Black-Box Neural System Identification and Differentiable Programming to Improve Earth System Model Predictions

Authors

Nathan Urban (BNL), Anthony DeGennaro (BNL), Yangang Liu (BNL)

Focal Area(s)

Focal Area 2: AI for predictive modeling, including AI-derived model components, and AI-enabled applications including parameter optimization, data assimilation, and uncertainty quantification.
Focal Area 1: AI-assisted data assimilation using numerical Earth system models.

Science Challenge

Earth system models have structural biases that lead to uncertain predictions, and their complexity and expense makes it difficult constraining the models with data or improved physical understanding.

Rationale

Earth system models (ESMs) are approaching exascale computing requirements, targeting DOE leadership-class high performance computing to achieve the highest resolution and process fidelity. Even when state-of-art algorithms are applied, the computational expense of ESMs creates a significant practical barrier to basic tasks like model tuning (parameter optimization), let alone advanced predictive artificial intelligence (AI) applications such as reduced order modeling (ROM), uncertainty quantification (UQ), data assimilation, or online training of hybrid physical/machine-learning (ML) models.

Many of the most advanced modern algorithms for these applications are ‘intrusive’ in nature, meaning they require access to the analytic governing equations describing Earth system dynamics, or at least model source code. These include projection-based ROMs, and adjoint-accelerated or differentiable programming data science methods that exploit model gradients or higher-order derivative information. An obstacle to applying advanced intrusive data science algorithms to Earth system models is that ESMs may be too complex to reliably apply methods that require the analytic manipulation of governing equations, or even source code tracing algorithms such as automatic differentiation (AD).

To apply more scalable intrusive AI methods to ESMs, we propose that ML-based system identification methods can be applied to learn to emulate the governing equations of an ESM, in a ‘black-box’ or ‘non-intrusive’ setting that only requires access to easily generated simulation output. The ML emulator can be reduced, differentiated, etc. even when the original ESM cannot. Intrusive model reduction or gradient-accelerated / adjoint-based algorithms then can be applied to the learned ML model much more easily, and possibly with greater computational efficiency, than they could to the original ESM.

DOI: https://doi.org/10.2172/1769681
DOE BER, AI4ESP 2021 #AI4ESP1134 (https://www.ai4esp.org)

Black-Box Neural System Identification and Differentiable Programming to Improve Earth System Model Predictions

Narrative

Reduced order modeling methods provide simplified approximations to an expensive model that can, in principle, be applied in a hierarchy of models to accelerate applications in UQ, model tuning, and data assimilation, or embedded within multisector dynamics models. Projection-based ROMs are able to exploit physical knowledge of the system by formally projecting the governing equations of a model onto a low-order reduced basis. This technique is intrusive, requiring an analytic description of the system dynamics. Many of the most scalable optimization and uncertainty quantification algorithms rely on being able to compute derivatives of model outputs with respect to inputs, such as stochastic gradient descent, Newton optimization, Hamiltonian Monte Carlo, variational inference, and variational data assimilation. Derivative information is obtained intrusively via analytically differentiation for simple models, or via reverse-mode AD applied to the source code of more complex models, which are also intrusive methods.

Scientific machine learning is moving toward hybrid simulations that embed data-driven machine learning models inside physics codes to improve their accuracy (see the whitepaper "Building an AI/ML/HPC-enhanced modeling framework to address multiscale predictability challenges" by Yangang Liu et al.). Examples of complex and uncertain model structures that are good targets for data-driven improvements abound within atmospheric models, including cloud microphysics, planetary boundary layer and turbulence closure schemes, and 3D radiative transfer. Typically, hybrid simulations are developed via offline training of a neural network to some higher-fidelity reference data, such as cloud resolving model output. However, some online training may be required to produce a stable simulation when the physics and ML components are coupled, or when there does not exist a reference data set (for example, learning an unknown correction term buried in a parameterization code that does not correspond directly to any observable quantity). Online training of hybrid models also requires intrusive gradient methods, such as end-to-end differentiable programming, which computes model loss gradients via a combination of backpropagation through the neural network and adjoint simulation through the physics code. Backpropagation-enabled online training also facilitates quantification of model structural uncertainties within hybrid models, via the gradient-enhanced UQ techniques described earlier.

The difficulty in applying intrusive methods to Earth system models is their great complexity. While projection-based model reduction has been applied to simple fluid dynamics systems for which the governing equations are easy to write down and manipulate, much of the complexity in an Earth system model is contained in its sub-grid scale physics parameterizations, which can involve hundreds of thousands of lines of code and do not yield a tractable, analytically manipulable form. Similarly, while reverse-mode automatic differentiation systems have existed for decades, it can be difficult to apply them to the complex mix of scalar and vector code with complicated conditional logic that exists within large ESM codebases. This has contributed to adjoint modeling being considered onerous and fragile in ESM applications, despite its theoretical computational advantages.

Black-Box Neural System Identification and Differentiable Programming to Improve Earth System Model Predictions

To facilitate the application of intrusive methods to ESMs, it may be necessary to treat the ESM as a black box, assuming that the internals of a production code are too complex to manipulate directly. Even in a black box setting, it may be possible to infer (a statistical approximation to) the equations the ESM is solving from the ESM's output alone. An ESM's equations can be very complicated in discretized form, including all its numerics and sub-grid parameterizations, but essentially remains a PDE solver with a very complicated right-hand side that in principle can be inferred from data.

System identification is a discipline of applied mathematics concerned with recovering the equations of motion of a system from observations of its solution. Some forms of system identification are concerned with exactly recovering the original equations in analytic form. This is not likely to be useful in the setting discussed here; the problem is not that the equations are not known (they are, at least in the form of their source code implementation), but that they are too complex to manipulate. Other forms of system identification, however, only attempt to recover a statistical approximation to the original equations. For example, we can consider treating the right-hand side of a PDE as an unknown function, and attempt to approximate that function using a deep neural network. Such a neural network can be trained by asking it to reproduce the action of the ESM upon its state variables over the course of a model timestep.

Once in statistical, neural form, the ESM's governing equations can be directly manipulated similarly to how the underlying analytic equations can be. They can be projected onto a low-dimensional basis set to construct ROMs. They can be differentiated, through standard neural network backpropagation techniques provided by AD packages such as TensorFlow and PyTorch, for use in gradient-accelerated or adjoint-based techniques. Differentiable programming for the online training of hybrid physical/ML models becomes pure backpropagation, when the physical model and its ML component are both in neural form. Long-term research challenges are similar to those in other emerging areas of scientific ML, including ensuring that the identified system dynamics are sufficiently accurate, numerically stable, respect known physical constraints, and do not unphysically extrapolate beyond their training set.

References

A. DeGennaro, N.M. Urban, B.T. Nadiga, and T. Haut, "Model structural inference using local dynamic operators", *International Journal on Uncertainty Quantification* **9**, 59 (2019).

A. DeGennaro (with N.M. Urban), "LDO inference from a quantity of interest using the adjoint-gradient optimization", Los Alamos Tech Report LA-UR-18-24884 (2018).

V.M. Krasnopolsky, M.S. Fox-Rabinovitz, and D.V. Chalikov, "New approach to calculation of atmospheric model physics", *Monthly Weather Review* **133**, 1370 (2005).

P. Melland, E.J. Albright, and N.M. Urban, "Differentiable programming for online training of a neural artificial viscosity function within a staggered grid Lagrangian hydrodynamics scheme", *Machine Learning: Science & Technology*, in press (2020). (DOI: 10.1088/2632-2153/abd644)