

---

# HECT: High-Dimensional Ensemble Consistency Testing for Climate Models

---

Niccolò Dalmaso<sup>\*,1</sup>, Galen Vincent<sup>\*,1</sup>, Dorit Hammerling<sup>2</sup> and Ann B. Lee<sup>1</sup>

<sup>1</sup>Department of Statistics & Data Science, Carnegie Mellon University

<sup>2</sup>Department of Applied Mathematics and Statistics, Colorado School of Mines

## Abstract

Climate models play a crucial role in understanding the effect of environmental and man-made changes on climate to help mitigate climate risks and inform governmental decisions. Large global climate models such as the Community Earth System Model (CESM), developed by the National Center for Atmospheric Research, are very complex with millions of lines of code describing interactions of the atmosphere, land, oceans, and ice, among other components. As development of the CESM is constantly ongoing, simulation outputs need to be continuously controlled for quality. To be able to distinguish a “climate-changing” modification of the code base from a true climate-changing physical process or intervention, there needs to be a principled way of assessing statistical reproducibility that can handle both spatial and temporal high-dimensional simulation outputs. Our proposed work uses probabilistic classifiers like tree-based algorithms and deep neural networks to perform a statistically rigorous goodness-of-fit test of high-dimensional spatio-temporal data.

## 1 Introduction

Climate change adaptation has been identified as one of the two main action items in addressing climate change [22]. The main goals in climate change adaptation are to anticipate the consequences of extreme climate events and understand the effects on climate of environmental and man-made changes. In this context, forecasts from climate models are essential for informing adaptation strategies, both at a local and national level (such as the IPCC reports [18, 19]). Also, climate models can be used to estimate the impact on the climate of a specific action or policy. In a recent example, scientists managed to estimate the effect of COVID19-related lockdown and travel restrictions on global warming (which turned out to be negligible) using the FaIR v1.5 carbon cycle emulator [24].

Overall, climate models are divided into two categories, General Circulation Models (GCMs) or Earth System Models (ESMs) [7]. Within the latter category, the Community Earth System Model (CESM [9]) is a state-of-the-art “virtual laboratory” for studying past, present, and future global climate states. Developed mainly at the National Center for Atmospheric Research (NCAR), CESM consists of different simulation components covering various aspects of the climate system, including atmosphere, land, river runoff, land-ice, oceans, and sea-ice. CESM is a fully coupled system, meaning that the differential equations within each component are all solved at the same time by a custom computing architecture [3]. About 2,900 academic publications have used CESM and CESM data, and the code base currently contains around 1.5 million lines of code.<sup>2</sup>

Such high-capacity and complex climate models are in a constant state of development, and frequent checks need to be in place for “quality assurance” [17]; that is, to detect and reduce errors which could adversely affect the simulation results and potentially erode the model’s scientific credibility.

---

<sup>\*</sup>Equal Contribution

<sup>2</sup>Source: <http://www.cesm.ucar.edu/> and Hammerling et al., 2020 Joint Statistical Meeting




Figure 1: Schematics of the ECT methodology for quality assurance of climate models. A series of simulations with different initial conditions are run on a “trusted” machine (blue lines) to form a baseline ensemble which defines the forecast uncertainty (grey area). Test runs (red line) are then evaluated against this initial condition ensemble. The current state-of-the-art is to compress the data first by PCA, which may lead to estimation bias and loss of power. In this work, we propose a new testing framework (HECT) based on probabilistic classifiers that provides diagnostics and allows the data scientist to work directly with high-dimensional spatial and temporal outputs.

CESM is a deterministic simulation model, i.e., the output from two simulations are bit-for-bit (BFB) identical if they are run with the same initial conditions and on the same computing infrastructure. However, because of the chaotic nature of the equations underlying the simulation process, running the same code on new machine architecture, or running a tweaked version of the original code expected to return the same results, could produce output that is not BFB identical but still represents the same climate model. How can we make sure that non-BFB identical simulation outputs are a result of this natural variation rather than a “climate-changing” error we introduced in the code? *Ensemble consistency testing* (ECT, [1, 15, 16]) provides a statistical answer to such question. Figure 1 provides a schematic representation of ECT. Multiple CESM runs are initialized using a “trusted” version of the climate model (in blue), where in each run the initial temperatures are perturbed on the order of  $O(10^{-14})$ . The ensemble of outputs from these slightly perturbed runs provides a baseline which approximates the natural variability in accepted simulation output (represented by the grey area). This ensemble is contrasted to outputs from a series of “test runs”, i.e. runs from the modified version of the climate model (in red), and a pass or fail is issued based on the first 50 sample PCA scores for global climate variable averages at the last time step of the simulation. Currently, ECT is implemented as an automatic Python tool as part of the official CESM release.

While PCA-based testing has proven successful, its main limitations are (i) a loss of statistical power from data compression and (ii) lack of theoretical guarantees on type I and II error. For (i), PCA-based testing only uses a subset of the PCA scores, compressing the data even further than averaging across the spatial and temporal resolution of simulation outputs. This loss of information makes the test more likely to fail in flagging truly “climate-changing” modifications (loss of statistical power), as well as biasing the PCA scores estimates (which are ultimately used to issue a pass/fail). In this paper, we propose a novel *High-dimensional Ensemble Consistency Testing* (HECT) approach, which aims to address the limitations mentioned above. HECT is based on a two-sample test [11, 4], which leverages probabilistic classifiers and hence does not require a prior dimension reduction step, global averaging of climate variables, or limiting the test to the last time step; see Section 3 for specific examples of its applicability. In addition, the performance of the probabilistic classifier is shown to directly connect to the type I and type II error of the test [4], making it straightforward to select which classifier to use in each setting. Finally, HECT provides diagnostics by identifying statistically significant spatial and/or temporal differences between trusted and test runs. While our proposal is specific to CESM, we envision our methodology to apply to other climate models as well.

## 2 Data

In [16], which represents the latest development in CESM ECT, the trusted run ensemble is composed of 350 high-resolution runs, while there are much fewer test runs (usually less than five). Each run simulates 134 climate variables across the globe (approximated in 3-D by a  $1^\circ$  square grid over 30 vertical levels) for each time step of the simulation (which [16] set to 9 as a result of an analysis on ensemble variability). In practice, each run has approximately 1.7 billion entries,<sup>3</sup> which cannot be used for comparison as is. The ECT in [16] is performed by averaging all variables globally, only considering the last time step and removing variables that are redundant, have zero variance, or do not show enough variability over all time steps. For the remainder of this work we will refer to a run as  $\mathbf{X}$ , but implicitly we assume some level of spatial averaging, temporal averaging or feature selection.

## 3 ECT and Diagnostics via High-Dimensional Goodness-of-Fit Testing

Let  $P_0$  and  $P_1$  represent the distributions of trusted and test runs respectively, and suppose  $\mathcal{S}_0 = \{\mathbf{X}_1^0, \dots, \mathbf{X}_m^0\} \stackrel{i.i.d.}{\sim} P_0$  and  $\mathcal{S}_1 = \{\mathbf{X}_1^1, \dots, \mathbf{X}_n^1\} \stackrel{i.i.d.}{\sim} P_1$  where typically  $m \gg n$ . To test the hypothesis  $H_0 : P_1 = P_0$  versus  $H_1 : P_1 \neq P_0$ , we introduce a binary random variable or class label  $Y$ , and interpret  $P_i$ ,  $i = 0, 1$ , as the class-conditional distributions of  $\mathbf{X}$  given  $Y = i$ . The two-sample problem then becomes equivalent to testing for independence between  $\mathbf{X}$  and  $Y$ , i.e.,  $H_0 : P_1 = P_0$  is equivalent to  $H_0 : \mathbb{P}(Y = 1 | \mathbf{X} = \mathbf{x}) = \mathbb{P}(Y = 1) \quad \forall \mathbf{x} \in \mathcal{X}$ . Motivated by this, [11] propose a two-sample test based on the test statistic

$$\hat{\mathcal{T}} = \frac{1}{n+m} \sum_{i=1}^{n+m} (\hat{r}(\mathbf{X}_i) - \hat{\pi}_1)^2, \quad (1)$$

where  $\hat{r}(\mathbf{x})$  is an estimate of  $\mathbb{P}(Y = 1 | \mathbf{X} = \mathbf{x})$  based on  $\{(\mathbf{X}_i, Y_i)\}_{i=1}^{n+m}$ , and  $\hat{\pi}_1 = \frac{1}{n+m} \sum_{i=1}^{n+m} I(Y_i = 1)$  is an estimate of  $\mathbb{P}(Y = 1)$ . By the above reformulation, we have converted two-sample testing to the well-studied problem of estimating the ‘‘class posterior’’  $\mathbb{P}(Y = 1 | \mathbf{X} = \mathbf{x})$ , where we have access to an arsenal of high-dimensional regression (or probabilistic classification) algorithms for computing  $\hat{r}(\mathbf{x})$ . In settings where  $m \gg n$  (i.e., there are many more trusted than test runs), such as in the CESM ECT example, we may gain power by using the test statistic in equation (1) in a goodness-of-fit test rather than in a two-sample test; for details see [4] (Algorithm 5, Section D). In a nutshell, we generate several different ensemble runs  $E$  of size  $m_e$  to produce a set of realized test statistics  $\{\hat{\mathcal{T}}^{(e)}\}_{e=1}^E$ ; such set is then used to define a null distribution for testing the hypothesis that trusted and test runs are statistically indistinguishable.

**Contribution.** The main benefit of HECT is that one now can eliminate most of the intermediate steps between simulation output and ECT (spatial averaging, global averaging, variable selection and dimensionality reduction). Such bottlenecks could potentially be avoided by leveraging probabilistic classifiers within HECT which can handle high-dimensional spatial and temporal data. For example:

- We could leverage tree-based algorithms (e.g, Random Forest [2] and Gradient Boosted trees [6]), which are robust to highly correlated features and implicitly perform feature selection;
- Grid-level global data can be used as input to convolutional neural networks [14, 12, 8] without a prior dimension reduction step, thereby detecting local differences between trusted and test runs otherwise masked by global averaging input;
- Similarly, multivariate time-series can be directly input into suitable classifiers such as recurrent neural networks [5] and related recent developments [26, 20, 21, 13], thereby taking the entire simulated sequence of climate variables into account instead of only the last time step as in [16];
- Finally, we could potentially leverage spatio-temporal deep neural networks [10, 23] to compare runs that are only averaged over the vertical level of the atmosphere, hence analyzing information at a global level over all time steps of the simulation.

For all settings above, we plan on exploring deep transfer learning techniques to aid training in cases where a limited amount of simulation runs is available (see [25] for a review). Corollary

<sup>3</sup>46, 802 grid cells across 30 vertical levels for 134 variables over 9 time steps.

2 of Theorem 2 in [4] provides a connection between the mean integrated squared error (MISE) and the type I and type II error of the proposed HECT, with no restrictions on the classifier form (i.e., also including neural networks). What this means in practice is that we now have a concrete way of deciding which regression or probabilistic classification algorithm to implement in different settings based on (anticipated) highest statistical power. In addition, the test statistic (1) provides *diagnostics* by identifying statistically significant differences (if any) between runs in a potentially high-dimensional spatio-temporal feature space; [11] showcase the capability of diagnostics in feature space by comparing multivariate distributions of galaxy morphology images.

**Planned Experiments and Evaluation Metrics.** We plan on first using the ensemble runs of [16] (taking the global averaged at the 9<sup>th</sup> time step for each run) and reproduce the common experiments among [16, 15, 1] to explore whether HECT issues a pass or fail in agreement with current ECT methodologies. We also propose to look at the performance of HECT when we include more time steps and/or global-level data (i.e., averaged only on the vertical levels). Finally, we will determine the optimal ensemble size for different settings of HECT, in similar fashion as what was done by [16], Section 4.

## Acknowledgments

This work is supported in part by the NSF AI Planning Institute for Data-Driven Discovery in Physics, NSF PHY-2020295.

## References

- [1] A. H. Baker, D. M. Hammerling, M. N. Levy, H. Xu, J. M. Dennis, B. E. Eaton, J. Edwards, C. Hannay, S. A. Mickelson, R. B. Neale, D. Nychka, J. Shollenberger, J. Tribbia, M. Vertenstein, and D. Williamson. A new ensemble-based consistency test for the Community Earth System Model (pyCECT v1.0). *Geoscientific Model Development*, 8(9):2829–2840, 2015.
- [2] Leo Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, October 2001.
- [3] Anthony P. Craig, Mariana Vertenstein, and Robert Jacob. A new flexible coupler for earth system modeling developed for CCSM4 and CESM1. *The International Journal of High Performance Computing Applications*, 26(1):31–42, 2012.
- [4] Niccolò Dalmaso, Ann Lee, Rafael Izbicki, Taylor Pospisil, Ilmun Kim, and Chieh-An Lin. Validation of approximate likelihood and emulator models for computationally intensive simulations. In *International Conference on Artificial Intelligence and Statistics*, pages 3349–3361. Proceedings of Machine Learning Research, 2020.
- [5] Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.
- [6] Jerome H. Friedman. Greedy function approximation: A gradient boosting machine. *Ann. Statist.*, 29(5):1189–1232, 10 2001.
- [7] Hugues Goosse, P-Y Barriat, Wouter Lefebvre, Marie-France Loutre, and Violette Zunz. *Introduction to climate dynamics and climate modeling*. 01 2009.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [9] James W. Hurrell, M. M. Holland, P. R. Gent, S. Ghan, Jennifer E. Kay, P. J. Kushner, J.-F. Lamarque, W. G. Large, D. Lawrence, K. Lindsay, W. H. Lipscomb, M. C. Long, N. Mahowald, D. R. Marsh, R. B. Neale, P. Rasch, S. Vavrus, M. Vertenstein, D. Bader, W. D. Collins, J. J. Hack, J. Kiehl, and S. Marshall. The Community Earth System Model: A Framework for Collaborative Research. *Bulletin of the American Meteorological Society*, 94(9):1339–1360, 09 2013.
- [10] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of International Computer Vision and Pattern Recognition (CVPR 2014)*, 2014.

- [11] Ilmun Kim, Ann B. Lee, and Jing Lei. Global and local two-sample tests via regression. *Electron. J. Statist.*, 13(2):5253–5305, 2019.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
- [13] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long- and short-term temporal patterns with deep neural networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '18, page 95–104, New York, NY, USA, 2018. Association for Computing Machinery.
- [14] Yann LeCun and Yoshua Bengio. *Convolutional Networks for Images, Speech, and Time Series*, page 255–258. MIT Press, Cambridge, MA, USA, 1998.
- [15] Daniel J. Milroy, Allison H. Baker, Dorit M. Hammerling, John M. Dennis, Sheri A. Mickelson, and Elizabeth R. Jessup. Towards Characterizing the Variability of Statistically Consistent Community Earth System Model Simulations. *Procedia Computer Science*, 80:1589 – 1600, 2016. International Conference on Computational Science 2016, ICCS 2016, 6-8 June 2016, San Diego, California, USA.
- [16] Daniel J. Milroy, Allison H. Baker, Dorit M. Hammerling, and Elizabeth R. Jessup. Nine time steps: ultra-fast statistical consistency testing of the Community Earth System Model (pyCECT v3.0). *Geoscientific Model Development (Online)*, 11(2), 2 2018.
- [17] William L. Oberkampf and Christopher J. Roy. *Verification and Validation in Scientific Computing*. Cambridge University Press, USA, 1st edition, 2010.
- [18] Intergovernmental Panel on Climate Change. *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [O. Edenhofer, R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlomer, C. von Stechow, T. Zwickel, J.C. Minx, (eds.)]*. 2014.
- [19] Intergovernmental Panel on Climate Change. *Global warming of 1.5 C. An IPCC special report on the impacts of global warming of 1.5 C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty [V. Masson-Delmotte, P. Zhai, H. O. Portner, D. Roberts, J. Skea, P.R. Shukla, A. Pirani, Y. Chen, S. Connors, M. Gomis, E. Lonnoy, J. B. R. Matthews, W. Moufouma-Okia, C. Pean, R. Pidcock, N. Reay, M. Tignor, T. Waterfield, X. Zhou (eds.)]*. 2018.
- [20] Yao Qin, Dongjin Song, Haifeng Cheng, Wei Cheng, Guofei Jiang, and Garrison W. Cottrell. A dual-stage attention-based recurrent neural network for time series prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI'17, page 2627–2633. AAAI Press, 2017.
- [21] Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7785–7794. Curran Associates, Inc., 2018.
- [22] David Rolnick, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Luccioni, Tegan Maharaj, Evan D. Sherwin, S. Karthik Mukkavilli, Konrad P. Kording, Carla Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer Chayes, and Yoshua Bengio. Tackling Climate Change with Machine Learning. *arXiv e-prints*, page arXiv:1906.05433, June 2019.
- [23] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, page 802–810, Cambridge, MA, USA, 2015. MIT Press.

- [24] C. J. Smith, P. M. Forster, M. Allen, N. Leach, R. J. Millar, G. A. Passerello, and L. A. Regayre. Fair v1.3: a simple emissions-based impulse response and carbon cycle model. *Geoscientific Model Development*, 11(6):2273–2297, 2018.
- [25] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. In Věra Kůrková, Yannis Manolopoulos, Barbara Hammer, Lazaros Iliadis, and Ilias Maglogiannis, editors, *Artificial Neural Networks and Machine Learning – ICANN 2018*, pages 270–279, Cham, 2018. Springer International Publishing.
- [26] Rose Yu, Stephan Zheng, Anima Anandkumar, and Yisong Yue. Long-term Forecasting using Higher Order Tensor RNNs. *arXiv e-prints*, page arXiv:1711.00073, October 2017.