

Development of Explainable, Knowledge-Guided AI Models to Enhance the E3SM Land Model Development and Uncertainty Quantification

Dali Wang¹, Shih-Chieh Kao¹, and Danial Riccuito¹

¹*Oak Ridge National Laboratory, Oak Ridge, TN, USA*

Focal Area(s)

(2) Predictive modeling using AI techniques and AI-derived model components; use of AI and other tools to design a prediction system comprising of a hierarchy of models. (3) Insight gleaned from complex data (both observed and simulated) using AI, big data analytics, and other advanced methods, including explainable AI and physics- or knowledge- guided AI.

Science Challenge

The Energy Exascale Earth System Model (E3SM) is a fully coupled, state-of-the-science Earth system model that uses code optimized for DOE's advanced computers to address the most critical scientific questions facing our nation and society (Golaz et al., 2019). The E3SM Land model (ELM) is designed to understand how the changes in terrestrial land surfaces will interact with other Earth system components and has been used to understand hydrologic cycles, biogeophysics, and ecosystem dynamics.

In spite of great successes, the ELM has several known issues that restrain rapid improvements. For example, the ELM uses equilibrium models to simulate dynamic land-climate interactions and it requires long model spin-up time to identify suitable initial conditions for transient simulations. The ELM lacks built-in uncertainty mechanisms that can improve the robustness of model predictions. The ELM is a holistic, deterministic model system with a rigid design, and in many situations, it is hard to modify the ELM system to incorporate new theory/hypothesis and new data across scales to address emerging science problems (such as predicting the impacts of water cycle extremes). In addition, The ELM is technically optimized for traditional CPU-centric computers and it cannot fully utilize the current and incoming leadership computers for model simulations and uncertainty quantification (UQ).

The success of artificial intelligence (AI) has inspired scientists to use AI models to discover intrinsic features from simulation data (Chattopadhyay et al., 2020) and observational data (Reichstein et al., 2019) to gain further process understanding of Earth science problems. However, autonomous AI model training through deep learning usually requires a huge amount of annotated data. To overcome the limitations from the data and computing resources, knowledge-guided AI models are necessary where human-knowledge is ingested in model construction (Banino et al., 2018) and training process (Silver et al., 2016) for efficient learning.

Herein, we present a new way that leverages the process understanding from the ELM to guide AI model development for the ELM enhancement and UQ. We hope this study can inspire further Earth and environmental system model developments and transformations.

Rationale

Essentially, we treat neural networks as a way to represent complex mathematical functions between datasets, similar to computational graphs that use nodes and edges to present associations. We then construct and deploy biology- and physics-inspired neural networks (i.e., Hochreiter and Schmidhuber, 1997, Krizhevsky et al, 2012, Banino et al, 2018, Rassi et al., 2019) and apply backpropagation representation learning procedure (Rumelhart et al., 1986) to reveal meaningful

intrinsic structures of high-dimensional data and to discover multiple levels of abstraction among datasets (LeCun et al, 2015).

In this study, we use the conceptual model and key ELM functions to create Prototype AI Models (PAMs), with customized neural network architectures, to represent the current understanding of Earth system processes. These PAMs will provide their own explanation and be faithful to what the current ELM actually computes. We have developed a novel method (Wang et al, in review) that adopts a hierarchical deep reinforcement learning (HDRL) (Kulkarni et al., 2016, Mnih et al., 2015), combined with a greedy approach to introduce subgoals across scales, to build a new model system from pre-trained AI models. This method has been applied to biology and discovered a novel mechanism with small training set, minimal labeling and simple rules/constraints. We would like to further develop this method and use HDRL as an unsupervised new model formation platform to build Knowledge-guided AI Models (KAMs) that can alleviate certain modeling restrictions caused by incomplete process knowledge, limited data, and uncertainty, etc. Our method contains two phrases:

PAM creation with supervised learning: Current key ELM functions present mathematical formulas (equations or logics) that approximate the observed behavior of given terrestrial processes through iterative scientific methods (observation, model, and experiment). We first design neural networks and construct PAMs as an approximation of key ELM functions via supervised learning with data collected from the ELM simulations. Our knowledge of these ELM functions will direct the neural network design in the PAM construction and also introduce better inductive bias in supervised training processes. The mathematical formula of the ELM functions can also help to explain the PAM behaviors over observation data.

KAM formation with HDRL: Following the current ELM design, we can construct KAMs with these PAMs. Considering the hierarchical structure of the KAMs, we will develop a multi-level HDRL (such as Vezhnevets et al., 2017) model that uses rewards associated with multi-level subgoals that represent local and longer-term feedback from the underlying ecosystem processes. Backpropagation of the rewards through time via HDRL will fine-tune the PAMs simultaneously for essential feature capture and to form KAMs with much relaxed assumptions on the underlying mechanism of terrestrial ecosystem processes. The HDRL-trained KAMs can be applied to ELM development and UQ. These HDRL-trained PAMs can also be transferred to other applications for further understanding of terrestrial ecosystem processes.

Narratives

Here, we present some technical details on application of our method to build PAMs and to form KAMs for the ELM development and UQ. We then propose several scenarios to test whether the HDRL-trained PAMs can be transferred into new applications to address emerging science questions. The development procedure (Figure 1) contains the following four major steps.

1. ELM data conversion and utility function development: ELM uses derived data types and customized functions to store and manipulate data and model parameters. While AI models use neural networks to process data in a “tensor” format, which can be easily understood as high-dimensional arrays. Data conversion from current ELM datatypes into tensors is necessary. Basic functions (sampling, subgrid, filter, atm2lnd) for the forcing data and the hierarchical landscape data (grid, column, PFTs) can be implemented as computing neuros.

2. PAMs creation with supervised learning: The ELM simulates the interactions and processes of terrestrial systems and uses mathematical questions, logics, and environmental and biological

variables to present system dynamics. Following the ELM conceptual model, PAMs can be created to represent key ELM functions (such as soil hydrology, soil temperature, canopy, photosynthesis) and to track the change of key variables. Neural networks will be designed to emphasize and track the signature patterns of the key ELM variables to improve model’s explainability. The ELM simulation data, collected from the Functional Unit Testing Framework (Wang et al., 2014), can be used as labelled data for supervised training. We expect the trained PAMs to reveal meaningful features of high-dimensional ELM data (such as these state and flux variables of energy, water, and carbon). Considering the unique settings of neural networks (i.e., all the neurons in each network layer share the same parameters), we expect higher computational efficiency, when the PAMs are deployed on high-end computing facilities with accelerators.

3. *KAM construction based on the ELM design:* Based on the ELM conceptual model and software design, we can build KAMs (such as Hydrology and Flux) from these individual PAMs. The connections between PAMs are highly flexible and reconfigurable, since their input and output can be customized by either data masking or neural network weight manipulation. The spin-up process of the KAMs will be rapidly accelerated using gradient descent and regularizations (such as dropout, disturbance ingestion, and weight decay).

4. *KAM formation with HDRL:* Considering the hierarchical structure of the KAMs, we will develop a multiple-level HDRL, combined with a greedy approach to introduce subgoals across scales, to train KAMs from these PAMs. The knowledge on the ELM functions can be used to create a weak form of subgoals (with appropriate reward functions) to constrain the interactions between individual PAMs. Available observation data across scales can also be used to provide sparse and delayed feedback into the HDRL training procedure. The robustness of the KAMs can be enhanced by introducing uncertainty in several ways, such as data and parameter disturbance, feature dropout, noise ingestion into model responses, as well as the modification of AI model connections.

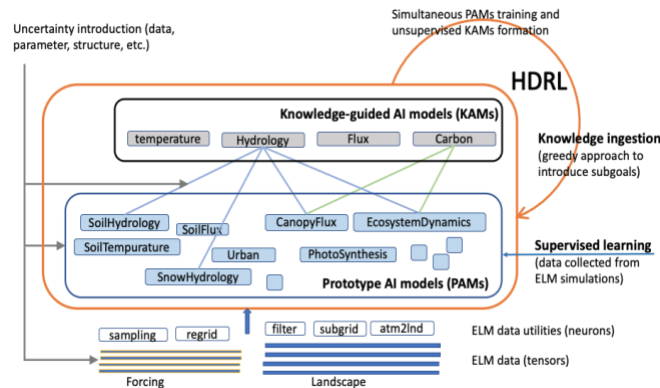


Figure 1. KAMs development from the PAMs for ELM diagnostics and UQ.

KAMs provide unique valuables and flexibilities for ELM diagnostics and UQ. For example, we can use KAMs to reveal meaningful features from landscape data over multiple spatial resolutions and from forcing data with variable lengths. We can use KAMs to generate validation datasets for ELM subcomponents. KAMs can also be used to quantify the uncertainty from a given subcomponent of the ELM. In addition, we would like to test whether the KAMs can be applied to address emerging science questions. For that purpose, we would construct a regional water cycle KAM to assess the impact of seasonal variability of surface water budget and hydroclimatic trend, or we could construct a water respiration KAM to estimate the canopy-level transpiration with the [FLUXNET data](#).

Reference

- Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., Pritzel, A., Chadwick, M. J., Degris, T., Modayil, J., and Wayne, G. 2018. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705), pp.429-433.
- Chattopadhyay, A., Hassanzadeh, P., and Pasha, S. 2020. Predicting clustered weather patterns: A test case for applications of convolutional neural networks to spatio-temporal climate data. *Scientific reports*, 10(1), pp.1-13.
- Golaz, J. C., Caldwell, P. M., Van Roekel, L. P., Petersen, M. R., Tang, Q., Wolfe, J. D., Abeshu, G., Anantharaj, V., Asay-Davis, X. S., Bader, D. C., and Baldwin, S. A. 2019. The DOE E3SM coupled model version 1: Overview and evaluation at standard resolution. *Journal of Advances in Modeling Earth Systems*, 11(7), pp.2089-2129.
- Hochreiter, S., and Schmidhuber, J. 1997. LSTM can solve hard long time lag problems. *Advances in neural information processing systems*, pp.473-479.
- Kratzert, F., Klotz, D., Shalev, G., Klambauer, G., Hochreiter, S., and Nearing, G. 2019. Benchmarking a catchment-aware long short-term memory network (LSTM) for large-scale hydrological modeling. *Hydrology and Earth System Sciences Discussions*, pp.1-32.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*. 1: 1097–1105.
- Kulkarni, T. D., Narasimhan, K., Saeedi, A., and Tenenbaum, J. 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. *Advances in neural information processing systems*, 29, pp.3675-3683.
- LeCun, Y., Bengio, Y., and Hinton, G. 2015. Deep learning. *nature*, 521(7553), pp.436-444.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., and Petersen, S. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540), pp.529-533.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, pp.686-707.
- Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., and Carvalhais, N. 2019. Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743), pp.195-204.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. 1986. Learning representations by back-propagating errors. *nature*, 323(6088), pp.533-536.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., and Dieleman, S. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587), pp.484-489.
- Vezhnevets, A. S., Osindero, S., Schaul, T., Heess, N., Jaderberg, M., Silver, D., and Kavukcuoglu, K. 2017. FeUdal Networks for Hierarchical Reinforcement Learning. In ICML.
- Wang, D., Xu, Y., Thornton, P., King, A., Steed, C., Gu, L., and Schuchart, J. 2014. A functional test platform for the Community Land Model. *Environmental Modelling & Software*, 55, 25-31.

Wang, Z., Xu, Y., Wang, D., and Bao, Z. (in review). Hierarchical deep reinforcement learning of time-lapse images reveals novel mechanism of cell movement, A preprint is available at <https://docs.google.com/document/d/1ubDcWWvNABKmf1iHAHNTYmnsderNEvhj1tp0eZ-pNmY/edit?usp=sharing>