The Usage of Observing System Simulation Experiments and Reinforcement Learning to Optimize Experimental Design and Operation

Authors

Joseph C. Hardin, PNNL; William I. Gustafson Jr., PNNL; Adam Varble, PNNL; Jim Marquis, PNNL; Jennifer Comstock, PNNL; Andrew Geiss, PNNL; Sam Silva, PNNL; Po-Lun Ma, PNNL, Scott E. Giangrande, BNL; Christine Sweeney, LANL; V. Chandrasekar, CSU

Focal Area(s)

The focal area is data acquisition and assimilation enabled by machine learning, in particular focusing on network design and optimization, as well as insights gained from complex data.

Science Challenge

Various organizations regularly conduct field campaigns across the globe designed to probe and improve atmospheric process understanding. However, these campaigns are mostly designed ad-hoc, and rely on anecdotes and knowledge of what was successful in previous campaigns. Such ad-hoc experiment design results in sub-optimal instrument siting and operation strategies. Inadequate targeted data collection of extreme weather is a major limitation to Earth and Environmental Systems Science Division (EESSD)'s goals of data model integration (4.5.3) [8], which limits the knowledge gained through the collected observations, holding back significant advances in predictability. While advances in instrument design are always chipping away at these limitations, we are not optimally utilizing the instrumentation we currently possess.

For relatively infrequent but highly impactful severe weather, it is critical to maximize chances of observing details in a way that optimizes process-level understanding and leads to accurate predictions. For many similar challenges in the integrated water cycle, where multiple instrument data-streams must be synthesized together, this ad-hoc uncoordinated deployment of instruments is problematic. A more holistic design framework is needed to optimize accurate, detailed measurements and predictions. Currently, limitations in our data capture is one of the fundamental limiting factors in our understanding of atmospheric science.

More recently, investigators have started to incorporate Observing System Simulation Experiments (OSSEs) [1], [2], where we combine Earth system model runs and instrument simulators to understand the effect of deployment decisions to capture details of the atmospheric state. However, OSSEs typically employ designs constrained by typical instrument deployment concepts; thus, they often do not explore innovative observing strategies that could maximize utility. We use OSSE's to reduce the cost and effort of trying different configurations in the field, a process which can be expensive and risk failure of an

DOI: https://doi.org/10.2172/1769782
DOE BER, AI4ESP 2021 #AI4ESP1057 (https://www.ai4esp.org)

entire campaign if incorrect choices are made. However, current implementation of OSSE's using trial and error is unable to fully explore the design space.

Principal investigators spend years proposing, advocating for, and analyzing data from these field campaigns representing an investment of 100's of millions of dollars each year. **The solution is to develop an approach for using machine learning, in particular reinforcement learning [3], to leverage OSSEs to optimize the selection, deployment, and operation of field campaign instrumentation to best achieve mission goals for a field campaign**. This encompasses three distinct but related goals:

1. **Adaptive Operations:** Given a layout and set of instruments, how do we best react to changing conditions to capture phenomena
2. **Campaign Design:** Given a set of instrumentation, constraints, and science goals, how do we optimize the choice, location, and static operational strategies of these instruments.
3. **Continual Improvement:** Furthermore, as strategies are deployed in the field, How do we combine Reinforcement Learning with Edge Computing in distributed networks to allow algorithms to improve and adapt over time to these changing environmental conditions.

Rationale

1. Field campaign design is still more art than science and we do not target measurements of processes in an optimal or objective way. Similarly we do not have strategies for exploring possible experimental design spaces in an efficient and productive manner.
2. Instrument operation strategies (for instance radar scanning strategies) are constrained by past experience, without much experimentation, due to the high cost of field campaigns and the difficulty of testing out new strategies.
3. Earth system model evaluation and improvement is limited by the lack of targeted observations of key processes.

The set of decisions about how to design and operate a field campaign can be thought of as a set of policies. Machine learning has been making large strides in optimizing policy based learning through the field of reinforcement learning.

Progress in modeling and understanding the atmosphere has been limited by the challenge of obtaining adequate amounts of high quality observations of targeted phenomena. A large part of our understanding is driven by "case-study" observations rather than robust statistical descriptions. The case study approach is particularly limiting for extreme weather where a relative lack of sampling caused by less frequent occurrence and non-optimal sampling results in shortcomings in process understanding and validation of model predictions. Moreover, static observing strategies lack the capability to target high impact weather, while the design of adaptive systems has been slow to develop. By optimizing the way we collect these measurements and the design of field campaigns, we have the opportunity to significantly improve Earth system predictability through improved characterization of integrated water cycle processes via better evaluation of models and subsequent improvements in their ability to make predictions.

Narrative

Every year, vast sums of money are spent on field campaigns with a goal of improving understanding and prediction of atmospheric processes, but experimental designs are manual and mostly ad-hoc with reliance on designs of previous field campaigns and intuition. This results in sub-optimal experimental design that may ultimately under-sample the targeted phenomenon. The need for adaptive operations is widespread with examples including adaptive radar scanning for mesoscale convective system tracking, the timing and location of sonde launches for convective initiation, and placement and timing of aerial assets for in-situ operations (tethered balloon systems, unmanned aerial systems, etc). The set of decisions that go into experimental design (choice of instruments, siting) and their operating instructions (scan strategies, launch timing, adaptive operations, network latency) can be formulated as a set of policy decisions. The field of reinforcement learning (RL) has developed an extensive set of tooling for policy optimization. Fundamentally, we propose to optimize two components of field campaign design using a combination of observing system simulation experiments and reinforcement learning techniques.

1. The first goal is to optimize the layout and "static" operation of a field campaign. Generally this consists of the choice of instruments, their geographical siting, and selection from an a-priori set of operation strategies. Here, the OSSE (model + instrument simulators) forms the environment, while the instruments and strategies are the agents in a set of policy decisions to be optimized.
2. The second goal is to optimize the operation of individual instruments or sets of instrumentation in an adaptive manner by reacting to the environment around them. While adaptive operations are not new [3], the traditional implementation is via rule based systems and the design of these systems are usually major research efforts. This design can be done easier, and more comprehensively, using OSSE and RL techniques.

RL has been making large inroads in solving ever more complex problems. For instance, newer systems have managed to learn and optimize over systems for which they are not given explicit rules[4], [5], instrument controls [6], and building controls based on edge computing [7]. This would fundamentally change how we are able to design experimental field campaigns, speeding up scientific discovery by improving data quality and usefulness, and facilitating more effective usage of research funding. In addition, the OSSE provides direct linkages between measurements provided by a given operational strategy and critical processes of interest that are simulated but not observable. Thus, this framework also provides a means to more objectively and fruitfully evaluate parameterized processes in Earth system models.

Suggested Partners/Experts

This paper has a natural synergy with "Framework for an adaptive integrated observation system using a hierarchy of machine learning approaches" led by Dr. Comstock. While this paper addresses the design and real time control of adaptive scanning systems, the work in Comstock's paper would provide useful guidance on the phenomenology that we wish to optimize over.

References

[1] S.-A. Boukabara *et al.*, “Community Global Observing System Simulation Experiment (OSSE) Package (CGOP): Description and Usage,” *J. Atmospheric Ocean. Technol.*, vol. 33, no. 8, pp. 1759–1777, Aug. 2016, doi: 10.1175/JTECH-D-16-0012.1.

[2] X. Zeng *et al.*, “Use of Observing System Simulation Experiments in the United States,” *Bull. Am. Meteorol. Soc.*, vol. 101, no. 8, pp. E1427–E1438, Sep. 2020, doi: 10.1175/BAMS-D-19-0155.1.

[3] R. S. Sutton, A. G. Barto, and others, *Introduction to reinforcement learning*, vol. 135. MIT press Cambridge, 1998.

[4] OpenAI *et al.*, “Dota 2 with Large Scale Deep Reinforcement Learning,” *ArXiv191206680 Cs Stat*, Dec. 2019, Accessed: Dec. 28, 2020. [Online]. Available: http://arxiv.org/abs/1912.06680.

[5] “Mastering Atari, Go, chess and shogi by planning with a learned model | Nature.” https://www.nature.com/articles/s41586-020-03051-4 (accessed Dec. 28, 2020).

[6] M. G. Bellemare *et al.*, “Autonomous navigation of stratospheric balloons using reinforcement learning,” *Nature*, vol. 588, no. 7836, Art. no. 7836, Dec. 2020, doi: 10.1038/s41586-020-2939-8.

[7] A. Gupta, Y. Badr, A. Negahban, and R. G. Qiu, “Energy-efficient heating control for smart buildings with deep reinforcement learning,” *J. Build. Eng.*, p. 101739, Aug. 2020, doi: 10.1016/j.jobe.2020.101739.

[8] U.S. DOE. 2018. Climate and Environmental Sciences Division Strategic Plan 2018–2023, DOE/SC–0192, U.S. Department of Energy Office of Science (science.osti.gov/-/media/ber/pdf/workshop-reports/2018_CESD_ Strategic_Plan.pdf)