
Addressing Deep Learning Model Uncertainty in Long-Range Climate Forecasting with Late Fusion

Ken C. L. Wong*
IBM Research
San Jose, CA, USA

Hongzhi Wang
IBM Research
San Jose, CA, USA

Etienne E. Vos
IBM Research
Johannesburg, GP, South Africa

Bianca Zadrozny
IBM Research
Rio De Janeiro, RJ, Brazil

Campbell D. Watson
IBM Research
Yorktown Heights, NY, USA

Tanveer Syeda-Mahmood
IBM Research
San Jose, CA, USA

Abstract

Global warming leads to the increase in frequency and intensity of climate extremes that cause tremendous loss of lives and property. Accurate long-range climate prediction allows more time for preparation and disaster risk management for such extreme events. Although machine learning approaches have shown promising results in long-range climate forecasting, the associated model uncertainties may reduce their reliability. To address this issue, we propose a late fusion approach that systematically combines the predictions from multiple models to reduce the expected errors of the fused results. We also propose a network architecture with the novel denormalization layer to gain the benefits of data normalization without actually normalizing the data. The experimental results on long-range 2m temperature forecasting show that the framework outperforms the 30-year climate normals, and the accuracy can be improved by increasing the number of models.

1 Introduction

Global warming leads to the increase in frequency and intensity of climate extremes [7]. High-impact extreme events such as heat waves, cold fronts, floods, droughts, and tropical cyclones can result in tremendous loss of lives and property, and accurate predictions of such events benefit multiple sectors including water, energy, health, agriculture, and disaster risk reduction [9]. The longer the range of an accurate prediction, the more the time for proper preparation and response. Therefore, accurate long-range forecasting of the key climate variables such as precipitation and temperature is valuable.

Numerical models for weather and climate prediction have a long history of producing the most accurate seasonal and multi-annual climate forecasts, but they come with the cost of large and expensive physics-based simulations (e.g. [6, 1]). With the recent advancements in machine learning such as deep learning, the use of machine learning for climate forecasting has become more popular [3, 14, 12, 11], and some machine learning approaches can outperform numerical models in certain tasks [3]. Nevertheless, depending on the machine learning algorithm and data availability, different degrees of model uncertainties exist. In deep learning, models trained with the same data and hyperparameters are usually not identical. This is caused by the random processes in training such as weight initialization and data shuffling. Such model uncertainties can be more prominent for climate forecasting given the limited data, and this can reduce the reliability of the models especially with large lead times. Even though reducing the randomness in training (e.g., using fixed weight initialization) may reduce the model uncertainties, the chance of getting better models is also reduced.

*Corresponding author (c1wong@us.ibm.com)

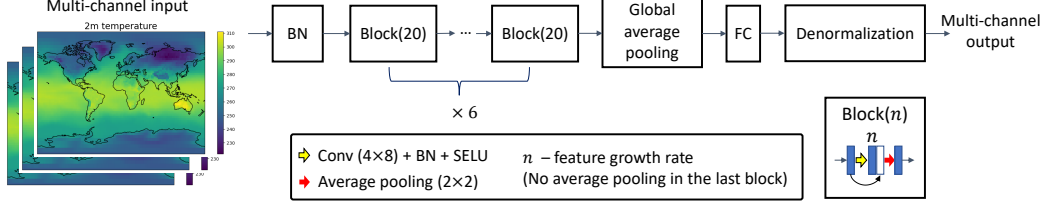


Figure 1: Network architecture for 2m temperature forecasting. BN and FC represent batch normalization and fully connected layers, respectively. The number of input channels is the input horizon, and the number of output channels is the number of predictand locations.

In this paper, our goal is to reduce model uncertainties and improve accuracy in seasonal climate forecasting. By modifying the late fusion approach in [13] to adapt to deep learning regression, predictions from different models trained with identical hyperparameters are systematically combined to reduce the expected errors in the fused results. We demonstrate its applicability on long-range 2m temperature forecasting. Furthermore, we propose a novel denormalization layer which allows us to gain the benefits of data normalization without actually normalizing the data.

2 Methodology

2.1 Network Architecture with Denormalization

The proposed network architecture is shown in Fig. 1. Given a multi-channel input tensor formed by stacking the input maps of 2m temperature spanning a fixed input horizon, the network predicts the 2m temperatures at multiple locations with a fixed lead time. The network comprises six dense blocks [5], each with a convolutional layer and a growth rate of 20. A batch normalization layer is used right after the input layer for data normalization. Furthermore, although we found that normalizing the predictands allows the use of simpler architectures with better accuracy, the resulting model can only provide normalized predictions and postprocessing is required to recover the original values. To address this issue, we introduce a *denormalization* layer after the final fully connected layer to obtain:

$$x_o(c) = x_i(c)\sigma(c) + m(c) \quad (1)$$

with c the channel index, and x_o , x_i the output and input features, respectively. σ and m are the standard deviation and mean value computed from the training data. Using this denormalization layer, the final fully connected layer only needs to predict normalized values, thus removing the need of predictand normalization. With this architecture, data normalization in training and forecast denormalization in inference are unnecessary.

2.2 Late Fusion

We modified the late fusion approach in [13] for regression. The method combines predictions from multiple models using weighted average. To compute the weights, the pairwise correlations between different models in terms of how likely they will make correlated errors are estimated, which are then used to compute the weights that reduce the expected error in the fused result. Let $f^j(s_i)$ be the prediction by the j_{th} model for input s_i and $t(s_i)$ the true value. The late fusion result for s_i is $\sum_j w^j f^j(s_i)$ with $\sum_j w^j = 1$. The pairwise correlation between model j_1 and j_2 is:

$$M[j_1, j_2] = \sum_i [f^{j_1}(s_i) - t(s_i)] [f^{j_2}(s_i) - t(s_i)] \quad (2)$$

Then the weights are computed by:

$$\mathbf{w} = \arg \min_{\mathbf{w}} \mathbf{w}^T \mathbf{M} \mathbf{w} = \frac{\mathbf{M}^{-1} \mathbf{1}_K}{\mathbf{1}_K^T \mathbf{M}^{-1} \mathbf{1}_K} \quad (3)$$

with K the number of models and $\mathbf{1}_K$ a vector with K ones. \mathbf{M} and \mathbf{w} are computed using the validation data. This procedure is applied on each output channel.

Table 1: Locations at low or high latitudes where the 2m temperatures are predicted.

Low	Honolulu (21.3°N, 157.9°W), Panama City (9.0°N, 79.5°W), Singapore (1.4°N, 103.8°E), Mid Pacific Ocean (4.4°N, 167.7°W)
High	Moscow (55.8°N, 37.6°E), London (51.5°N, 0.1°W), Christchurch (43.5°S, 172.6°E), Perth (32.0°S, 115.9°E)

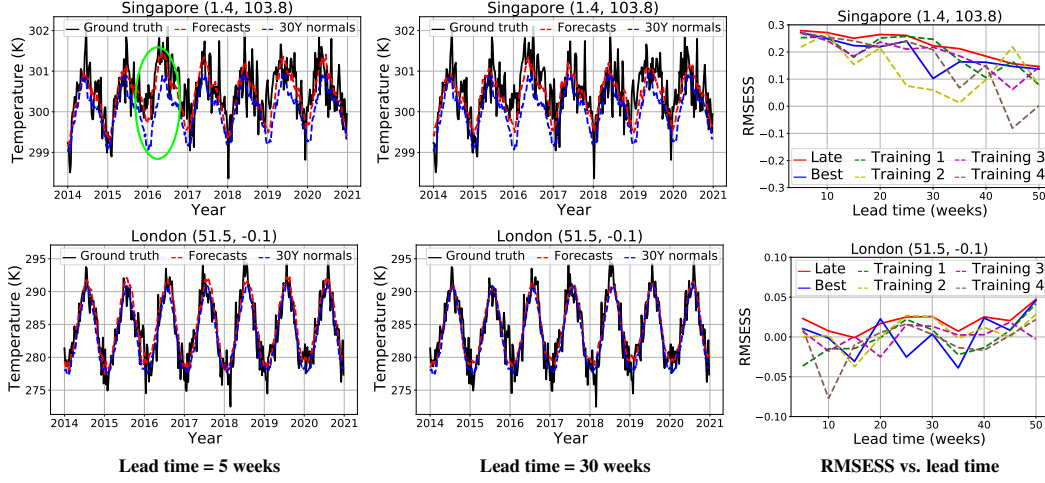


Figure 2: Left two: examples of forecasts at different lead times, with the green circle highlighting the hottest year (2016) on record. Right: RMSESS of models trained with identical hyperparameters (dashed lines) compared with the late fusion and the best model frameworks with 20 models per lead time (solid lines). Top: Singapore. Bottom: London.

2.3 Training Strategy

The 2m temperature maps of the ERA5 reanalysis data [4] were partitioned for training (1979 – 2007), validation (2008 – 2011), and testing (2012 – 2020). Each data map was resampled from the original spatial resolution of $0.25^\circ \times 0.25^\circ$ to $1^\circ \times 1^\circ$. The data were also aggregated over time from hourly to weekly. An input horizon of six weeks was used with 10 forecast lead times (5 to 50 weeks with a stride of 5 weeks). Each model was trained for 200 epochs with the batch size of 32. The Nadam optimizer [2] was used with the cosine annealing learning rate scheduler [8], with the minimum and maximum learning rates as 10^{-4} and 10^{-2} , respectively. The mean absolute error was used as the loss function.

3 Experiments

To study model uncertainties, we trained 20 models with identical hyperparameters per lead time (i.e., 200 models in total). Each model was used to predict temperatures from four low-latitude and four high-latitude locations (Table 1). Two frameworks were compared:

- **Late fusion:** the framework that combines the predictions of different models at each lead time.
- **Best model:** at each lead time, the model with the smallest root mean square error (RMSE) on the validation data was chosen to provide the predictions.

For evaluation, the RMSE skill score ($\text{RMSESS} \in [-\infty, 1]$) that compares between the model forecasts and the 30-year climate normals was used:

$$\text{RMSESS} = 1 - \frac{\text{RMSE}_{\text{model}}}{\text{RMSE}_{\text{clim}}} \quad (4)$$

with $\text{RMSE}_{\text{model}}$ computed between the forecasts and the true values, and $\text{RMSE}_{\text{clim}}$ computed between the 30-year climate normals and the true values. A 30-year climate normal is the 30-year average of a predictand at a given time point, which is a generally accepted benchmark for comparison.

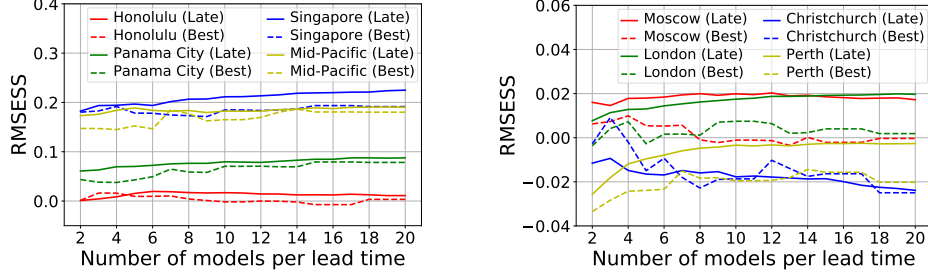


Figure 3: Comparison between the late fusion and the best model frameworks. The y-axis shows the average RMSESS over the lead times. Left to right: low-latitude and high-latitude locations.

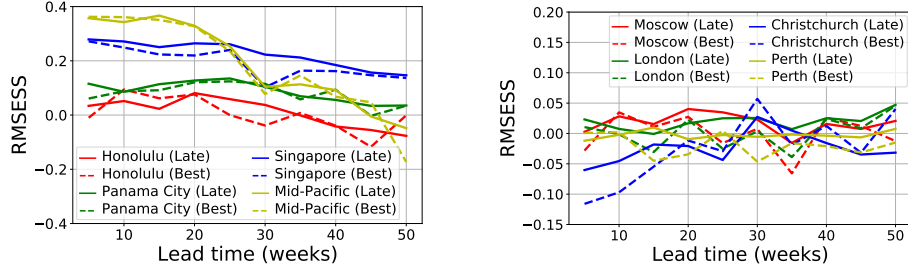


Figure 4: Performance of the late fusion and best model frameworks with 20 models per lead time.

3.1 Results

Fig. 2 (left) shows examples of forecasts on the testing data before applying the frameworks. In Singapore, with a lead time of five weeks, the forecasts closely followed the ground truth and outperformed the climate normals especially in 2016. In fact, 2016 was the hottest year on record [10] and the proposed model was able to forecast this anomalous event. However, as expected, the accuracy decreased with the increase of the lead time. In London, both the forecasts and the climate normals were very similar to the ground truth regardless of the lead time, probably because of the larger range in temperature.

In the RMSESS plot of Singapore in Fig. 2 (right), the mostly positive scores indicate that the forecasts outperformed the climate normals, though the scores decreased when the lead time increased. In London, the forecasts and climate normals were very similar, and the discrepancies among models were less obvious. Both plots show that although identical hyperparameters were used in training, the models performed differently especially with large lead times. By combining these models, the late fusion framework outperformed the best model framework and had the best overall results.

Fig. 3 shows comparison between the late fusion and the best model frameworks. The late fusion framework outperformed the best model framework in general. When the number of models per lead time increased, the late fusion framework improved smoothly in most locations and gradually converged with around 16 models. In contrast, the best model framework performed less well and may not benefit from a larger number of models. This is because the late fusion framework systematically reduced the expected errors from all models, while the best model framework only chose a single model that had the overall minimal RMSE on the validation data. Fig. 4 compares the two frameworks with 20 models per lead time. The late fusion framework outperformed the best model framework at most locations.

4 Conclusion

The results show that the models trained by the proposed architecture and training strategy can forecast large deviations from climate normals that attribute to climate change. Nevertheless, the models trained with identical hyperparameters may perform differently especially with large lead times. Using the late fusion approach, predictions from different models are combined systematically

to provide forecasts with reduced expected errors, and the results can be better than using a single model with the least validation error. As late fusion also improves forecasts with large lead times which associate with large model uncertainties, it is valuable for long-range climate forecasting.

References

- [1] Takeshi Doi, Swadhin K Behera, and Toshio Yamagata. Improved seasonal prediction using the SINTEX-F2 coupled model. *Journal of Advances in Modeling Earth Systems*, 8(4):1847–1867, 2016.
- [2] Timothy Dozat. Incorporating Nesterov momentum into Adam. In *ICLR Workshop*, 2016.
- [3] Yoo-Geun Ham, Jeong-Hwan Kim, and Jing-Jia Luo. Deep learning for multi-year ENSO forecasts. *Nature*, 573:568–572, 2019.
- [4] Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.
- [5] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4700–4708, 2017.
- [6] Stephanie J Johnson, Timothy N Stockdale, Laura Ferranti, Magdalena A Balmaseda, Franco Molteni, Linus Magnusson, Steffen Tietsche, Damien Decremier, Antje Weisheimer, Gianpaolo Balsamo, et al. SEAS5: the new ECMWF seasonal forecast system. *Geoscientific Model Development*, 12(3):1087–1117, 2019.
- [7] Timothy M Lenton, Johan Rockström, Owen Gaffney, Stefan Rahmstorf, Katherine Richardson, Will Steffen, and Hans Joachim Schellnhuber. Climate tipping points — too risky to bet against. *Nature*, 575:592–595, 2019.
- [8] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017.
- [9] William J Merryfield, Johanna Baehr, Lauriane Batté, Emily J Becker, Amy H Butler, Caio AS Coelho, Gokhan Danabasoglu, Paul A Dirmeyer, Francisco J Doblas-Reyes, Daniela IV Domeisen, et al. Current and emerging developments in subseasonal to decadal prediction. *Bulletin of the American Meteorological Society*, 101(6):E869–E896, 2020.
- [10] NOAA. 2020 was Earth’s 2nd-hottest year, just behind 2016. <https://www.noaa.gov/news/2020-was-earth-s-2nd-hottest-year-just-behind-2016>. Accessed: September 13, 2021.
- [11] Eduardo Rodrigues, Bianca Zadrozny, Campbell Watson, and David Gold. Decadal forecasts with ResDMD: a residual DMD neural network. In *ICML 2021 Workshop on Tackling Climate Change with Machine Learning*, 2021.
- [12] Etienne E Vos, Ashley Gritzman, Sibusisiwe Makhanya, Thabang Mashinini, and Campbell D Watson. Long-range seasonal forecasting of 2m-temperature with machine learning. In *NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning*, 2020.
- [13] Hongzhi Wang, Vaishnavi Subramanian, and Tanveer Syeda-Mahmood. Modeling uncertainty in multi-modal fusion for lung cancer survival analysis. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1169–1172. IEEE, 2021.
- [14] Meng-Hua Yen, Ding-Wei Liu, Yi-Chia Hsin, Chu-En Lin, and Chii-Chang Chen. Application of the deep learning for the prediction of rainfall in Southern Taiwan. *Scientific Reports*, 9:1–9, 2019.