Machine learning to extend and understand the sources and limits of water cycle predictability on subseasonal-to-decadal timescales in the Earth system

Katherine Dagon[1], Maria J. Molina[1], Gerald A. Meehl[1], Jadwiga H. Richter[1], Elizabeth A. Barnes[2], Judith Berner[1], Julie M. Caron[1], Will Chapman[3], Gokhan Danabasoglu[1], David John Gagne[1], Sasha Glanville[1], Sue Ellen Haupt[1], Aixue Hu[1], Zane Martin[2], Kirsten Mayer[2], Kathy Pegion[4], Kevin Raeder[1], Isla Simpson[1], Aneesh Subramanian[5], and Steve Yeager[1]

[1]National Center for Atmospheric Research, Boulder, Colorado, [2]Colorado State University, Fort Collins, Colorado, [3]Scripps Institution of Oceanography, San Diego, California, [4]George Mason University, Fairfax, Virginia, [5]University of Colorado, Boulder, Colorado

Focal Area

This white paper provides initial insight on how artificial intelligence (AI) and machine learning (ML), including interpretability and explainable AI (XAI) methods, can be leveraged to glean insight from complex data for a paradigm-changing improvement in Earth system predictability on subseasonal-to-seasonal (S2S) and seasonal-to-decadal (S2D) timescales. The application of AI to extend and improve predictability, in combination with causal inference and uncertainty quantification, could lead to a transformative understanding of the integrative water cycle and associated extremes.

Science Challenge

The water cycle is a vital process for life on Earth, yet remains a poorly predicted phenomenon in modeling systems. Can we harness ML to understand the sources and limits of S2S and S2D predictability and extend the prediction skill of the integrative water cycle and its associated extremes?

Rationale

The integrative water cycle is a key component of the Earth system and is essential to sustaining humans and ecosystems. Water can be evaporated, transported, and precipitated via physical mechanisms that have teleconnections to large-scale modes of variability, such as the El Niño-Southern Oscillation (ENSO) and the Madden Julian Oscillation (MJO). However, prediction skill of the water cycle and associated extremes (e.g., flash floods and droughts) on S2S/S2D timescales from Earth system models remains poor (Pegion et al. 2019, Pendergrass et al. 2020, Meehl et al. 2021). Much of the S2S/S2D forecasting uncertainty associated with the water cycle arises from chaotic atmospheric dynamics and limited understanding of the upstream sources of predictability (Robertson et al. 2020). The growing complexity of Earth system models, which generate large amounts of data, as well as the many associated physical mechanisms increase the difficulty of identifying the processes that provide or limit predictability skill.

The power of XAI lies in enabling humans to understand ML predictions, increasing transparency and trust in complex ML models. Recent advances in XAI can be harnessed to provide insight into precursor mechanisms associated with prediction skill, which potentially arise from known and unknown modes of variability. However, there are challenges associated with XAI. For example, the degree of causality is not output from the analysis, nor is a measure of uncertainty. Moreover, the limited observational record on decadal or longer timescales poses challenges in a supervised learning framework, where we do not have sufficient labeled data (or "ground truth") to train ML models with. These challenges can be addressed using the proposed approaches detailed in this white paper, along with collaborations between

DOI: https://doi.org/10.2172/1769744
DOE BER, AI4ESP 2021 #AI4ESP1032 (https://www.ai4esp.org)

domain experts and ML scientists, which can advance Earth system predictability through a deeper understanding of the integrative water cycle.

Narrative

By training ML models using a supervised learning framework, with a combination of observations (e.g, ARM precipitation), reanalysis (e.g., NCEP-DOE), and/or Earth system models (e.g., E3SM), windows of forecast opportunities can be identified, along with sources and limits of predictability. The limits of predictability related to the water cycle can be detected by training an ML model to predict climate model errors, which can help isolate model components or processes responsible for the largest biases. The success of ML to extend predictability on S2S/S2D timescales has been demonstrated by several recent studies. For example, Ham et al. (2019) trained a convolutional neural network (CNN; LeCun et al. 2015) to predict ENSO, first using model simulations, and then using reanalysis data, in order to address challenges with limited observational data availability. This successful application of transfer learning produced predictions out to a lead time of one and a half years with skill that exceeds current state-of-the-art dynamical forecast systems. Here we describe one example of how, in theory, ML models and interpretability techniques can be used to extend S2S predictions of the water cycle generated with Earth system models (e.g., E3SM). A CNN can be trained to skillfully predict precipitation anomalies or associated modes of variability (e.g., MJO) with lead times of 2-4 weeks using S2S process-oriented fields (Kim et al. 2021). Precursor mechanisms that contribute to CNN prediction skill can then be identified using XAI methods (Barnes et al. 2020, Mayer and Barnes 2020), such as saliency maps (Simonyan et al. 2013, McGovern et al. 2019) and Layer-wise Relevance Propagation (Bach et al. 2015, Toms et al. 2020), and these results can be used to generate hypotheses that can be tested using dynamical models. XAI could also be approached using knowledge distillation (Liu et al. 2018), where a simple model (e.g., decision trees) can be trained using output from a skillful deep learning model, providing improved interpretability with limited reduction in accuracy.

ML can also be applied to identify and predict S2D modes of variability with teleconnections to the water cycle, using XAI to detect the relevant upstream signals. Within long E3SM simulations and building upon the work of Ham et al. (2019), fields other than sea surface temperatures and heights can be incorporated into ML forecasts of ENSO, to exploit any additional predictive power they may have. Fields relevant to other modes of variability, such as the MJO and Pacific Decadal Oscillation, could also be used to help illuminate relationships among climate signals (Tang et al. 2008, Meehl et al. 2010, Kapur and Zhang 2012). It is then possible to test whether incorporating other modes of variability further extends or limits the predictability of ENSO, which has implications for the integrative water cycle via teleconnections. Synthetic data can also be generated using ML applied to observations and reanalysis to extend the availability of labeled data to train a CNN on S2D timescales. Dasgupta et al. (2020) recently demonstrated this technique with a CNN-based approach that improved historical reconstructions of the MJO, which was then used to explore its decadal variability. This type of semi-supervised approach can extend S2D predictability by increasing the availability of labeled data. Generative adversarial networks (GANs; Goodfellow et al. 2014) can also provide further skill in the augmentation of synthetic data that resemble observations or model output. GANs can also transform more widely available datasets to match the statistical properties and structure of sparser observations

(Wang et al. 2021). These models have the advantage of not requiring complex loss functions and can be used to extract salient properties within spatial data by visualizing the discriminator output (Bau et al. 2018), though further study is needed to assess the limitations of this approach (i.e., are we getting the right answer for the right reason; Borji 2018).

The identified sources and limits of predictability can be further explored using causal inference, such as Granger causality (McGraw and Barnes 2018) and causal hypothesis testing (Runge et al. 2019). These methods are advantageous because they can help determine the direction of causality and disentangle whether a signal is due to a direct causal relationship or simply autocorrelation. These techniques would be particularly useful in situations where the identified source or limit of predictability is not a recognizable climate pattern or is from an unknown mode of variability. Causal graphs, which are probabilistic graphical models that mimic domain knowledge graphs (Nowack et al. 2020), are another viable causal discovery approach. The temporal causality discovery framework (Nauta et al. 2019) uses an attention-based CNN to output a causality graph structure with time lag. This causal inference method could help differentiate between spurious correlations and direct causal associations in multivariate S2S/S2D time series information.

Uncertainty quantification is also an important consideration for ML predictability problems. Epistemic uncertainty, or uncertainty in the model parameters, can be addressed using Bayesian neural networks that specify a distribution over the network weights (Dusenberry et al. 2020). To address aleatoric uncertainty, or inherent noise in the data, data augmentation techniques can be used to transform existing data to create new samples for training (Wang et al. 2019). Ensembling methods also address this form of uncertainty; for example, training a set of ML models with different random initializations or optimizing an ML model to output a distribution instead of a single class or value (Kendall and Gal 2017, Foster et al. 2021).

We note that there is a tradeoff between physics-informed ML and knowledge discovery. A significant advantage of ML is its ability to learn from data, but providing stringent constraints can hamper knowledge discovery. However in certain cases, simple physical constraints are truly necessary (e.g., constraining precipitation to be nonnegative values). To encourage knowledge discovery and data driven approaches, it may be best to compare results using a range of physical constraints. In a prediction context, this relates to using ML to discover unknown modes of variability or interactions between modes that are associated with the integrative water cycle.

An additional goal of the approaches described above is the creation of automated ML workflows that can be easily shared across disciplines. Automated training, including automated hyperparameter tuning (e.g., Optuna; Akiba et al. 2019), is a powerful approach to solving "big data" problems and model optimization. Efficient testing allows the user to explore a wide range of ML approaches, lowers the entry barrier for non-ML experts, and encourages interdisciplinary collaborations. We anticipate these workflows will have potential benefits across ML areas of interest to DOE EESSD that extend beyond predictability and the water cycle. Overall, we foresee that ML for Earth system predictability will be characterized by achieving higher skill in prediction of the integrative water cycle, coupled with a deeper understanding of the physical mechanisms, causal relationships, and associated uncertainty in Earth system models on S2S/S2D timescales.

Machine learning to extend and understand the sources and limits of water cycle predictability on subseasonal-to-decadal timescales in the Earth system

Data Availability

All the code and tools generated in the process of addressing these science questions will be made available following FAIR principles. We will leverage tools such as zenodo and github repositories to share code, model output, and other processed datasets.

Suggested Partners/Experts

The following groups contain partners and experts in machine learning and Earth system predictability that could potentially provide expertise at a related webinar or workshop.

NCAR Computational and Information Systems Laboratory Analytics and Integrative Machine Learning group (Dr. David John Gagne, NCAR), National Science Foundation AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography (Dr. Amy McGovern, University of Oklahoma; Dr. Elizabeth Barnes, Colorado State University; and Dr. Imme Ebert-Uphoff, Colorado State University), and the NCAR CESM and Climate and Global Dynamics Laboratory Earth System Prediction Working Groups (Dr. Yaga Richter, NCAR; Julie Caron, NCAR).

References

Akiba, T., Sano, S., Yanase, T., Ohta, T. and Koyama, M., 2019. Optuna: A Next-generation Hyperparameter Optimization Framework. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19), https://doi.org/10.1145/3292500.3330701.

Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R. and Samek, W., 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. PloS ONE, 10(7), e0130140, https://doi.org/10.1371/journal.pone.0130140.

Barnes, E.A., Mayer, K., Toms, B., Martin, Z. and Gordon, E., 2020. Identifying Opportunities for Skillful Weather Prediction with Interpretable Neural Networks. NeurIPS AI4EARTH 2020, https://arxiv.org/abs/2012.07830.

Bau, D., Zhu, J.Y., Strobelt, H., Zhou, B., Tenenbaum, J.B., Freeman, W.T. and Torralba, A., 2018. GAN dissection: Visualizing and understanding generative adversarial networks. arXiv preprint, arXiv:1811.10597.

Borji, A., 2018. Pros and Cons of GAN Evaluation Measures. arXiv preprint, arXiv:1802.03446.

Dasgupta, P., Metya, A., Naidu, C.V., Singh, M. and Roxy, M.K., 2020. Exploring the long-term changes in the Madden Julian Oscillation using machine learning. Scientific Reports, 10, 18567, https://doi.org/10.1038/s41598-020-75508-5.

Dusenberry, M.W., Jerfel, G., Wen, Y., Ma, Y.-A., Snoek, J., Heller, K., Lakshminarayanan, B. and Tran, D., 2020. Efficient and Scalable Bayesian Neural Nets with Rank-1 Factors. arXiv preprint, arXiv: 2005.07186.

Foster, D., Gagne, D.J. and Whitt, D., 2021. Probabilistic Machine Learning Estimation of Ocean Mixed Layer Depth from Dense Satellite and Sparse In-Situ Observations. Journal of Advances in Modeling Earth Systems, in review, https://doi.org/10.1002/essoar.10505859.1.

Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and

Bengio, Y., 2014. Generative adversarial networks. arXiv preprint, arXiv:1406.2661.

Ham, Y.G., Kim, J.H. and Luo, J.J., 2019. Deep learning for multi-year ENSO forecasts. Nature, 573(7775), 568-572, https://doi.org/10.1038/s41586-019-1559-7.

Kapur, A. and Zhang, C., 2012. Multiplicative MJO forcing of ENSO. Journal of Climate, 25(23), 8132-8147, https://doi.org/10.1175/JCLI-D-11-00609.1.

Kendall, A. and Gal, Y., 2017. What uncertainties do we need in Bayesian Deep Learning for Computer Vision? arXiv preprint, arXiv:1703.04977.

Kim, H., Ham, Y.G., Joo, Y.S. and Son, S.W., 2021. Deep Learning for bias correction of MJO prediction, in revision.

LeCun, Y., Bengio, Y. and Hinton, G., 2015. Deep learning. Nature, 521, 436-444, https://doi.org/10.1038/nature14539.

Liu, X., Wang, X. and Matwin, S., 2018. Improving the interpretability of deep neural networks with knowledge distillation. 2018 IEEE International Conference on Data Mining Workshops (ICDMW), 905-912, https://doi.org/10.1109/ICDMW.2018.00132.

Mayer, K.J. and Barnes, E.A., 2020. Subseasonal Forecasts of Opportunity Identified by an Interpretable Neural Network. Submitted to Geophysical Research Letters, preprint available at https://www.essoar.org/doi/10.1002/essoar.10505448.1.

McGovern, A., Lagerquist, R., Gagne, D.J., Jergensen, G.E., Elmore, K.L., Homeyer, C.R. and Smith, T., 2019. Making the black box more transparent: Understanding the physical implications of machine learning. Bulletin of the American Meteorological Society, 100(11), 2175-2199, https://doi.org/10.1175/BAMS-D-18-0195.1.

McGraw, M.C. and Barnes, E.A., 2018. Memory matters: A case for Granger causality in climate variability studies. Journal of Climate, 31(8), 3289-3300, https://doi.org/10.1175/JCLI-D-17-0334.1.

Meehl, G.A., Hu, A. and Tebaldi, C., 2010. Decadal prediction in the Pacific region. Journal of Climate, 23(11), 2959-2973, https://doi.org/10.1175/2010JCLI3296.1.

Meehl, G. A. et al., 2021. Initialized Earth system prediction from subseasonal to decadal timescales. Nature Reviews Earth and Environment, Accepted.

Nauta, M., Bucur, D. and Seifert, C., 2019. Causal discovery with attention-based convolutional neural networks. Machine Learning and Knowledge Extraction, 1(1), 312-340, https://doi.org/10.3390/make1010019.

Nowack, P., Runge, J., Eyring, V. and Haigh, J.D., 2020. Causal networks for climate model evaluation and constrained projections. Nature Communications, 11, 1415, https://doi.org/10.1038/s41467-020-15195-y.

Pegion, K., Kirtman, B.P., Becker, E. et al., 2019. The Subseasonal Experiment (SubX): A Multimodel Subseasonal Prediction Experiment, Bulletin of the American Meteorological Society, 100(10), 2043-2060, https://doi.org/10.1175/BAMS-D-18-0270.1.

Pendergrass, A.G., Meehl, G.A., Pulwarty, R. et al., 2020. Flash droughts present a new challenge for subseasonal-to-seasonal prediction. Nature Climate Change 10, 191-199, https://doi.org/10.1038/s41558-020-0709-0.

Robertson, A.W., Vitart, F. and Camargo, S.J., 2020. Subseasonal to Seasonal Prediction of Weather to Climate with Application to Tropical Cyclones. Journal of Geophysical Research: Atmospheres, 125(6), e2018JD029375, https://doi.org/10.1029/2018JD029375.

Runge, J., Bathiany, S., Bollt, E. et al., 2019. Inferring causation from time series in Earth system sciences. Nature Communications, 10, 2553, https://doi.org/10.1038/s41467-019-10105-3.

Simonyan, K., Vedaldi, A. and Zisserman, A., 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv preprint, arXiv:1312.6034.

Tang, Y., Deng, Z., Zhou, X., Cheng, Y. and Chen, D., 2008. Interdecadal variation of ENSO predictability in multiple models. Journal of Climate, 21(18), 4811-4833, https://doi.org/10.1175/2008JCLI2193.1.

Toms, B.A., Barnes, E.A. and Ebert-Uphoff, I., 2020. Physically Interpretable Neural Networks for the Geosciences: Applications to Earth System Variability. Journal of Advances in Modeling Earth Systems, 12(9), e2019MS002002, https://doi.org/10.1029/2019MS002002.

Wang, C., Tang, G. and Gentine, P., 2021. PrecipGAN: Merging Microwave and Infrared Data for Satellite Precipitation Estimation using Generative Adversarial Network. Geophysical Research Letters, 48, https://doi.org/10.1029/2020GL092032.

Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S. and Vercauteren, T., 2019. Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. Neurocomputing, 338, 34-45, https://doi.org/10.1016/j.neucom.2019.01.103.