Autonomous reinforcement learning agents for improving predictions and observations of extreme climate events

André Goncalves[1], Donald Lucas[2]
1. andre@llnl.gov, Computational Engineering Division, LLNL
2. ddlucas@llnl.gov, Atmospheric, Earth, and Energy Division, LLNL

Primary Focus Area: This proposal addresses focus area 2, "Predictive modeling through the use of AI techniques."

Science Challenge: Extreme climate events associated with severe weather, coastal and inland flooding, droughts, heat waves and wildfires are expected to increase in frequency and severity in the future. Due to the complexity and chaotic behavior of the climate system, accurately predicting and observing extreme climate events requires a tremendous amount of human intervention to run predictive climate simulations and deploy measurement systems. Extreme events often unfold very quickly, leaving little time to iterate on simulations or re-position instruments. Through reinforcement learning, autonomous AI agents can be designed to make real-time decisions to characterize extreme climate events more efficiently through adaptive models and targeted observations.

Rationale: Climate is a complex interconnection of highly nonlinear dynamical systems including atmosphere, land surface, snow and ice, oceans, and other bodies of water. Given its complexity, the mechanisms by which extreme climate events occur are not yet fully understood. Reinforcement learning (RL) is a subfield of machine learning that is concerned with learning an intelligent agent that can make optimal sequential decisions in an unknown and noisy environment, so it maximizes its performance on a given task. RL has achieved breakthrough results in autonomous robotics, some of the most difficult games invented by humans, such as Chess and Go [Silver *et al.*, 2016], and real-time strategy, multi-player games including StarCraft II [Vinyals *et al.*, 2019]. Due to its nature, many research problems in the climate sciences, including the study of extreme climate events, can be posed as a sequential decision-making process. We highlight how some of these problems can be automated and tackled from a RL perspective.

Narrative: Extreme weather/climate-related events are the result of a combination of several dynamical processes interacting at different temporal and spatial scales. Therefore, learning how these systems relate at different time and space resolutions is essential for better predicting and observing extreme events. We envision at least four potential areas in climate-related research on extreme events that could benefit from reinforcement learning:

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

DOI: https://doi.org/10.2172/1769680
DOE BER, AI4ESP 2021 #AI4ESP1052 (https://www.ai4esp.org)

1. Automated and dynamic refinement of meshes used in climate and weather models to target and track evolving events and regions of interest.
2. Adaptive and dynamic selection of ensemble members for climate forecasts to capture variability and uncertainty in a computationally efficient way.
3. Optimal collection of high-value climate observations and measurements using efficient sequential data collection processes. An application of RL to collect data in the stratosphere autonomously was recently realized [Bellemare *et al.*, 2020].
4. Automated calibration of climate and weather models to match historical data and improve predictability of future projections.

As an example, we highlight in more detail the ways in which RL and intelligent agents can help automate and solve complex climate model mesh refinement problems for extreme events and dynamically evolving systems. In dynamical downscaling, climate and weather models needed for local scale forecasts require a high-resolution mesh, as shown in Figure 1, which are computationally intensive and even prohibitive for larger areas. Researchers usually use a high-resolution mesh only on a targeted small region and use a global coarse resolution mesh to account for other long-range climate phenomena that might be relevant for the prediction. Adaptive mesh refinement (AMR) strategies have been developed to refine resolution in limited regions without requiring a fine grid over the entire model domain. While AMR is useful, it can be tedious to setup and is prone to numerical issues when mesh cells become overly distorted or entangled in highly dynamical systems.

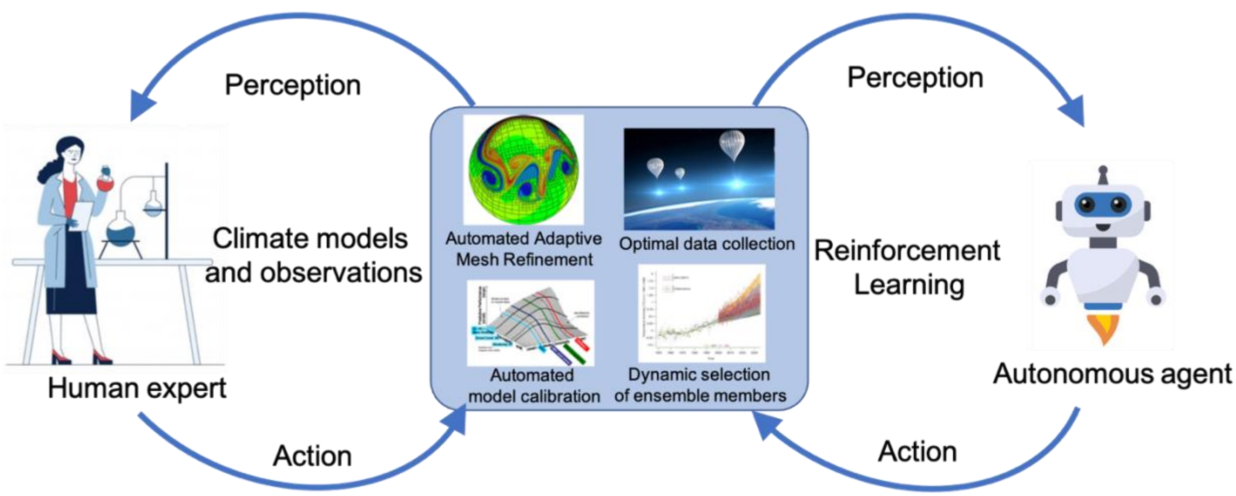

Figure 1 Current scientific approach on the left: human expert-centered; and proposed scientific process on the right: based on autonomous intelligent agents that employs reinforcement learning strategies for knowledge discovery.

An AI system of intelligent agents powered by reinforcement learning can overcome these issues by automatically and simultaneously tracking multiple dynamical climate/weather phenomena that may lead to extreme events such as flooding, fires and hurricanes in targeted regions. Once ready (trained) the system will act on top of climate models to determine where and when to refine the mesh to maximize extreme event predictability. Relevant areas identified by the AI system will be simulated at higher resolution than other less relevant areas. In this way, computational resources are directed to local areas with stronger impacts on the target area.

Moreover, intelligent RL agents can solve refinement problems that are beyond the capabilities of current AMR methods. If the origin of an extreme event or disturbance is located far from the target area through a teleconnection, RL agents can be designed to search out the source and refine the mesh on both sides of the connection. As multiple climate phenomena are expected to interfere with the local target area, the AI system can consist of a collection of intelligence agents that work jointly and collaboratively. A possible way to implement such an AI system is via multi-agent reinforcement learning [Da Silva *et al.*, 2019]. In this case, multiple intelligent agents will work collaboratively to explore the global climate system looking for and tracking phenomena that affect extreme events at the local target area. As all agents are acting and learning in the same climate system, each agent is influenced by the joint actions of all agents. Every reinforcement agent will be in charge of tracking a particular dynamical process and acting in a sequential manner: given the current climate situation, decide the next area to be refined. All agents may have a common goal, which is provide information to local scale extreme event prediction.

Like the mesh refinement benefits outlined above, we also expect similar breakthroughs and benefits to occur by applying RL techniques to the other climate modeling and observing topics, including optimal data collection, ensemble model selection, and automated calibration and tuning.

Computational resources and software: The proposed approach can be integrated with existing climate modeling systems, such as the DOE's Energy Exascale and Earth System Model (E3SM), with measurement and data acquisition systems used by the DOE's Atmospheric System Research program and Atmospheric Radiation Measurement (ARM) facilities and other observational systems and projects, and in the leadership class high performance computers used by DOE researchers for simulating climate change at ultra-high resolutions. Embedded within these systems, the automated RL agents will be able to automatically refine simulation meshes, collect the most useful observations, and make efficient use of supercomputing resources. Associated software suites will be made available to the community as an open-source tool.

Partners/Experts:

- This white paper is supported by Prof. Stuart Russell from U.C. Berkeley. Prof. Russell is co-author of the widely used textbook “Artificial Intelligence: A Modern Approach” and has expressed an interest in participating in future AI4ESP workshops.

References:

[Bellemare *et al.*, 2020] Bellemare, M.G., Candido, S., Castro, P.S. *et al.* Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature* 588, 77–82 (2020).

[Da Silva *et al.*, 2019] Da Silva, F.L., Costa, A.H. A survey on transfer learning for multiagent reinforcement learning systems. *Journal of Artificial Intelligence Research* 64:645-703 (2019)

[Silver *et al.*, 2016] Silver, D., Huang, A., Maddison, C. *et al.* Mastering the game of Go with deep neural networks and tree search. *Nature* 529**,** 484–489 (2016).

[Vinyals *et al.*, 2019] Vinyals, O., Babuschkin, I., Czarnecki, W.M. *et al.* Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575**,** 350–354 (2019).