TRACKING CLIMATE MODELS

Claire Monteleoni[1], Shailesh Saroha[2], and Gavin Schmidt[3]
Columbia University

Introduction:

With the increased threats of climate change, such as global warming, and the increasing severity of storms and natural disasters, the study of climate science is an international effort of growing urgency. A fundamantal tool used in predicting climate is the use of large-scale mathematical models run as computer simulations. Geophysical experts, including climate scientists and meteorologists, encode their knowledge of a myriad of processes into highly complex mathematical models. One climate model will include the modeling of such processes as sea-ice melting, cloud formation as a function of increased pollution in the atmosphere, cyclone formation, and the creation, depletion and transport of many atmospheric gases. These are just a few of the processes modeled in one model; each climate model is a highly complex system.

The global effort on climate modeling started in the 1970s, and the models have evolved over time, becoming extremely complex. There are currently 20 laboratories across the world whose climate models inform the Intergovernmental Panel on Climate Change (IPCC), a panel established by the United Nations, that was recognized for its work on climate change with the 2007 Nobel Peace Prize (shared with former US Vice President Al Gore). The state-of-the-art is that there is very high variance among the predictions of the 20 models. It was observed however, that the average prediction over all the models is a more consistent predictor (for example, of global mean temperature), than any one model. While one model may predict well during a certain time period, there is no one model out of the 20 that is better over all time periods, than the average over all models.

Our goal is to provide a machine learning algorithm that, given the predictions of each model, in the online setting, produces predictions that are better than any one model used for the entire sequence.[4] We will demonstrate that we can in fact achieve a much stronger result: predictions that are strictly better than the average prediction of the 20 models, *on every observation*. This is an impactful result because to date, the average of all models' predictions had been the benchmark.

Methods, results, and future work:

We are not aware of other applications of machine learning to the problem of tracking climate models. We applied the Learn-α algorithm of Monteleoni and Jaakkola [1] to track a shifting sequence of temperature values with respect to the predictions of "experts," which we instantiate in this case with climate models. Their work extends the literature on algorithms to track a sequence of observations with respect to the predictions of a set of experts, due to Herbster and Warmuth [2], and others.[5]

We used historical temperature data from 1900 through 2008, as well temperature predictions from 20 different climate models from 1900 to 2098. These 20 climate models were developed at research laboratories worldwide, and their predictions inform the IPCC. It is important to emphasize that these are not data-driven models but rather complex mathematical models based on geophysical and metereological principles. In particular they are not "trained" on data as is done with machine learning models. Therefore it is valid to run them predictively on past data.

Figure 1a plots the squared error between predicted and observed annual mean temperature anomalies, by year. We plot the best model *per year* (green) since it is the unattainable benchmark which detects every change-point and may switch models as often as every observation. The algorithm (black) consistently predicts better than the average (red) of all the climate models (averaged over all their runs) every year, often by significant margins. We also plot the worst model per year (blue) to show the range of variability across models. Our performance with respect to the average global models curve is a break-through; as that was the current state-of-the-art.

[1]Center for Computational Learning Systems, cmontel@ccls.columbia.edu
[2]Department of Computer Science, ss3515@columbia.edu
[3]NASA Goddard Institute for Space Studies, gschmidt@giss.nasa.gov
[4]We use online learning because the eventual goal is to make both real-time and future predictions.
[5]There has also been recent work related to [1], in [3].

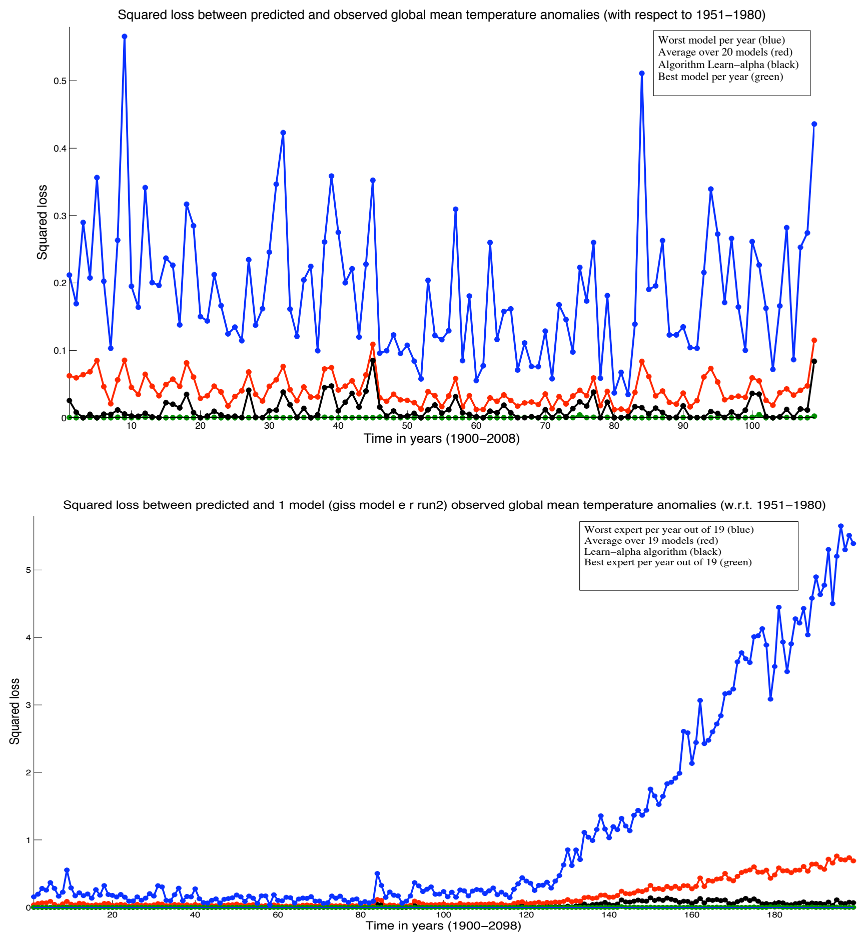

Figure 1: Top figure: a. Squared loss between predicted and observed global mean temperature anomalies; Bottom figure: b. Tracking the predictions of one model, using the predictions of the remaining 19 as input, with no true temperature observations.

Figure 1b demonstrates that our algorithm is very successful at predicting one model's predictions (we found similar results for using 19 climate models to predict the output of *any* one model, as the 20th), for future predictions up to the year 2098. This is notable, as the future predictions vary widely among the climate models. The (blue) curve indicating the worst model per year (with respect to predicting the model in question) varies increasingly into the future, whereas our algorithm tracks the best model per year very closely.

The exciting challenge begged by our encouraging results, is how to track climate models when predicting *future* climates. Our goal is to track models in unsupervised, or semi-supervised settings. One possible approach is to treat the model predictions themselves as observations, or partial feedback (see Figure 1b). We also to plan predict other climate quantities such as carbon dioxide.

References

[1] Monteleoni, C. and Jaakkola, T.: Online Learning of Non-stationary Sequences. *NIPS*, 2003.
[2] Herbster M. and Warmuth M.: Tracking the Best Expert. *Machine Learning Journal*, 32, pp. 151–178, 1998.
[3] de Rooij and S. van Erven, T.: Learning the Switching Rate by Discretising Bernoulli Sources Online. *AISTATS*, 2009.

Topic: learning algorithms, applications, climate
Preference: poster