Title:

Earth System Model Improvement Pipeline via Uncertainty Attribution and Active Learning

Authors:

Khachik Sargsyan (Sandia National Laboratories, Livermore, CA)
Daniel Ricciuto (Oak Ridge National Laboratory, Oak Ridge, TN)
Cosmin Safta (Sandia National Laboratories, Livermore, CA)

Focal Area:

Primary focal area 2 (Predictive Modeling via AI): We develop methods to formally quantify uncertainties in Earth System models for the land-atmosphere coupled system.

Science Challenge:

Earth system models still have significant biases in historical predictions of the intensity and frequency of water cycling extremes (*e.g.,* droughts and flood events), leading to low confidence in future projections. Uncertainties arise from incomplete understanding of land and atmospheric processes, and insufficient observational constraints on model parameters. Many observations, including those from key DOE investments such as ARM and AmeriFlux, are used to evaluate model performance but have not been used to formally quantify model uncertainty because of the expense of running ESM simulations. An efficient pipeline engaging cutting-edge machine learning (ML) and uncertainty quantification (UQ) methods is needed to improve the predictive understanding of water cycle extremes in the Earth system.

Rationale:

Several challenges currently hinder earth system model (ESM) predictability improvement. The heterogeneity of observational water cycle data in spatial and temporal scales makes it extremely difficult to make generalizable improvements to the models. Most frequently, climate modeling resorts to skill tables that describe the degree of agreement between a nominal model simulation and a list of observational datasets of various quantities of interest. Perturbing or tuning nominal model simulations may improve the model skill for one dataset while necessarily making it worse for another dataset. A significant reason behind this is structural uncertainty in models due to simplifying physical assumptions and parameterizations. **This work will account for structural uncertainties using embedded model error approach coupled with surrogate construction in order to remain non-intrusive to the code source.** These surrogates, otherwise called metamodels or emulators, are the necessary precursor of model improvement pipeline for computational models of complex multiphysics phenomena. Surrogate construction is a supervised ML procedure in which an ensemble of training simulations is performed before optimizing a parametric model ranging from polynomials to radial basis expansions to highly overparameterized neural networks. When it comes to coupled climate models, the major challenge is the expense of a single simulation, making an accurate surrogate construction an extremely difficult task. **As such, we see efficient and accurate surrogate modeling as the major opportunity for improving predictive capabilities of ESMs through model-data integration.** We will leverage recent ML advances that have led to substantial improvements in a wide range of disciplines but have not been explored and enhanced in an ESM context. Our vision towards improving model development via ML methods is depicted in Figure 1. The systematic implementation of these existing technical components in a coupled ESM will provide a coherent

DOI: https://doi.org/10.2172/1769699
DOE BER, AI4ESP 2021 #AI4ESP1117 (https://www.ai4esp.org)

framework with a groundbreaking potential to improve ESM predictability by integrating land and atmosphere water cycle observations, such as spatiotemporal patterns of precipitation, latent heat fluxes, runoff, and soil moisture from AmeriFlux. The framework is non-intrusive and hence flexible for use in future model versions as well as amenable to new hardware architectures.

Narrative:

The proposed work will hinge on **uncertainty estimation and attribution for key water cycle quantities of interest** at all stages of model development. It will inform decision making, optimal design of experimental campaigns (*e.g.,* new AmeriFlux towers) and help evaluate risks associated with extreme events. In this regard, the Bayesian probabilistic viewpoint paves the way for formal reasoning about uncertainties that stem from heterogeneous sources. In the proposed workflow, the Bayesian methods make an appearance twice: (a) when calibrating ESMs with observational data, which necessitates *a priori* construction of surrogates due to ESM expense, and (b) when building those surrogates themselves via, *e.g.,* highly parameterized neural network (NN) and tensor network (TN) models. Computational challenges associated with augmenting NN/TN surrogates with uncertainties will be alleviated by approximate and empirical methods, such as variational inference, deep ensembles, Bayesian last-layers. Uncertainty attribution will enable targeted model improvement steps. For example, a classical variance-based decomposition of uncertain ESM predictions, also called global sensitivity analysis (Sargsyan, 2014; Ricciuto, 2018), will attribute model prediction uncertainties to various uncertain components of the modeling pipeline, including parametric uncertainties, data noise, structural errors and uncertainties due to surrogate approximations to the model. In the proposed framework, uncertainties will play a crucial role in selecting the next batch of model evaluations via active learning techniques, described later in this text.

At the core of the proposed work is the **embedded model structural error estimation,** which is critical for producing accurate confidence intervals for water cycle extremes. While addressing parametric uncertainties has become relatively routine in computational science and climate modeling specifically, the estimation of structural uncertainties is largely unresolved. External statistical model corrections often do not preserve the physical laws and are always specific to output quantities of interest, thus they do not generalize meaningfully to the full set of model outputs. In this regard, embedded model error approaches provide a rigorous framework to augment models internally with statistical corrections, targeting components that are believed to be leading contributors to structural errors, informed either by domain scientists or by sensitivity analysis (Sargsyan, 2015; Sargsyan, 2019). Note that embedded does not mean intrusive source code editing: as well as for a typical model tuning exercise, embedded model corrections are performed with pre-constructed surrogate approximations of the model outputs.

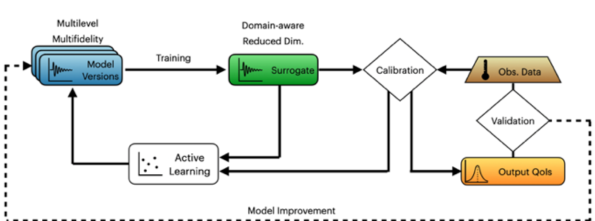

Figure 1: A schematic of the proposed path towards improved ESM predictability

We will perform **active learning (AL) of ESM surrogate** to develop model approximations with fewest possible model simulations, enabling efficient use of computational resources. AL has gradually become an integral part of many practical ML pipelines. It has been well-developed mainly in classification context, *e.g.*, to minimize the amount of labeling by a human in image libraries (Settles, 2009). Such labeling is the equivalent of an expensive climate simulation. The typical objective function for selecting the next batch of model simulations is driven by approximate or empirical uncertainties of the current ML surrogate form. We will augment classical uncertainty-based objective functions with constraints driven by the comparison with observational data, *e.g.*, gridded benchmarks and relationships between variables from ILAMB (Collier, 2018). This will allow efficient exploitation of the parameter space leading to accurate and usable surrogates with minimal amount of expensive coupled model simulations.

To further enhance surrogate accuracy under computational budget constraints, we will deploy **multilevel-multifidelity (MLMF) methods** that have recently been very successful for constructing model approximations for UQ (Huan, 2018). The MLMF methods benefit from the availability of model hierarchies, both in terms of multiple fidelities (submodel ML vs model parameterization vs resolved physics) and multiple levels (in an ESM context, spatial resolution). ESMs such as E3SM already support multiple resolutions, while both land and atmosphere model parameterizations vary by orders of magnitude in computational expense. Relying on the underlying premise of correlations between different levels or fidelities, we will achieve much more accurate high-fidelity model approximations while balancing the number of model evaluations at different fidelities/resolutions.

Furthermore, we hypothesize that using generally ***nonlinear*** **autoencoders (AE)** will allow a much more efficient dimensionality reduction of high-dimensional spatio-temporal fields of climate QoIs. We have achieved significant dimensionality reduction from 500,000 grid cells to 180 eigendirections without significant loss of accuracy in global simulations of the E3SM land model (Ricciuto, 2020). Earth system modeling has historically benefited from such linear dimensionality reduction techniques via principal component analysis, Karhunen-Loève expansions and empirical orthogonal functions (Monahan, 2009), which are all examples of linear autoencoders. Nonlinear AE techniques will lead to drastic improvements when it comes to extracting fingerprints from spatio-temporal model outputs and observational data. The requisite dimensionality reduction will be achieved by working in the latent, encoded space for more accurate and trainable surrogates, as well as for constructing model-vs-data decorrelated comparison metrics for subsequent calibration and validation. Furthermore, variational AEs will provide the necessary framework for augmenting the dimensionality reduction with uncertainty estimates (Weiling, 2019).

Finally, we will embed **domain-aware constraints** to ensure ML surrogates adhere to fundamental constraints and relations between variables that are inherent in ESMs (*e.g.*, hydrologic balance of precipitation, runoff, evapotranspiration and water storage in a watershed). Such consistency will be achieved via (a) soft constraints, augmenting ML training objectives with terms measuring the constraint violations, and (b) hard constraints, which reconfigure the ML parametric forms by physics-driven network topologies (*e.g.,* maintaining a positive relationship between leaf area and transpiration) or handcrafted functional terms (*e.g.*, use an exponential last layer to assert transpiration is never negative). This will ensure that predictions from the trained ML models are reliable and grounded in physics, leading to improved decision support and facilitating scientific discovery.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

References:

Collier, N., F.M. Hoffman, D.M. Lawrence, G. Keppel-Aleks, C.D. Koven, W.J. Riley, M. Mu, J.T. Randerson. "The international land model benchmarking (ILAMB) system: Design, theory and implementation. Journal of Advances in Modeling Earth Systems, 10 (11), 2731-2754.

Ghahramani, Z. (2015), "Probabilistic machine learning and artificial intelligence". Nature, 521, 452–459.

Huan X., G. Geraci, C. Safta, M. Eldred, K. Sargsyan, Z. Vane, J. Oefelein and H. Najm (2018), "Multifidelity Statistical Analysis of Large Eddy Simulations in Scramjet Computations", AIAA Non-Deterministic Approaches Conference. AIAA 2018-1180.

Monahan, A. H., J. Fyfe, M. Ambaum, D. Stephenson and G. North (2009), "Empirical Orthogonal Functions: The Medium is the Message", Journal of Climate, 22(24), 6501-6514.

Ricciuto, D., K. Sargsyan and P. Thornton (2018). The impact of parametric uncertainties on biogeochemistry in the E3SM land model. Journal of Advances in Modeling Earth Systems, 10, 297– 319.

Ricciuto D., D. Lu, K. Sargsyan (2020), "Surrogate Modeling and Machine Learning for Uncertainty Quantification", https://e3sm.org/surrogate-modeling-and-machine-learning-for-uncertainty-quantification.

Sargsyan K., C. Safta, H. Najm, B. Debusschere, D. Ricciuto and P. Thornton (2014), "Dimensionality reduction for complex models via Bayesian compressive sensing", International Journal for Uncertainty Quantification, 4(1): 63—93.

Sargsyan, K., H. Najm and R. Ghanem (2015), "On the Statistical Calibration of Physical Models". International Journal of Chemical Kinetics, 47: 246-276.

Sargsyan, K, X. Huan, H. Najm (2019), "Embedded model error representation for Bayesian model calibration". International Journal for Uncertainty Quantification, 9 (4), 365-394.

Settles, B. (2009), "Active learning literature survey". Computer Sciences Technical Report 1648, University of Wisconsin– Madison.

Welling, M., D. Kingma (2019), "An Introduction to Variational Autoencoders". Foundations and Trends in Machine Learning, 12 (4): 307–392.