
Meta-modeling strategy for data-driven forecasting

Dominic J. Skinner
Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139-4307
dskinner@mit.edu

Romit Maulik
Argonne Leadership Computing Facility
Argonne National Laboratory
Lemont, IL 60439
rmaulik@anl.gov

Abstract

Accurately forecasting the weather is a key requirement for climate change mitigation. Data-driven methods offer the ability to make more accurate forecasts, but lack interpretability and can be expensive to train and deploy if models are not carefully developed. Here, we make use of two historical climate data sets and tools from machine learning, to accurately predict temperature fields. Furthermore, we are able to use low fidelity models that are cheap to train and evaluate, to selectively avoid expensive high fidelity function evaluations, as well as uncover seasonal variations in predictive power. This allows for an adaptive training strategy for computationally efficient geophysical emulation.

1 Introduction

While numerical weather forecasting dates back over a century [10], climate change is expected to significantly alter the predictability of the atmosphere, increasing the error of weather forecasts [18]. Simultaneously, climate mitigation requires more accurate forecasts of weather events for power grid optimization [1], and extreme event prediction [20], from flooding [9], to heatwaves [6]. Even the largest climate simulations have $O(10\text{ km})$ separation between grid points, meaning dynamics of smaller scales can not be explicitly resolved [14, 17, 15], and often ad-hoc closure conditions are postulated to account for the unresolved dynamics. Machine learning has the potential to find data-driven closure conditions [16] and parameterize sub-grid scale modeling [3, 7]. Convolutional neural networks can forecast weather [4], deep learning can predict extreme weather events [5], and neural network architectures can be optimized automatically to enhance the quality of forecasts [11, 2]. However, these models can become expensive to train and deploy, especially when ensemble predictions are required [8], and often their predictions lack interpretability. Here, we use a combination of proper orthogonal decomposition (POD), and long-short term memory (LSTM) recurrent neural networks to create forecasts for two real-world data sets. In addition to the LSTM network, which is our high-fidelity (HF) model, we introduce simpler or low-fidelity (LF) models, which do not have the same predictive power, but are faster to deploy. By studying the prediction differences of the LF models, we are able to selectively avoid HF function evaluations, as well as uncovering seasonal variations in prediction accuracy. With the proliferation of data driven methods for climate forecasting, the approach introduced here, could be used to reduce computational resources, which is of particular relevance for ensemble forecasting [8], as well as to identify uncertainties in neural network predictions.

2 Proper Orthogonal Decomposition

To first reduce the dimensionality of the problem, we project spatially resolved fields, such as temperature, onto a set of principal modes which capture the salient features and then track the evolution of these modes in time. We use the technique of Proper Orthogonal Decomposition (POD)

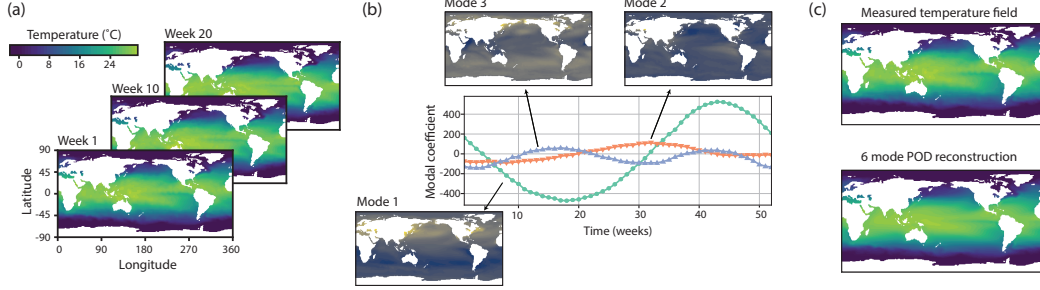


Figure 1: Proper orthogonal decomposition of sea surface temperature. (a) Historical climate data measuring ocean surface temperature weekly from 1981 to 2000. (b) Evolution of the first 3 modal coefficients, together with their respective basis elements. (c) Reconstructed temperature field from first the 6 modal coefficients against measured temperature field.

in order to find the dominant modes [21, 12, 13]. In short, suppose we are given spatial snapshots of the data at various times, $\theta_1, \dots, \theta_T$. We take a truncated set of orthonormal basis vectors v_1, \dots, v_M , that approximates the spatial snapshots optimally with respect to the L_2 norm, i.e. minimizing $\sum_i^T \|\theta_i - \phi_i\|^2$ over $\phi_i \in \text{Span}\{v_1, \dots, v_M\}$. Defining the snapshot matrix, $S = [\theta_1 | \dots | \theta_T]$, the optimal basis is given by the M eigenvectors of SS^T , with largest eigenvalues [21], which is readily found numerically. See appendix A, for further details.

3 NOAA sea surface temperature

Our first data set is the NOAA Optimum Interpolation SST V2 data set, containing the sea surface temperature weekly on a 1 degree grid across the period 1981-2018¹. The first 20 years were used as training data, the rest was reserved for testing. From this, we build data driven forecasts, predicting the next 5 weeks of temperature evolution from historical data. We project the system onto the first 6 POD modes, which is sufficient to approximate the temperature field (Fig. 1), and examine the evolution of these modal coefficients (Fig. 1a-b). From predictions of the modal coefficients, we can reconstruct the temperature field and test the temperature predictions at a sensor located at a specific coordinate, Fig. 2. The baseline to improve on is the climatology prediction; the historical average temperature at that specific location for the time of year.

For our purposes, we will use a bi-directional LSTM network [19] as our high-fidelity model, taking 5 weeks of historical data as input and making a 5 week forecast. This model is proto-typical for more complex machine learning models that could be deployed on data like this [11]. In short, a bi-directional LSTM is a recurrent neural network, that instead of acting purely sequentially on the data, has an additional pass backwards in time, Fig. 3c, in practice improving predictive power. In all cases, we take the first 20 years as training data, and use the remainder for testing. We compare the LSTM predictions against the recorded temperature and climatology baseline, Fig. 2, finding that while both the climatology and LSTM forecast captures the seasonal trends, the LSTM predictions have lower L_2 error and cosine norm closer to 1, as compared to the recorded temperature, than the climatology prediction for all sensors tested.

Our aim now is to investigate when and where the model breaks down, as well as identifying where expensive HF function evaluations can be avoided, and substituted for low-fidelity alternatives. To do so, we introduce two LF models. The first is linear regression, applied autoregressively to each mode individually, with an input window of 5 weeks, and making a single prediction. The best linear model is fitted to the training data for a particular mode, and then applied to the testing data. The second is a random forest regressor [22], which takes in all 5 weeks of input, and trains a model to predict each modal coefficient across the forecast. This works by training 100 decision trees on the input sequence, and then taking an average of the results as the output, Fig. 3d.

Unsurprisingly, the LF models perform worse on testing data than the HF one, Fig. 3(a,b). However, the low-fidelity models are extremely cheap to train and evaluate, we therefore seek to make use of

¹Available at <https://www.esrl.noaa.gov/psd/>




Figure 2: Sea surface temperature at 4 randomly chosen sensors (orange points) for the first 4 years of testing data (2002-2006), with POD-LSTM prediction and climatology baseline. After training, the LSTM model was applied non-autoregressively, making 5 weeks of predictions at a time. The spatial predictions were then reconstructed from the POD basis. Also shown is the climatology baseline.

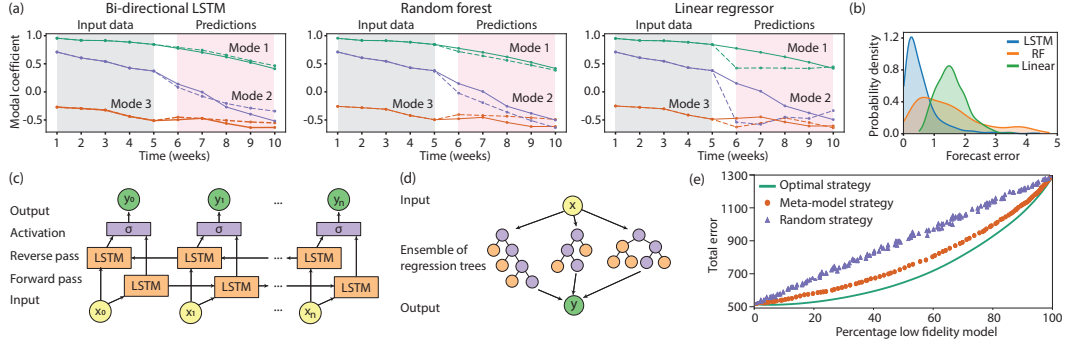


Figure 3: Comparison of high and low-fidelity models. (a) Input and predictions for a representative test example, true modal coefficients are solid lines, predictions are dashed. (b) Distributions of forecast errors for the testing data. (c) Architecture of the bi-directional LSTM, in our case $n = 10$. The input information travels forwards through time, and then backwards again to allow for a more accurate global prediction. (d) Architecture of the random forest; the input is sent to several decision trees, and their result is aggregated at the end for a prediction. (e) Total error of adaptive meta-modeling strategy against the optimal and random strategy

them whenever possible. Suppose now that we have a limited computational budget for evaluating a number of forecasts, and so are forced to use a mixture of HF and LF models. If we made the choice randomly with probability p for the HF model, we would expect the average error to be $p \times (\text{HF average error}) + (1 - p) \times (\text{LF average error})$. Similarly, we expect the computational cost to be $p \times (\text{HF cost}) + (1 - p) \times (\text{LF cost})$. We will use the two LF models to create a meta-modeling strategy that achieves lower error than the random strategy for the same computational cost. The strategy is as follows:

1. Take input data and perform the random forest and linear forecasts.
2. Calculate the difference between the predictions.
3. If this difference exceeds some threshold, evaluate the HF model. If not, stick with the LF model.

This simple strategy using two LF models significantly outperforms the random strategy, Fig. 3e. Moreover, it performs similarly to the theoretical optimum strategy, defined as the best possible choice of HF and LF evaluations, given their values and the true value in advance. This optimum strategy is clearly not viable in practice, yet the meta-model strategy performs similarly.

4 DayMet North America Daily Surface Temperature

We now validate the ideas we have introduced on an alternative data set. We consider the forecasting problem for the Daymet data set², containing the maximum daily temperature field across North

²available at <https://daymet.ornl.gov/>

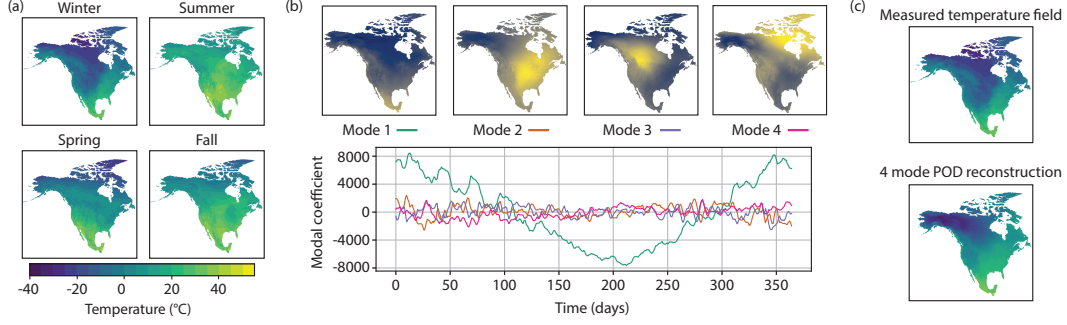


Figure 4: (a) Maximum daily temperature field across North America in 2011, representative seasonal snapshots shown. (b) Temporal evolution of the first four POD modes of the temperature field across one year, together with their representative basis elements. (c) Reconstruction of the temperature field from 4 modes against the true value.

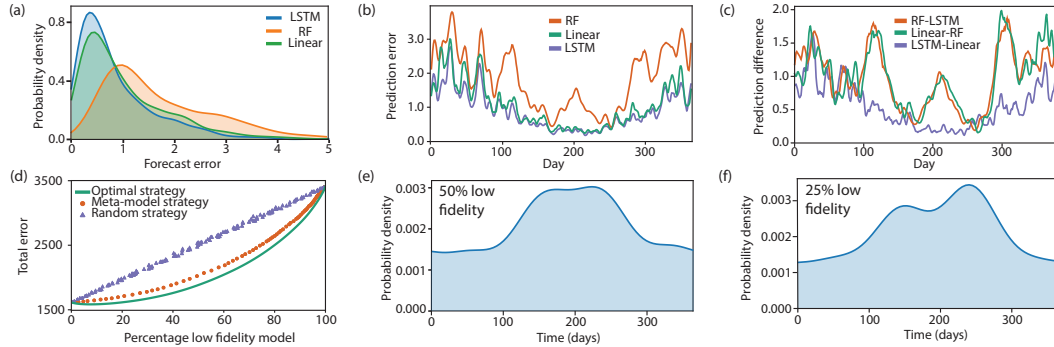


Figure 5: (a) Forecasting error distribution for HF and LF models. (b) Seasonal prediction error, averaged across 5 years of testing data and over a window of 5 days. (c) Seasonal prediction difference between models. Model predictions were made, and the norm of their difference was calculated and averaged as in (b). Difference between LSTM and linear models recovers seasonal trend in error. (d) Total error of meta-modeling strategy against the theoretical optimum and random strategy. (e-f) Seasonal distribution for where the LF model is evaluated for 25% and 50% LF model evaluations.

America for 2000-2015, Fig. 4. The first 11 years are used as training data, the rest is reserved for testing. The temperature over land is more variable than the sea temperature, so here we forecast 7 days from 7 days of historical data, and take only the first 4 POD modes.

The bi-directional LSTM has the lowest total prediction error, Fig. 5, and outperforms the climatology baseline for all sensors tested. The linear model performs almost as well, and both perform better than the random forest, Fig. 5a-b. Beyond overall error analysis, we aim to interrogate where predictions break down. Averaging the prediction error across 5 years of testing data, we find that error varies with season, with the worst predictions in winter, Fig. 5c. Examining the difference between models, we also find this seasonal variation, meaning we need not know the error to deduce the regions of low predictive power, Fig. 5d. On this data, the meta-model approach is almost as good as the theoretical optimum strategy, Fig. 5d. Under this strategy, we see that in the summer when the LF models agree, a LF model is used for prediction, whereas in the winter the HF model is used, Fig. 5e-f.

5 Summary

In this work, we considered the forecasting problem for two real-world geophysical data sets, the NOAA sea surface temperature and the NASA Daymet land surface temperature data for North America. We confirmed that the combination of dimensionality reduction through proper orthogonal decomposition and recurrent neural network based predictions outperforms a climatology baseline. While the neural network prediction outperforms simple random forest and linear predictions, we were able to extract information from these LF models nevertheless. Specifically, by using the two LF

models, we were able to devise a simple rule for deciding whether to use a HF or LF model, which allowed a near optimal decision strategy. When used as a strategy for complicated architectures, this could allow substantial computational savings, especially when ensemble predictions are required. We were further able to use the LF models to understand seasonal variation in predictive power, finding that our models predictive power is highest in the summer for land surface temperature, a valuable insight for climate modeling. In future, it will be of interest to apply this to more advanced network architectures.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357. This research was funded in part and used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. DS acknowledges support from the NSF-MSGI fellowship. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. DOE or the United States Government.

A Proper Orthogonal Decomposition

Proper orthogonal decomposition (POD) provides a systematic method to project dynamics of a high dimensional system onto a lower dimensional subspace. We suppose that a single snapshot of the full system is a vector in \mathbb{R}^N , where N could be the number of grid points at which the field is resolved. Observing the system across a number of time points gives us the snapshots $\theta_1, \dots, \theta_T$, with mean subtracted by convention. The aim of POD, is to find a small set of orthonormal basis vectors v_1, \dots, v_M , with $M \ll N$, which approximates the spatial snapshots,

$$\theta_t \approx \sum_{j=1}^M a_j(t) v_j, \quad t = 1, \dots, T, \quad (1)$$

and so allows us to approximate the evolution of the full N dimensional system, by considering only the evolution of the M coefficients $a_j(t)$. POD chooses the basis, v_j , to minimize the residual with respect to the L_2 norm,

$$R = \sum_{t=1}^T \left\| \theta_t - \sum_{j=1}^M a_j(t) v_j \right\|^2. \quad (2)$$

Defining the snapshot matrix, $S = [\theta_1 | \dots | \theta_T]$, the optimal basis is given by the M eigenvectors of SS^T , with largest eigenvalues, after which, the coefficients are found by orthogonal projection, $a_j(t) = \langle \theta_t, v_j \rangle$ [21].

For both of our data sets, we take only the training data snapshots, say D_1, \dots, D_T , from which we calculate the mean $\bar{D} = (1/T) \sum_t D_t$, hence defining the mean subtracted snapshots $\theta_t = D_t - \bar{D}$. We then create the snapshot matrix, S , and find numerically the M eigenvectors of SS^T with largest eigenvalues. From this, we train models, \mathcal{N} , to forecast the coefficients

$$a(t+1) \approx \hat{a}(t+1) = \mathcal{N}(a(t), a(t-1), \dots). \quad (3)$$

making predictions of future coefficients given previous ones.

To test the predictions on unseen data, E_1, \dots, E_k , we take the mean \bar{D} , and vectors v_j calculated from the training data to get that

$$a_j(t) = \langle E_t - \bar{D}, v_j \rangle, \quad j = 1, \dots, M, \quad (4)$$

which will be used by the model \mathcal{N} to make a prediction for future coefficients. The prediction for the coefficients \hat{a} , can be converted into predictions in the physical space by taking $\bar{D} + \sum_j \hat{a}_j v_j$. This procedure only makes use of testing data to pass into the model, not to train the model in any way. Crucially, to make a forecast of E_{t+1} , only previous measurements E_t, E_{t-1}, \dots are needed.

References

- [1] M. Auffhammer, P. Baylis, and C. H. Hausman. Climate change is projected to have severe impacts on the frequency and intensity of peak electricity demand across the united states. *Proceedings of the National Academy of Sciences*, 2017.
- [2] P. Balaprakash, R. Egele, M. Salim, S. Wild, V. Vishwanath, F. Xia, T. Brettin, and R. Stevens. Scalable reinforcement-learning-based neural architecture search for cancer deep learning research. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–33, 2019.
- [3] N. D. Brenowitz and C. S. Bretherton. Prognostic validation of a neural network unified physics parameterization. *Geophysical Research Letters*, 45(12):6289–6298, 2018.
- [4] A. Chattopadhyay, P. Hassanzadeh, and S. Pasha. Predicting clustered weather patterns: A test case for applications of convolutional neural networks to spatio-temporal climate data. *Scientific Reports*, 10(1):1317, Jan 2020.
- [5] A. Chattopadhyay, E. Nabizadeh, and P. Hassanzadeh. Analog forecasting of extreme-causing weather patterns using deep learning. *Journal of Advances in Modeling Earth Systems*, 12(2):e2019MS001958, 2020. e2019MS001958 10.1029/2019MS001958.
- [6] T. W. Ford, P. A. Dirmeyer, and D. O. Benson. Evaluation of heat wave forecasts seamlessly across subseasonal timescales. *npj Climate and Atmospheric Science*, 1(1):20, Oct 2018.
- [7] P. Gentine, M. Pritchard, S. Rasp, G. Reinaudi, and G. Yacalis. Could machine learning break the convection parameterization deadlock? *Geophysical Research Letters*, 45(11):5742–5751, 2018.
- [8] T. Gneiting and A. E. Raftery. Weather forecasting with ensemble methods. *Science*, 310(5746):248–249, 2005.
- [9] E. J. Kendon, N. M. Roberts, H. J. Fowler, M. J. Roberts, S. C. Chan, and C. A. Senior. Heavier summer downpours with climate change revealed by weather forecast resolution model. *Nature Climate Change*, 4(7):570–576, Jul 2014.
- [10] P. Lynch. The origins of computer weather prediction and climate modeling. *Journal of Computational Physics*, 227(7):3431 – 3444, 2008. Predicting weather, climate and extreme events.
- [11] R. Maulik, R. Egele, B. Lusch, and P. Balaprakash. Recurrent Neural Network Architecture Search for Geophysical Emulation. *arXiv e-prints*, Apr. 2020.
- [12] R. Maulik, A. Mohan, B. Lusch, S. Madireddy, P. Balaprakash, and D. Livescu. Time-series learning of latent-space dynamics for reduced-order model closure. *Physica D: Nonlinear Phenomena*, 405:132368, 2020.
- [13] M. Milano and P. Koumoutsakos. Neural network modeling for near wall turbulent flow. *Journal of Computational Physics*, 182(1):1 – 26, 2002.
- [14] R. Mizuta, H. Yoshimura, H. Murakami, M. Matsueda, H. Endo, T. Ose, K. Kamiguchi, M. Hosaka, M. Sugi, S. Yukimoto, et al. Climate simulations using mri-agcm3. 2 with 20-km grid. *Journal of the Meteorological Society of Japan. Ser. II*, 90:233–258, 2012.
- [15] W. Ohfuchi, H. Sasaki, Y. Masumoto, and H. Nakamura. Mesoscale resolving simulations of the global atmosphere and ocean on the earth simulator. *Eos, Transactions American Geophysical Union*, 86(5):45–46, 2005.
- [16] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, and A. Ramadhan. Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*, 2020.
- [17] M. Satoh, T. Matsuno, H. Tomita, H. Miura, T. Nasuno, and S. Iga. Nonhydrostatic icosahedral atmospheric model (nicam) for global cloud resolving simulations. *Journal of Computational Physics*, 227(7):3486 – 3514, 2008. Predicting weather, climate and extreme events.

- [18] S. Scher and G. Messori. How global warming changes the difficulty of synoptic weather forecasting. *Geophysical Research Letters*, 46(5):2931–2939, 2019.
- [19] M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [20] J. Sillmann, T. Thorarinsdottir, N. Keenlyside, N. Schaller, L. V. Alexander, G. Hegerl, S. I. Seneviratne, R. Vautard, X. Zhang, and F. W. Zwiers. Understanding, modeling and predicting weather and climate extremes: Challenges and opportunities. *Weather and Climate Extremes*, 18:65 – 74, 2017.
- [21] K. Taira, M. S. Hemati, S. L. Brunton, Y. Sun, K. Duraisamy, S. Bagheri, S. T. M. Dawson, and C.-A. Yeh. Modal analysis of fluid flows: Applications and outlook. *AIAA Journal*, 58(3):998–1022, 2020.
- [22] Tin Kam Ho. Random decision forests. In *Proceedings of 3rd International Conference on Document Analysis and Recognition*, volume 1, pages 278–282 vol.1, 1995.