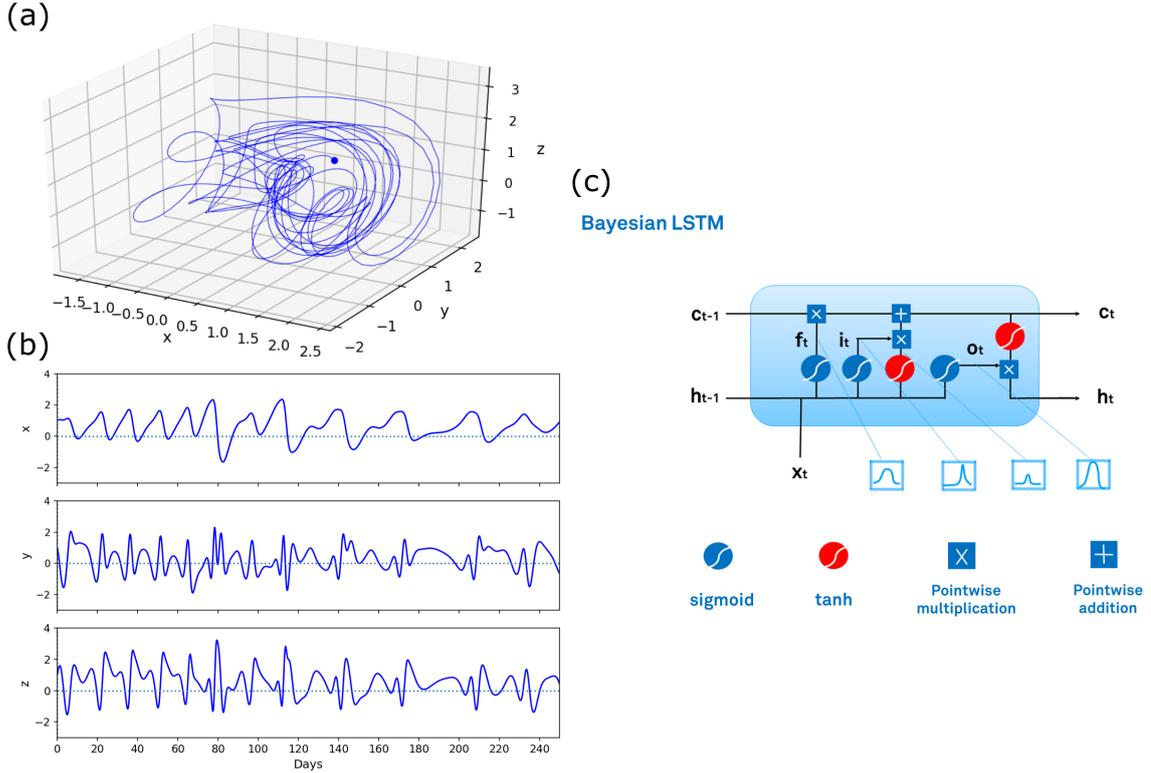


Figure 1. (a) Trajectory and (b) time series of each variable of the Lorenz 84 model with seasonal forcing. The sequences contain 250 days (1500 time steps) and the starting point is marked with a blue dot. (c) Structure of the Bayesian Long-Short Term Memory neural networks (Fortunato et al., 2017).

370 BayesLSTM is able to learn the propagation of Rossby waves in this atmospheric system,
 371 in terms of both the amplitude and phase. It further demonstrates that the BayesLSTM
 372 is able to replicate the interaction between the jet stream and large-scale eddies and thus
 373 the evolution of Rossby waves on a midlatitude jet. The forecasts get worse with increasing
 374 lead times due to the accumulation of errors, as expected.

375 The probabilistic forecast skill of BayesLSTM was analyzed and evaluated against
 376 persistence and a VAR model. We found that the BayesLSTM forecasts saturate around
 377 the model output considering both the sequences of each variable and the trajectory. In
 378 terms of the scores in the chosen metrics, the BayesLSTM shows better probabilistic forecast
 379 skill than persistence and the VAR model in the inspected lead days. It shows that the
 380 BayesLSTM is able to account for uncertainties relevant to the evolution of this simplified
 381 atmospheric circulation system, though the procedure differs from well-known NWP based
 382 approaches. Given the relatively low cost of ensemble forecasts compared to deterministic
 383 DNN and NWP systems, and the capacity in probabilistic forecasting, BayesLSTM, or
 384 in general BDL, is useful to produce fast and reliable probabilistic weather forecast and
 385 therefore is promising to enhance weather forecasting capabilities at short to medium-range
 386 timescales.

387 **Acknowledgments**

388 The authors gratefully acknowledge the support by the Netherlands eScience Center and

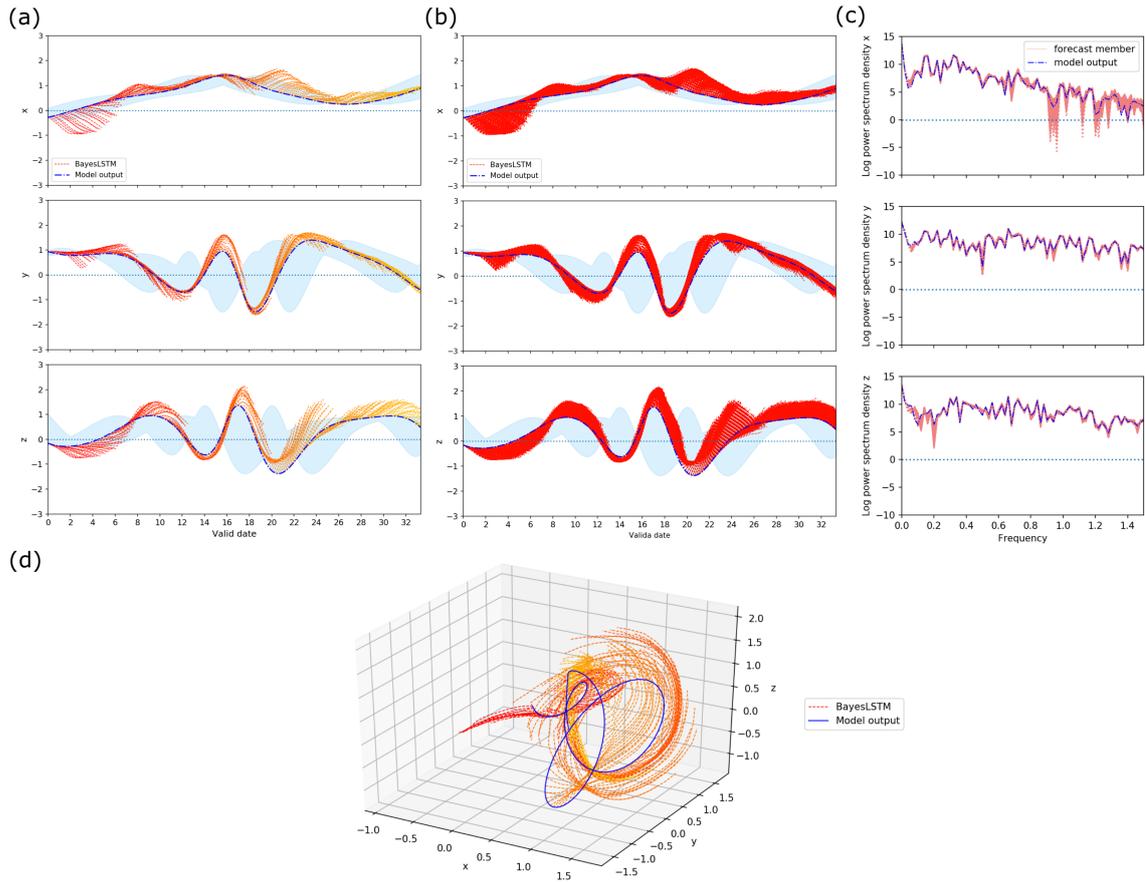


Figure 2. BayesLSTM retrospective forecasts up to a lead time of 3 days (18 time steps), with forecasts starting every time step (every 4 hours). (a) Time series of each variable (b) time series of a 20-member ensemble (c) logarithmic power spectrum and (d) trajectory in phase space. Except for (b) all the figures contain the results from a single BayesLSTM retrospective forecast. The Lorenz model output is included as reference (blue, labelled as "model output") and the blue shades indicate the range between the Lorenz model output persisting for 3 days, both lead (3 days forward) and lag (3 days backward).

389 Wageningen University. This study is supported by Blue Action project (European Union's
 390 Horizon 2020 research and innovation programme, grant number: 727852). We would like to
 391 thank SURFsara (Netherlands) for providing us their super computing infrastructure for our
 392 project. All data analyzed here are publicly available. The forecasts from BayesLSTM net-
 393 work can be accessed at a public github repository (<https://doi.org/10.5281/zenodo.4494116>,
 394 in the data folder), which also contains the forecasts from VAR model. The Lorenz 84
 395 model output is available through Lorenz (1984), with modifications described in the sec-
 396 tion Methodology.

397 References

398 Bannister, R. (2017). A review of operational methods of variational and ensemble-
 399 variational data assimilation. *Quarterly Journal of the Royal Meteorological Society*,
 400 *143*(703), 607–633.

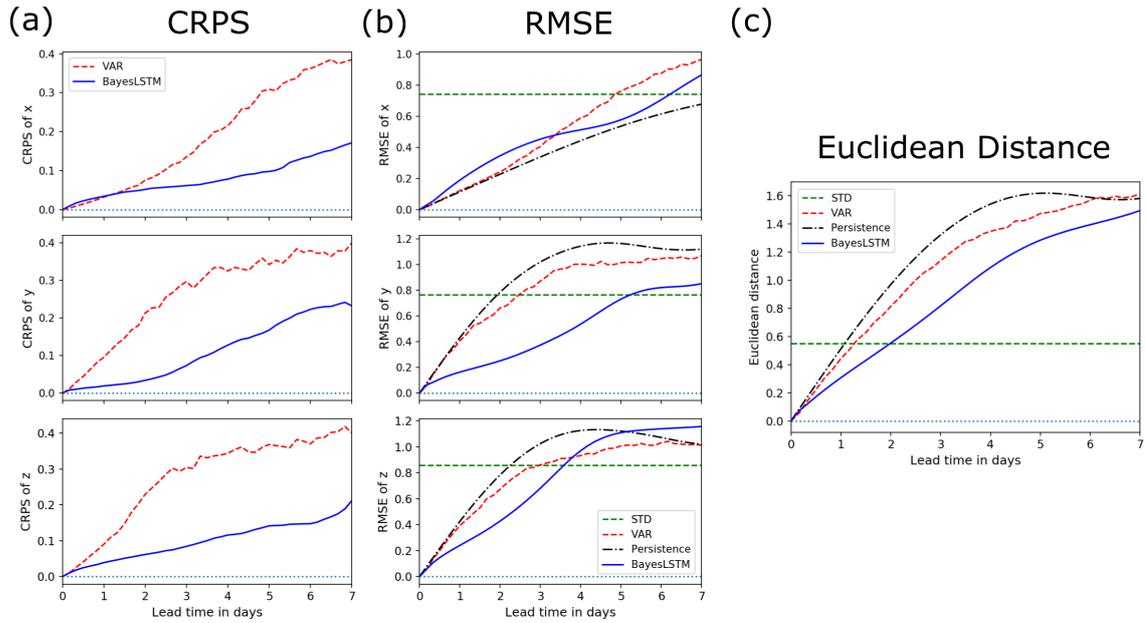


Figure 3. Skill evaluation of the BayesLSTM ensemble forecasts against VAR and persistence with (a) CRPS and (b) RMSE and (c) EuD, which are averaged over 200 forecasts starting every time step (4 hours). The standard deviation of the full time series of Lorenz model output (based on 250 days data) is included in (b) and (c) (green, labelled as "STD").

- 401 Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015). Weight uncertainty
 402 in neural networks. *arXiv preprint arXiv:1505.05424*.
- 403 Broer, H., Simó, C., & Vitolo, R. (2002). Bifurcations and strange attractors in the Lorenz-84
 404 climate model with seasonal forcing. *Nonlinearity*, 15(4), 1205.
- 405 Buizza, R., Leutbecher, M., & Isaksen, I. (2008). Potential use of an ensemble of analyses
 406 in the ECMWF ensemble prediction system. *Quarterly Journal of the Royal Meteorological
 407 Society: A journal of the atmospheric sciences, applied meteorology and physical
 408 oceanography*, 134(637), 2051–2066.
- 409 Chattopadhyay, A., Hassanzadeh, P., Palem, K., & Subramanian, D. (2019). Data-driven
 410 prediction of a multi-scale Lorenz 96 chaotic system using a hierarchy of deep learning
 411 methods: Reservoir computing, ANN, and RNN-LSTM. *arXiv preprint arXiv:1906.08829*.
- 412 Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model
 413 using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research:
 414 Oceans*, 99(C5), 10143–10162.
- 415 Fortunato, M., Blundell, C., & Vinyals, O. (2017). Bayesian recurrent neural networks.
 416 *arXiv preprint arXiv:1704.02798*.
- 417 Freire, J. G., Bonatto, C., DaCamara, C. C., & Gallas, J. A. (2008). Multistability, phase
 418 diagrams, and intransitivity in the Lorenz-84 low-order atmospheric circulation model.
 419 *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 18(3), 033121.
- 420 Ghil, M., & Malanotte-Rizzoli, P. (1991). Data assimilation in meteorology and oceanogra-
 421 phy. In *Advances in geophysics* (Vol. 33, pp. 141–266). Elsevier.
- 422 Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration
 423 and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Method-
 424 ology)*, 69(2), 243–268.
- 425 Gneiting, T., Larson, K., Westrick, K., Genton, M. G., & Aldrich, E. (2006). Calibrated
 426 probabilistic forecasting at the stateline wind energy center: The regime-switching

- 427 space–time method. *Journal of the American Statistical Association*, 101(475), 968–
 428 979.
- 429 Gneiting, T., Raftery, A. E., Westveld III, A. H., & Goldman, T. (2005). Calibrated
 430 probabilistic forecasting using ensemble model output statistics and minimum crps
 431 estimation. *Monthly Weather Review*, 133(5), 1098–1118.
- 432 Graves, A. (2011). Practical variational inference for neural networks. In *Advances in neural
 433 information processing systems* (pp. 2348–2356).
- 434 Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*,
 435 9(8), 1735–1780.
- 436 Houtekamer, P., & Zhang, F. (2016). Review of the ensemble kalman filter for atmospheric
 437 data assimilation. *Monthly Weather Review*, 144(12), 4489–4532.
- 438 Houtekamer, P. L., & Mitchell, H. L. (1998). Data assimilation using an ensemble kalman
 439 filter technique. *Monthly Weather Review*, 126(3), 796–811.
- 440 Kendall, A., & Gal, Y. (2017). What uncertainties do we need in bayesian deep learning
 441 for computer vision? In *Advances in neural information processing systems* (pp.
 442 5574–5584).
- 443 Kingma, D. P., Salimans, T., & Welling, M. (2015). Variational dropout and the local
 444 reparameterization trick. In *Advances in neural information processing systems* (pp.
 445 2575–2583).
- 446 LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, 521(7553), 436–444.
- 447 Leutbecher, M., Lock, S.-J., Ollinaho, P., Lang, S. T., Balsamo, G., Bechtold, P., ... others
 448 (2017). Stochastic representations of model uncertainties at ecmwf: State of the art
 449 and future vision. *Quarterly Journal of the Royal Meteorological Society*, 143(707),
 450 2315–2339.
- 451 Leutbecher, M., & Palmer, T. N. (2008). Ensemble forecasting. *Journal of computational
 452 physics*, 227(7), 3515–3539.
- 453 Liu, Y., Bogaardt, L., Attema, J., & Hazeleger, W. (2020). Extended range arctic sea
 454 ice forecast with convolutional long-short term memory networks. *Monthly Weather
 455 Review*. (under review)
- 456 Lorenz, E. N. (1984). Irregularity: A fundamental property of the atmosphere. *Tellus A*,
 457 36(2), 98–110.
- 458 Milinski, S., Maher, N., & Olonscheck, D. (2020). How large does a large ensemble need to
 459 be. *Earth System Dynamics*, 11, 885–901.
- 460 Navon, I. M. (2009). Data assimilation for numerical weather prediction: a review. In
 461 *Data assimilation for atmospheric, oceanic and hydrologic applications* (pp. 21–65).
 462 Springer.
- 463 Neal, R. M., & Hinton, G. E. (1998). A view of the em algorithm that justifies incremental,
 464 sparse, and other variants. In *Learning in graphical models* (pp. 355–368). Springer.
- 465 Palmer, T. (2002). Predicting uncertainty in numerical weather forecasts. In *International
 466 geophysics* (Vol. 83, pp. 3–13). Elsevier.
- 467 Parker, T., Woollings, T., & Weisheimer, A. (2018). Ensemble sensitivity analysis of
 468 greenland blocking in medium-range forecasts. *Quarterly Journal of the Royal Mete-
 469 orological Society*, 144(716), 2358–2379.
- 470 Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., et al.
 471 (2019). Deep learning and process understanding for data-driven earth system science.
 472 *Nature*, 566(7743), 195–204.
- 473 Scher, S., & Messori, G. (2018). Predicting weather forecast uncertainty with machine
 474 learning. *Quarterly Journal of the Royal Meteorological Society*, 144(717), 2830–2841.
- 475 Shridhar, K., Laumann, F., & Liwicki, M. (2018). Uncertainty estimations by softplus
 476 normalization in bayesian convolutional neural networks with variational inference.
 477 *arXiv preprint arXiv:1806.05978*.
- 478 Shridhar, K., Laumann, F., & Liwicki, M. (2019). A comprehensive guide to bayesian convo-
 479 lutional neural network with variational inference. *arXiv preprint arXiv:1901.02731*.
- 480 Vandal, T., Kodra, E., Dy, J., Ganguly, S., Nemani, R., & Ganguly, A. R. (2018). Quanti-
 481 fying uncertainty in discrete-continuous and skewed data with bayesian deep learning.

- 482 In *Proceedings of the 24th acm sigkdd international conference on knowledge discovery*
483 & *data mining* (pp. 2377–2386).
- 484 Wang, H., Yu, Y., & Wen, G. (2014). Dynamical analysis of the lorenz-84 atmospheric
485 circulation model. *Journal of Applied Mathematics, 2014*.
- 486 Wang, H.-z., Li, G.-q., Wang, G.-b., Peng, J.-c., Jiang, H., & Liu, Y.-t. (2017). Deep
487 learning based ensemble approach for probabilistic wind power forecasting. *Applied*
488 *energy, 188*, 56–70.
- 489 Wang, L., Yuan, X., Ting, M., & Li, C. (2016). Predicting summer arctic sea ice concentra-
490 tion intraseasonal variability using a vector autoregressive model. *Journal of Climate,*
491 *29(4)*, 1529–1543.
- 492 Zaier, I., Shu, C., Ouarda, T., Seidou, O., & Chebana, F. (2010). Estimation of ice
493 thickness on lakes using artificial neural network ensembles. *Journal of Hydrology,*
494 *383(3-4)*, 330–340.