Telescopes Drive Themselves: Optimizing Cosmic Survey Scheduling with Reinforcement Learning

Franco Terranova[12], Maggie Voetberg[2], Brian Nord[2345], Eric Neilsen[2]

New Perspectives 2023

June 2023

1. University of Pisa, Department of Information Engineering, Pisa, IT **2.** Fermi National Accelerator Laboratory, Scientific Computing Division, Batavia, IL **3.** Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL **4.** Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL **5.** Laboratory for Nuclear Science, MIT, Cambridge, MA

FERMILAB-SLIDES-23-143-CSAID

Clockwork Science Design

Use modern Machine Learning techniques to optimize what we investigate and how we do it

- Let the telescope decide what is interesting
  - Investigate that in non-human scheduled downtime
  - Respond to unexpected events without intervention
- Find novel questions to ask

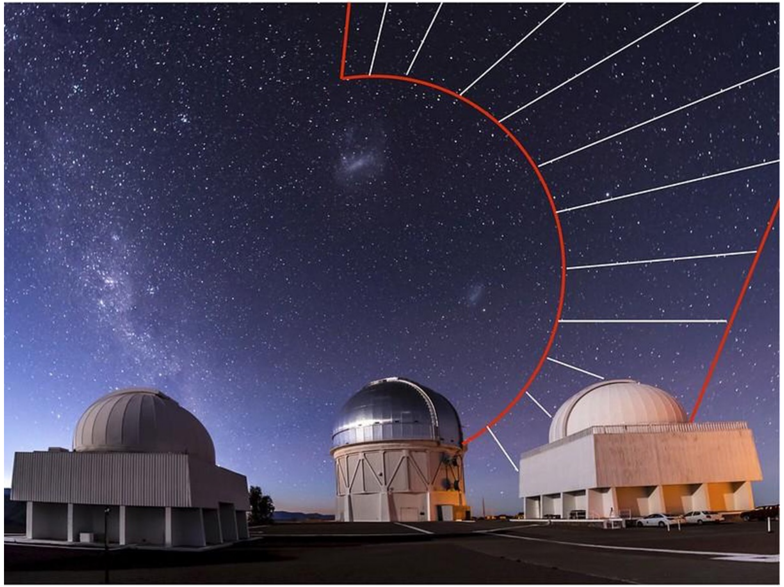

Dark Energy Survey, via Yuanyuan Zhang

FERMILAB-SLIDES-23-143-CSAID

The first step – Learning to schedule

- Start with a program that optimizes a sequential schedule based on a single parameter maximization
  - Give the model a set of parameters, receive the best possible next step in the schedule
- Working with Stone Edge Observatory in California for a proof-of-concept volume survey

Stone Edge Observatory

FERMILAB-SLIDES-23-143-CSAID

Learning to Learn – Reinforcement Learning

- Attempts to replicate how a human learns
- Conceptualize as a model trying to learn how to play a game
- The agent interacts with an environment, via an ‘action’, and gains ‘reward’ through the interaction

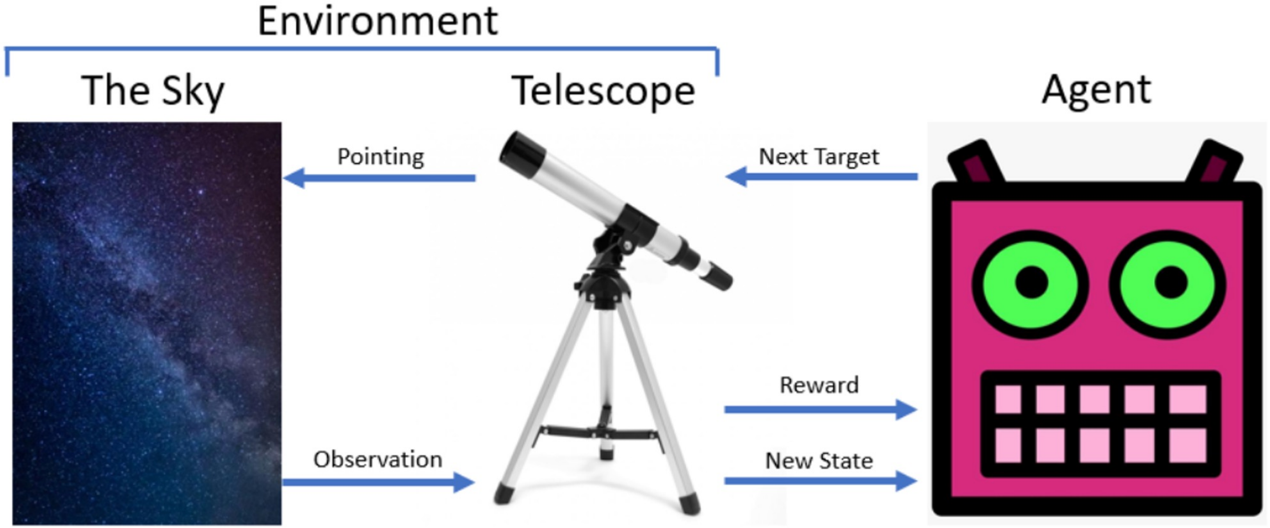

Using Value-Based Methods

- Imagine a look-up table of values that correspond to every state-action pairing, with the probability of the optimal reward
- Gets complicated very quickly for large problems – intractable for continuous problems

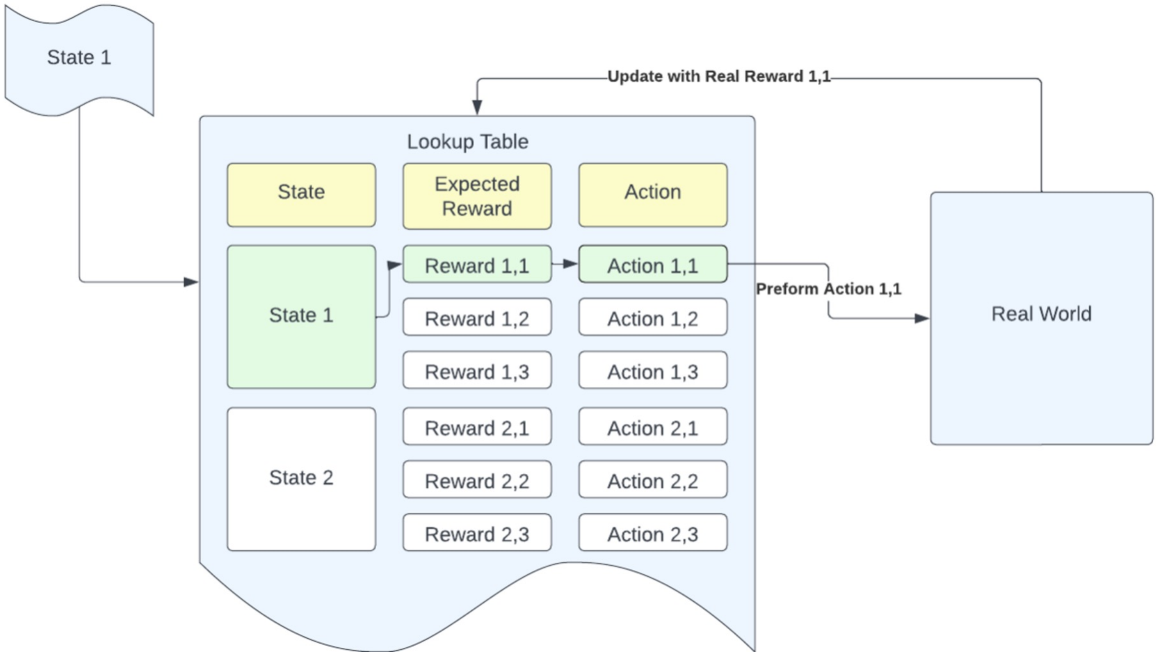

FERMILAB-SLIDES-23-143-CSAID

Enhance Value Methods with Offline Deep Learning

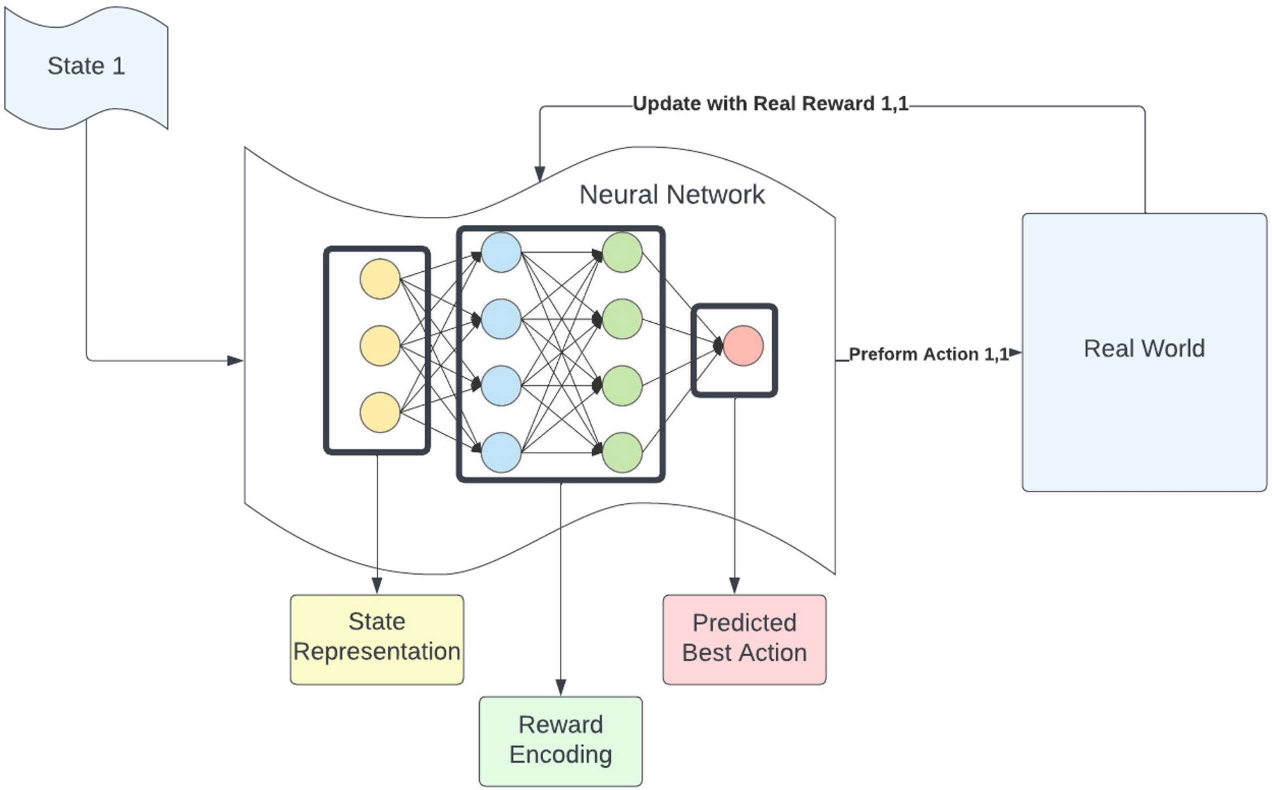

FERMILAB-SLIDES-23-143-CSAID

Data Preprocessing

ID	az	alt	….	sun_ ra	sun_ decl	….	moon_ airmass	moon_ angle	sky_ mag
ID_1	10.01	31.03	….	64.77	21.41	….	5.89	112.2	18.91
ID_2	10.01	31.03	….	64.77	21.41	….	5.89	112.2	18.91
….	….	….	….	….	….	….	….	….	….

Observation space

Right Ascension	Declination
79.0	324.0
24.0	221.0
….	….

Action space

T-Effective
38
29
….

Reward

New State
ID_29
ID_46
….

Next state

Normalization of the dataset involves adjusting the value of the data to a standardized range. This operation will ensure that the data is in a format easier for the RL algorithm to work with.

- Z-score normalization of the reward values intra-state
- Z-score normalization of the observation space columns

Data point Mean

$$z = \frac{(x - \mu)}{\sigma}$$

Standard deviation

Training

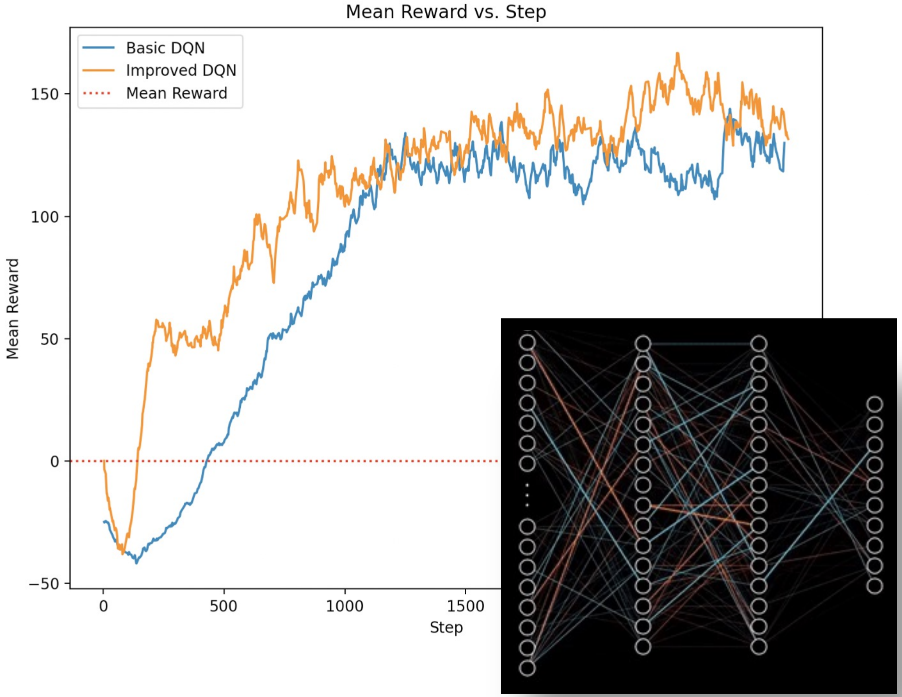

Improvements:

- Normalization of the observation space
- N-steps Bellman Unrolling (N=3)
- Dueling DQN
- Noisy Linear Layers

FERMILAB-SLIDES-23-143-CSAID

Effectiveness & Generalization - Comparison with a Random Agent

Our Agent

Random Agent

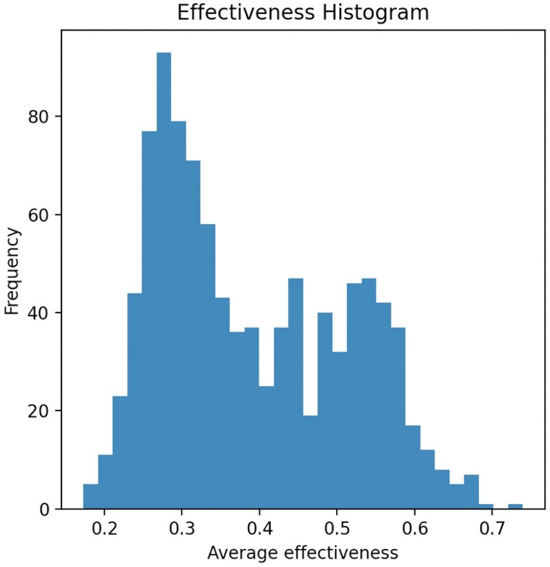

Mean: 0.39
Std: 0.12

Full environment

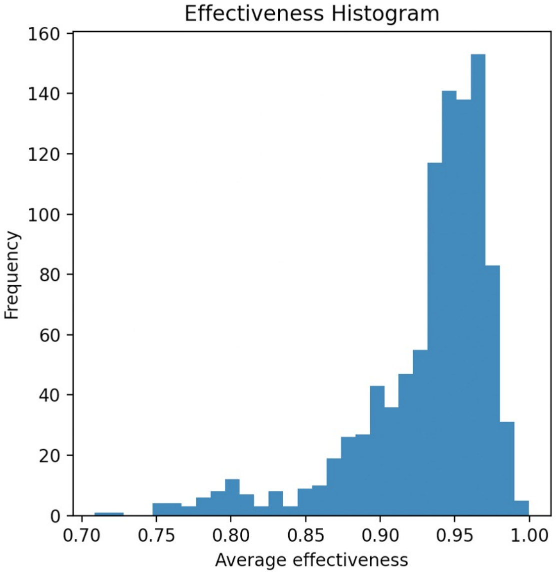

Mean: 0.93
Std: 0.05

Test set

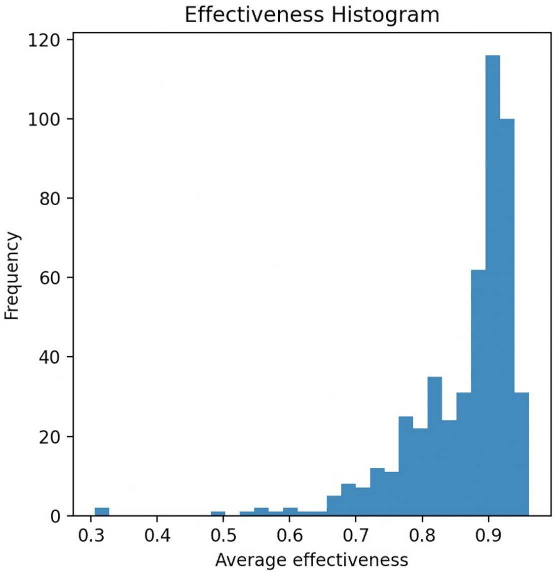

Mean: 0.86
Std: 0.08

FERMILAB-SLIDES-23-143-CSAID

Live execution - Current best Agent

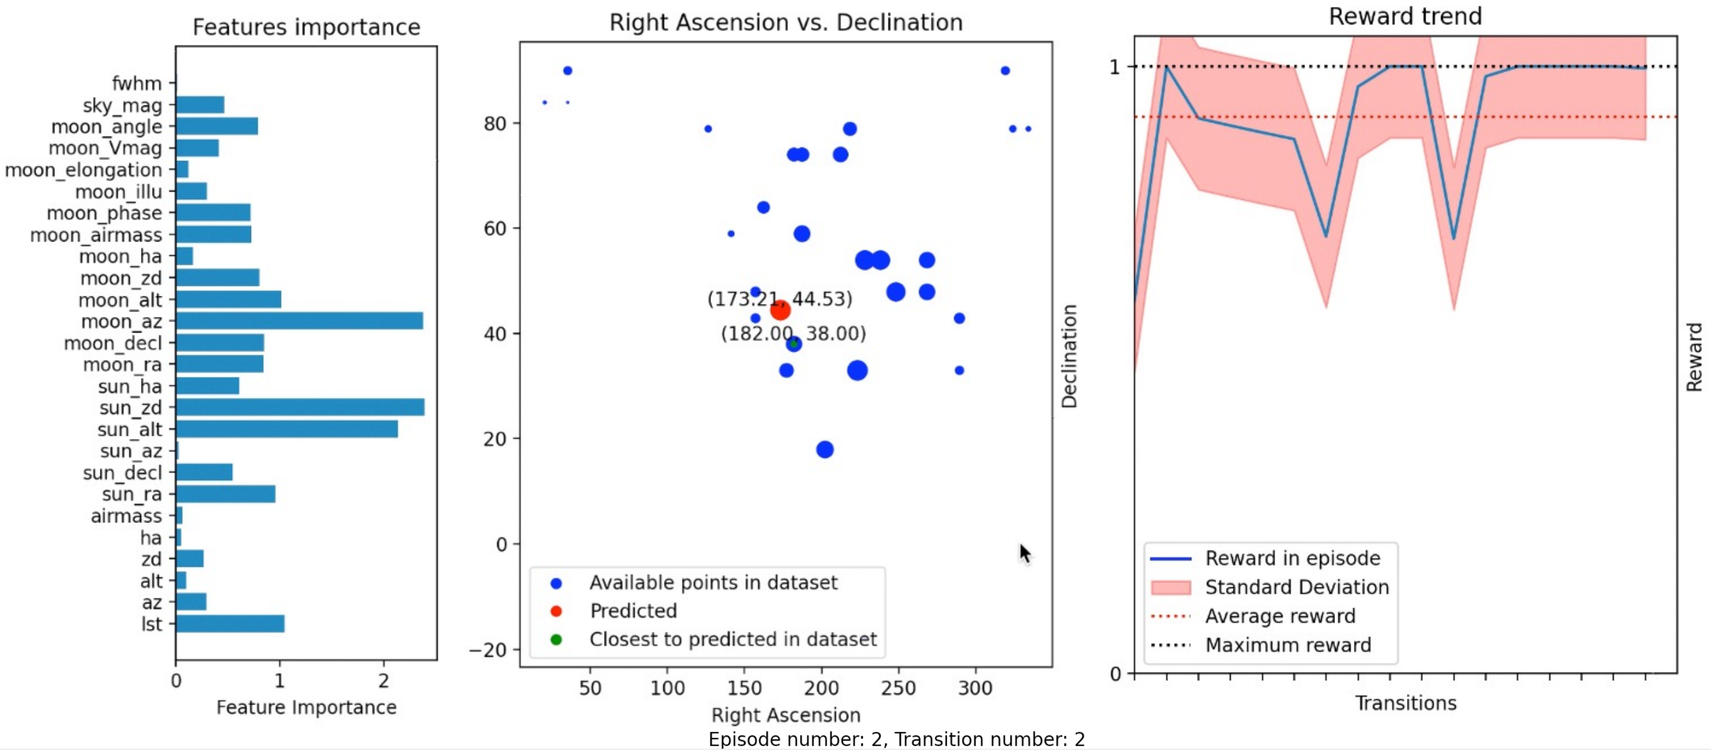

Fermilab

FERMILAB-SLIDES-23-143-CSAID

Conclusions

- Trained an agent able to optimize telescope observations
- Future deployment at the Stone Edge Observatory

"What if an autonomous self-driving telescope could find the solutions we've been missing?"

The Victor M. Blanco 4-meter Telescope at Cerro Totolo Inter-American Observatory

FERMILAB-SLIDES-23-143-CSAID