Forecasting El Niño with Convolutional and Recurrent Neural Networks

Ankur Mahesh* ClimateAi
Maximilian Evans ClimateAi
Garima Jain ClimateAi
Mattias Castillo ClimateAi
Aranildo Lima ClimateAi

Brent Lunghino ClimateAi
Himanshu Gupta ClimateAi
Carlos Gaitan ClimateAi
Jarrett K. Hunt ClimateAi

Omeed Tavasoli ClimateAi
Patrick T. Brown ClimateAi, San Jose State University
V. Balaji Geophysical Fluid Dynamics Laboratory

Abstract

The El Niño Southern Oscillation (ENSO) is the dominant mode of variability in the climate system on seasonal to decadal timescales. With foreknowledge of the state of ENSO, stakeholders can anticipate and mitigate impacts in climate-sensitive sectors such as agriculture and energy. Traditionally, ENSO forecasts have been produced using either computationally intensive physics-based dynamical models or statistical models that make limiting assumptions, such as linearity between predictors and predictands. Here we present a deep-learning-based methodology for forecasting monthly ENSO temperatures at various lead times. While traditional statistical methods both train and validate on observational data, our method trains exclusively on physical simulations. With the entire observational record as an out-of-sample validation set, the method's skill is comparable to that of operational dynamical models. The method is also used to identify disagreements among climate models about the predictability of ENSO in a world with climate change.

1 Introduction

The El Niño-Southern Oscillation (ENSO) is a cycle of warm and cold temperatures in the equatorial Pacific Ocean that influences weather patterns around the world. It impacts North American temperature and precipitation, (Ropelewski and Halpert) the Indian Monsoon (Kumar, Rajagopalan, and Cane), and hurricanes in the Atlantic (Pielke Jr and Landsea). Thus, it has consequences for agricultural planning, commodity prices, insurance terms, and energy availability (Yu et al. Anderson et al. Aryal et al.).

Traditionally, the European Center for Medium-Range Weather Forecasts (ECMWF) runs a physics-based seasonal forecasting model called SEAS5. With supercomputers, they forecast the Niño-3.4 index, an area in the tropical Pacific Ocean averaged from 5N–5S and 120–170W (Bamston, Chelliah, and Goldenberg). Existing machine learning ENSO forecasts have utilized long short-term memory (LSTM) neural networks (Broni-Bediako et al.) or a combination of autoregressive integrated moving average models and artificial neural networks (Nooteboom et al.). While these forecasts are trained exclusively on observations, we train a neural network only on simulations from Coupled Atmosphere–Ocean General Circulation Models (AOGCMs) and evaluate it on observations. This

*Corresponding author: mahesh.ankur10@gmail.com

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

serves as an implicit investigation of the ability of AOGCMs to simulate the climate accurately, and it allows us to keep the entire observational record as a validation set, ensuring the neural networks' predictions are robust.

AOGCMs are numerical simulations of atmospheric and ocean processes used to improve understanding of the climate system and its temporal behavior. The scientific community has created historical, present, and future simulations of the behavior of the climate under different scenarios of greenhouse gas emissions. By training neural networks on AOGCMs, we present three contributions:

1. Trained only on physical simulations, neural networks offer comparable performance to operational seasonal forecasting models.
2. To evaluate the interpretability of neural networks, we compare their saliency maps to global grids showing the R^2 value between each grid cell and the subsequent Niño-3.4 index.
3. We investigate the predictability of ENSO in a world with climate change and a world without climate change.

2 Forecasting El Niño with Deep Learning

2.1 Methodology

We train a neural network that takes as input a 24-month time series of monthly surface temperature on a global 192 x 96 grid, and it forecasts the Niño-3.4 index at a specified lead time. Our neural network is a convolutional and recurrent architecture: it encodes the spatial information of each global surface temperature grid using a 6-layer convolutional neural network, and then feeds the encoded information into an LSTM to learn from the temporal sequence. (See Appendix 5.1 for a detailed description of the architecture.)

Based on their accuracy in simulating ENSO (Voldoire et al.), we trained our neural network on the following AOGCMs (named after the modeling centers that produced them): CNRM-CM5 (800 years), MPI-ESM-LR (1000 years), NorESM1-M (500 years), HadGEM2-ES (500 years) and GFDL-ESM2G (500 years). To develop a set of ground truth Niño-3.4 targets, we calculate the Niño-3.4 indices reported in ERA5 (Hersbach), a gridded dataset of historical observations. Due to SEAS5 data availability, we use the years 1993-2016 as our validation set.

2.2 Results

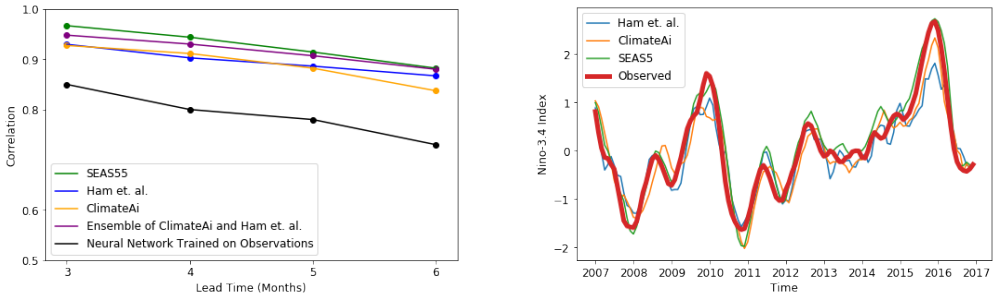

(a) (b)

Figure 1: (1a) Correlation of ClimateAi method, Ham et. al. method, Ensemble of deep learning methods, and SEAS5. (1b) Time series comparing ClimateAi Method, Ham et. al. Method, SEAS5, and Observed Nino3.4 Index

As shown in Figure 1a and 1b, neural networks trained on AOGCMs offer comparable performance to SEAS5. To show the full potential of deep learning, we include predictions from an ensemble of our proposed method and Ham et. al's method (Ham, Kim, and Luo). Ham et. al. recently proposed a method of training CNNs on AOGCMs; however, their method does not utilize LSTMs and uses an additional input predictor: ocean heat content. Also, in Figure 1a, we demonstrate that the proposed method outperforms another other benchmark: training a neural network purely on observations. We

conclude that neural networks learn better from abundant AOGCM simulations than from limited historical observations. In appendix section 5.3, we compare the performance of the neural networks to SEAS5 for extreme values of the Nino3.4 Index.

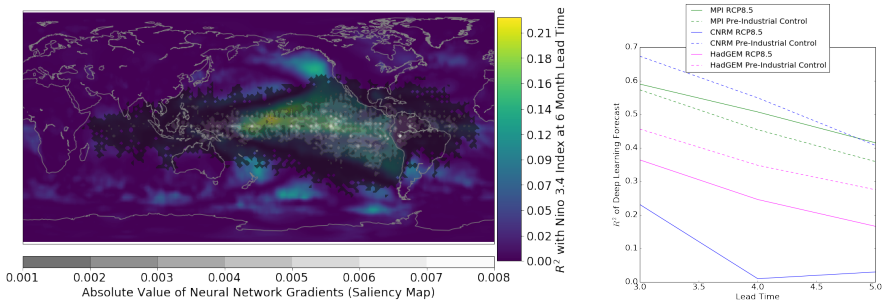

(a) (b)

Figure 2: (2a) Neural Network Saliency Map Overlaid With Niño-3.4 R^2 Values at 3 Month Lead Time. (2b) Performance of Neural Networks When Forecasting in Pre-Industrial Control and RCP8.5 scenario of MPI, HadGEM, and CNRM AOGCMs.

To investigate the geographic regions of most importance for the networks' prediction, we employ saliency maps. Saliency maps are a computer vision interpretability technique that calculate the gradient of the final prediction with respect to each pixel of an input image; they indicate the most important pixels to the neural network's prediction (Simonyan, Vedaldi, and Zisserman). In gridded climate datasets, each grid cell is analogous to a pixel in an image, so saliency maps can be directly compared to plots showing the R^2 coefficient between each grid cell and the Niño-3.4 index. In Figure 2a, our neural network's saliency map (averaged across all members of the validation set) indicates that it primarily bases its forecast on equatorial temperature in the Pacific Ocean. Its saliency map aligns with areas with high R^2 values, showing that the model primarily relies on the regions that are strong linear predictors of the Niño-3.4 Index. Still, the neural network also highlights regions with low R^2 in the Pacific and Indian Ocean as important predictors, indicating that the model does not purely rely on linear relationships to form its predictions.

3 Evaluating El Niño Predictability in a Warmer World

Existing work has explored the trend, magnitude, and variability of ENSO in future climate projections (Yeh et al.; Ashok and Yamagata; Van Oldenborgh, Philip, and Collins). Leveraging the light computational expense of our seasonal forecasting model, we investigate the *predictability* of ENSO in a warmer world.

We explore the degree to which it is possible to forecast the Niño-3.4 index at in three AOGCMS: CNRM, HadGEM, and MPI. Each AOGCM has a pre-industrial control run, in which greenhouse gas emissions stay at their pre-industrial levels, and a Representative Concentration Pathway 8.5 (RCP8.5) scenario, in which greenhouse gas (GHG) emissions increase substantially by 2100. To ensure an equitable comparison, we train neural networks at lead times of 3, 4 and 5 months on each scenario of each climate model independently; this accounts for intrinsic AOGCM differences in responding to GHG emissions. Each neural network is trained on 60 years and evaluated on 30 years of data. The models trained on RCP8.5 are only trained and validated on years after 2200 when emissions have stabilized.

As shown in Figure 2b, in the MPI and HadGEM AOGCMs, ENSO is similarly predictable in RCP8.5 and pre-industrial control. However, in CNRM, a warmer world is significantly less predictable than the corresponding pre-industrial control run. (The time series from which the R^2 values are calculated are included in the Appendix Section 5.2.) We conclude that ENSO predictability is a point of disagreement among current climate models. When applied to all ensemble members across all AOGCMs submitted to the Climate Model Intercomparison Project, our method can determine if there is a consensus in ENSO predictability.

4 Discussion

We demonstrate that neural networks trained purely on physical simulations can forecast ENSO with simular accuracy as SEAS5, the state-of-the-science seasonal forecasting dynamical model. This implicitly indicates that AOGCMs accurately resolve the behavior of ENSO and that neural networks learn this behavior embedded across AOGCMs. Finally, ENSO predictability can be a new metric used to evaluate climate models and to investigate differences between present and future climate.

Works Cited

Ropelewski, Chester F and Michael S Halpert. "North American precipitation and temperature patterns associated with the El Niño/Southern Oscillation (ENSO)". *Monthly Weather Review* 114.12 (1986): 2352–2362. Print.

Bamston, Anthony G, Muthuvel Chelliah, and Stanley B Goldenberg. "Documentation of a highly ENSO-related SST region in the equatorial Pacific: Research note". *Atmosphere-ocean* 35.3 (1997): 367–383. Print.

Kumar, K Krishna, Balaji Rajagopalan, and Mark A Cane. "On the weakening relationship between the Indian monsoon and ENSO". *Science* 284.5423 (1999): 2156–2159. Print.

Pielke Jr, Roger A and Christopher N Landsea. "La nina, el nino, and atlantic hurricane damages in the united states". *Bulletin of the American Meteorological Society* 80.10 (1999): 2027–2034. Print.

Van Oldenborgh, Geert Jan, SY Philip, and Matthew Collins. "El Niño in a changing climate: a multi-model study". *Ocean Science* 1.2 (2005): 81–95. Print.

Ashok, Karumuri and Toshio Yamagata. "Climate change: The El Niño with a difference". *Nature* 461.7263 (2009): 481. Print.

Yeh, Sang-Wook, et al. "El Niño in a changing climate". *Nature* 461.7263 (2009): 511. Print.

Simonyan, Karen, Andrea Vedaldi, and Andrew Zisserman. "Deep inside convolutional networks: Visualising image classification models and saliency maps". *arXiv preprint arXiv:1312.6034* (2013). Print.

Voldoire, Aurore, et al. "The CNRM-CM5. 1 global climate model: description and basic evaluation". *Climate Dynamics* 40.9-10 (2013): 2091–2121. Print.

Hersbach, Hans. "The ERA5 Atmospheric Reanalysis." *AGU Fall Meeting Abstracts*. 2016. Print.

Anderson, Weston, et al. "Trans-Pacific ENSO teleconnections pose a correlated risk to agriculture". *Agricultural and Forest Meteorology* 262 (2018): 298–309. Web.

Aryal, Yog N., et al. "Long term changes in flooding and heavy rainfall associated with North Atlantic tropical cyclones: Roles of the North Atlantic Oscillation and El Niño-Southern Oscillation". *Journal of Hydrology* 559 (2018): 698–710. Web.

Broni-Bediako, Clifford, et al. "El niño-southern oscillation forecasting using complex networks analysis of LSTM neural networks". Jan. 2018. Print.

Nooteboom, P. D., et al. "Using network theory and machine learning to predict El Niño". *Earth System Dynamics* 9.3 (2018): 969–983. Web.

Yu, Bin, et al. "ENSO and Sea Surface Temperature Anomalies in Association with Canadian Wheat Yield Variability". *Atmosphere-Ocean* 56.1 (2018): 28–39. eprint: https://doi.org/10.1080/07055900.2017.1416574. Web.

Ham, Yoo-Geun, Jeong-Hwan Kim, and Jing-Jia Luo. "Deep learning for multi-year ENSO forecasts". *Nature* 573.7775 (2019): 568–572. Print.

5 Appendix

5.1 Architecture Details

The 2D-CNN contains 6 convolution layers with ReLU non-linearity at each layer and one fully connected layer at the end.

The network details are the following: [Conv2D-> batch normalization->ReLU] * 5 -> [Conv2D -> ReLU] -> FC. The first convolutional layer uses 10 filters and then the number of filters double in every layer. Paddings and strides are defined so as to get the desired size reduction.

The output vector from the fully connected layer feeds into a RNN in sequences of successive months (here, 24 months). The RNN consists of a many to one architecture and has 2 LSTM layers, with each of them having 500 hidden units. At the end, the hidden state of the last time step is decoded to a real value using a fully connected layer, to output the predicted monthly Niño-3.4 sea surface temperature anomalies.

5.2 Time Series of RCP8.5 and Pre-Industrial Control performance

We include the time series overlaying the deep learning prediction and the target Niño-3.4 index in the RCP8.5 and pre-industrial control runs of CNRM and MPI.

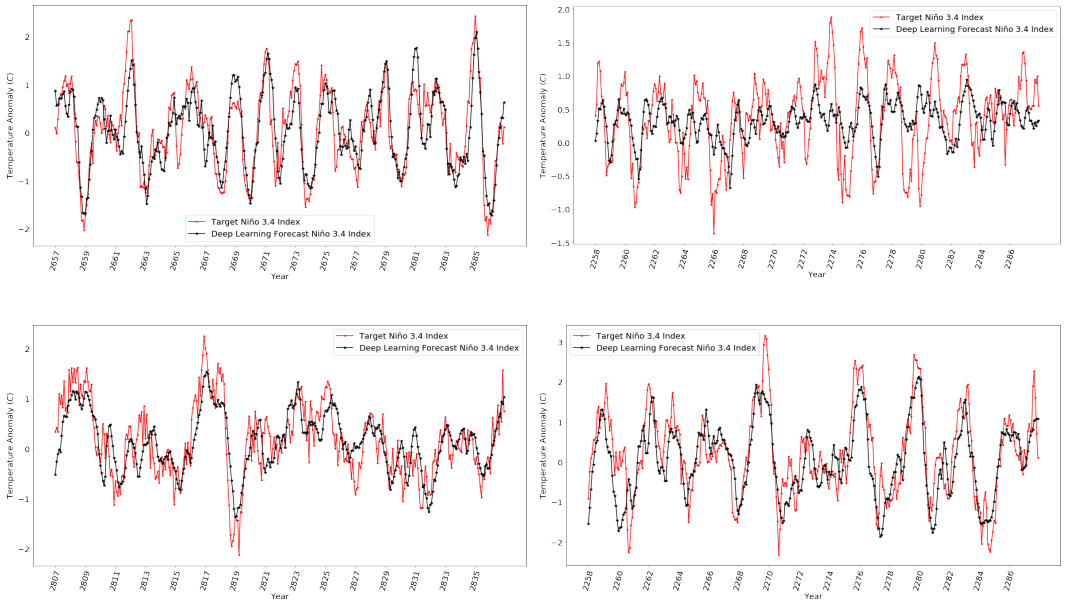

Figure 3: (Top Left) CNRM pre-industrial control (RMSE:0.481, R^2:0.700) (Top Right) CNRM RCP8.5 (RMSE:0.528, R^2:0.230)(Bottom Left) MPI pre-industrial control (RMSE:0.472, R^2:0.592). (Bottom Right) MPI RCP8.5 (RMSE:0.681, R^2:0.592).

5.3 Evaluation of Neural Networks at Extreme Values of the ENSO Index

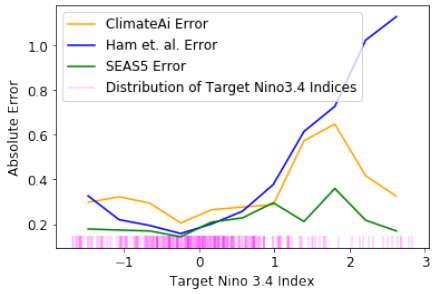

In the figure above, we evaluate the performance of our neural network with SEAS5 and with Ham et. al.'s neural network at a 6 month lead time. SEAS5 still offers better performance at extreme values of the Niño-3.4 index, indicating that an ensemble of neural networks still has room for improvement. We believe that this area offers the next challenge for improving the performance of neural networks.