Contribution ID: **40**   Type: **not specified**

A Checker-Board Sky: Reinforcement Learning Techniques for Telescope Scheduling

From Chess to Telescopes: Using reinforcement learning to automate the observation scheduling process

The size, complexity, and duration of telescope surveys are growing beyond the capacity of traditional methods for scheduling pointings and observations. Scheduling algorithms must have the capacity to balance multiple, often competing, observational and scientific goals, address both short-term and long-term considerations, and adapt to rapidly changing stochastic elements (e.g., weather). Reinforcement learning (RL) methods have the potential to significantly automate the scheduling and operation of telescope campaigns. In this work, we present the application of an RL-based scheduler, which uses a Markov decision process framework to construct scheduling policies in a way that is recoverable and computationally efficient for surveys that can include over a hundred observations.

We investigate and compare three RL policy optimizers: Proximal Policy Optimization (PPO), Evolutionary Solutions (EvoStrat), and Deep Q-Network (DQN). We show the success of EvoStrat, benefitting from its mutative and iterative method of exploration, which proves useful in a shallow loss landscape. Additionally, we examine how well an agent can learn a telescope's environment and produce results comparable to human-designed schedules, by comparing cumulative reward of different schedules and other metrics.

Primary authors: VOETBERG, Margaret (Fermi National Accelerator Laboratory);   ZHOU, Sophia (Mass Academy)

Co-authors: COHEN, Benjamin;   NEILSEN, Eric (Fermilab);   NORD, Brian (Fermilab)

Presenters: VOETBERG, Margaret (Fermi National Accelerator Laboratory);   ZHOU, Sophia (Mass Academy)

Session Classification: Cosmic Physics