Multi-Agent Radiation Localization

Student: Teresa Nguyen (they/them)
Mentors: Dr. Christof Teuscher (he/him) and Philippe Proctor (he/him)

Motivation/Project Goals

- Build a multi-agent model that localizes radiation sources in an urban environment.
- Each agent has a radiation detector and can communicate with its nearest neighbors only.
- Building off of pre-existing work by Philippe Proctor on single agent radiation detection.

Fig 1: (top) DARPA radiation identification image of an urban environment, (right) portable radiation detector, (bottom) drone radiation detector

How is my approach different/better? What is the secret sauce?

- I hypothesize that multi-agent radiation localization could lead to faster and more efficient detection in urban environments
- Secret Sauce
  - OpenAI (spinning up)
  - Proximal Policy Optimization (PPO)
  - Parallelization through Message Passing Interface (MPI)
    - Decreases training time
    - Increases diversity of environments agent is exposed to
  - Use Philippe's PPO plus multi-agent optimization

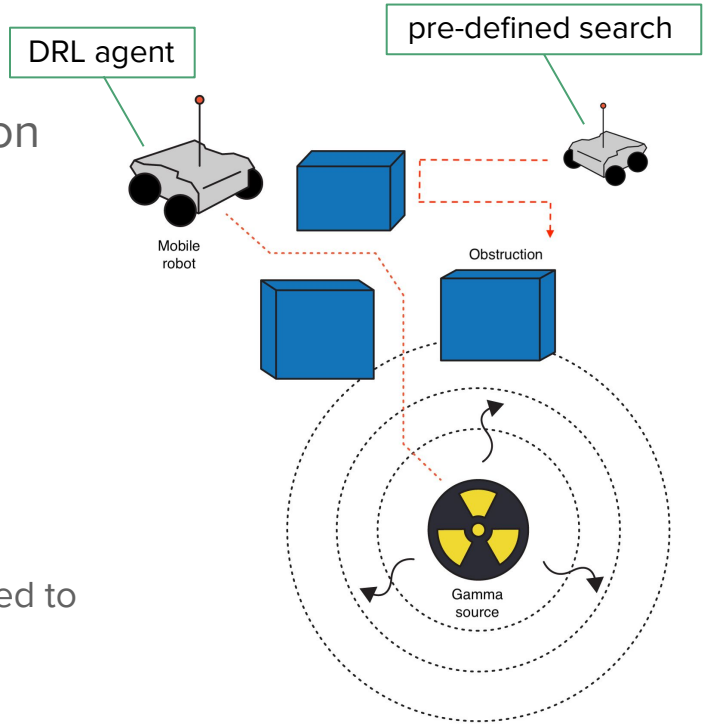

Fig. 2: autonomous mobile robots, Proctor 2021

How have others tried to solve this problem?

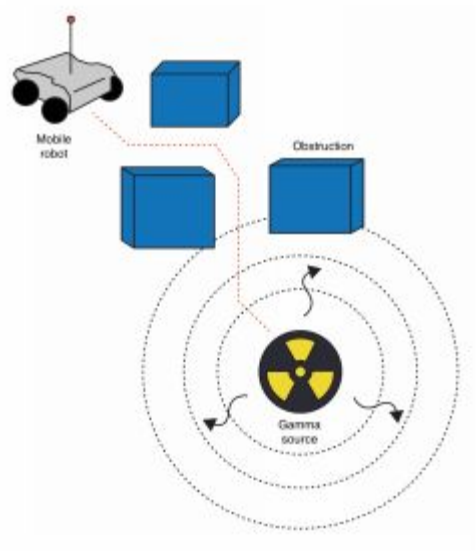

Fig 3: single agent radiation detection in a non-convex (with obstacles) environment (Proctor 2021)

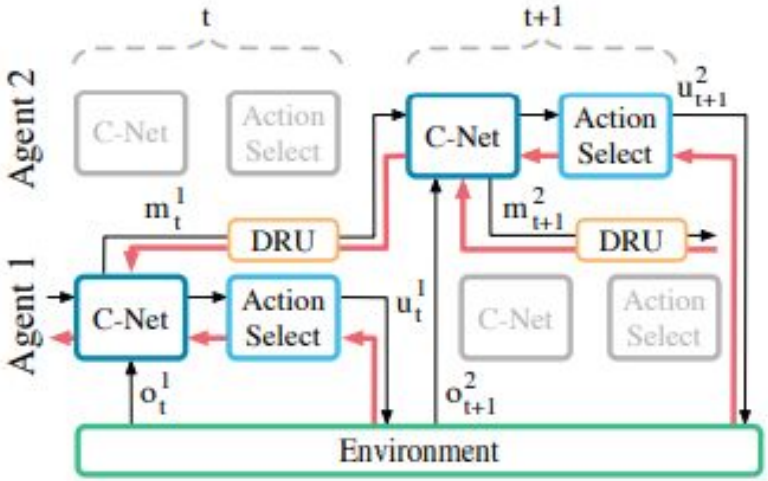

Fig 4: Differentiable Inter-Agent Radiation Localization (DIAL) diagram, where two agents operate independently with a shared objective (Foerster et al 2016)

What is being used

On-Policy Reinforcement Learning

Policy: An agent’s strategy to accomplish the task at hand

Reinforcement Learning: the agent's goal is to learn a policy that maximizes a reward signal in an environment

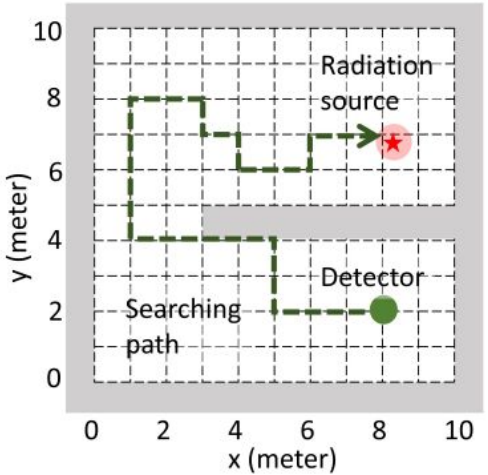

Fig 5. Illustration of radiation detection task. Liu & Abbaszadeh 2019 pg. 3

What is being used

Reinforcement Learning is: a subsection of machine learning that uses a policy (strategy) to maximize a cumulative reward system in an environment[1]

It is appropriate for this problem because:
The detector needs to find the most efficient path to the radiation source. The closer the detector is to its goal, the higher the reward

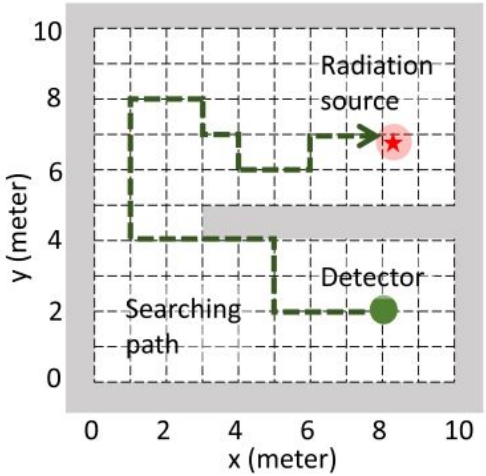

Fig 5. Illustration of radiation detection task. Liu & Abbaszadeh 2019 pg. 3

[1] Philippe Proctor, Christof Teuscher, Adam Hecht, and Marek Osinski. “Proximal Policy Optimization for Radiation Source Search.” *Sensors*, July 19, 2021, under review.

What have I accomplished so far?

- Making OpenAI compatible with Windows 10 (this took a while...)
- Deep research on machine learning and reinforcement learning

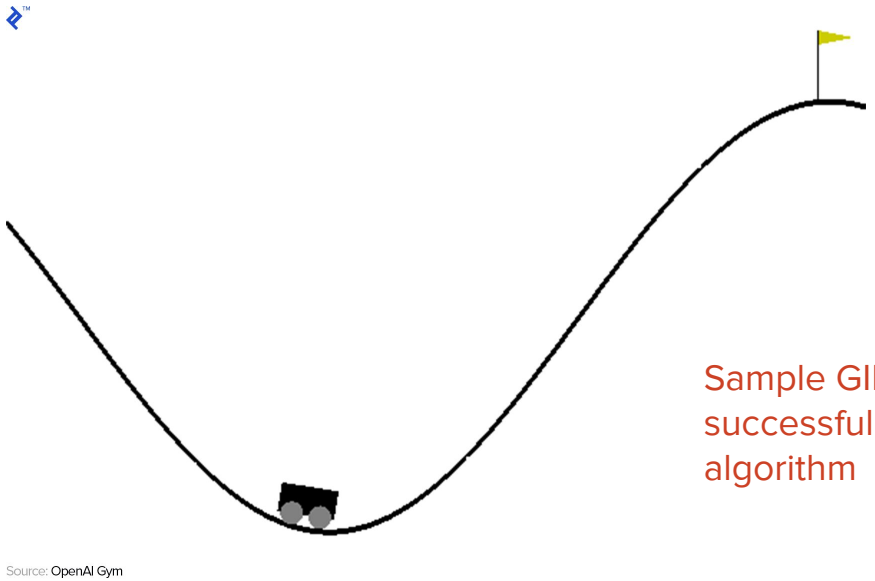

Sample GIF of a successful algorithm

Fig 6: OpenAI mountain car GIF. Goal is to teach agent to “learn” how to use force and momentum to push its way up the hill. OpenAI 2015

What have I learned from this?

- Machine learning has a wide variety of applications
- Formulate a clear plan of execution for intended goal
- Nuance within machine learning can make a significant difference in how a project is approached and completed

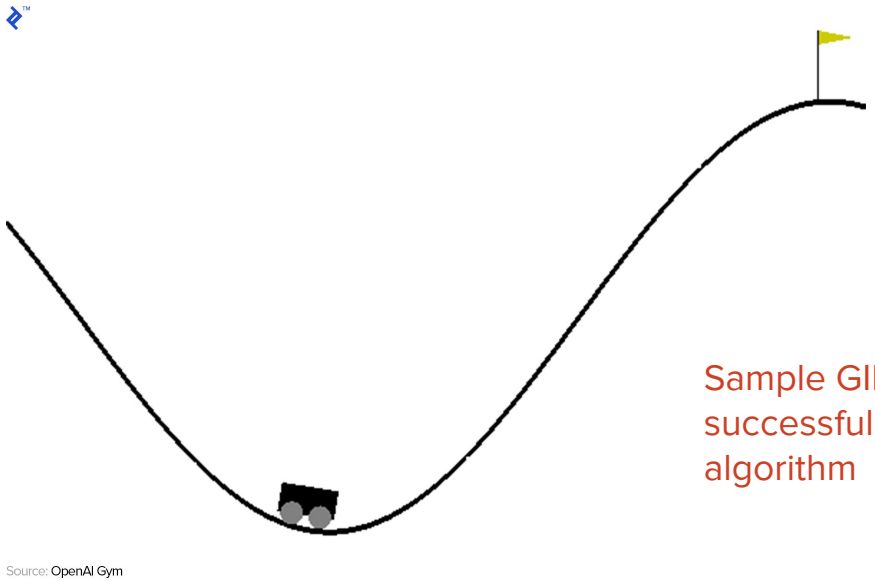

Sample GIF of a successful algorithm

Fig 6: OpenAI mountain car GIF. Goal is to teach agent to “learn” how to use force and momentum to push its way up the hill. OpenAI 2015

What will the algorithm do?

- Train agents to find radiation sources quickly and efficiently through Reinforcement Learning and Proximal Policy Optimization
- Allow agents to handle different real-world environments because of its training to move through random virtual environments

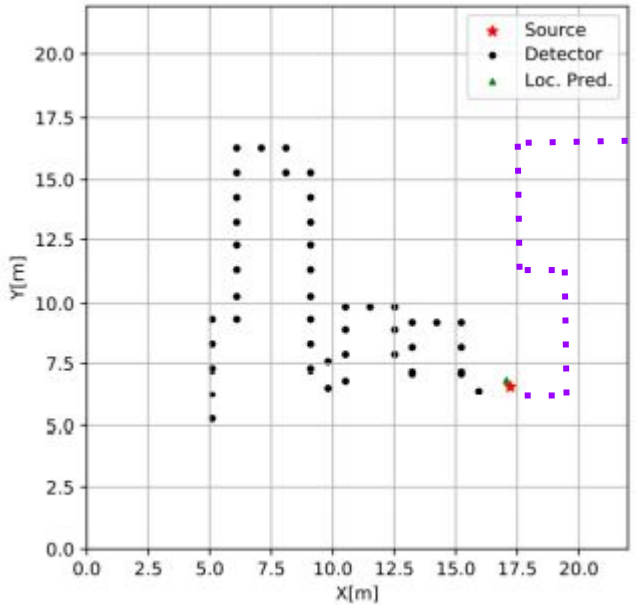

(a) Detector path.

Fig. 7: Edited graphical representation of a radiation detector choosing a path to get to its radiation destination. Proctor 2021

What SHOULD have happened?

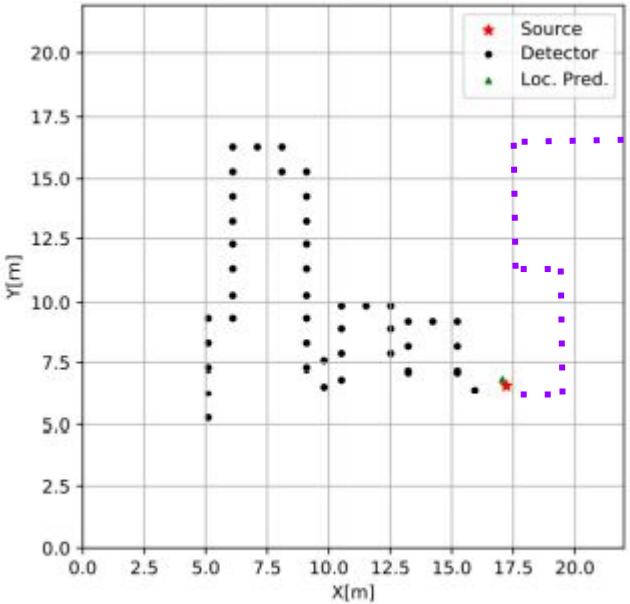

(a) Detector path.

Objective: Find radiation source 

Episode 1: (Fig 7) Agents travel through environment using PPO to locate radiation source

Episode 2 (3, 4, etc): Agents share findings from their first episode, refine their path search to radiation source

Complete when: Agents consistently use most efficient route based on reward system designed

Fig. 7: Edited graphical representation of a radiation detector choosing a path to get to its radiation destination. Proctor 2021

What are the next steps?

- Continue studying Philippe's work on single agent radiation detection
- Figure out how to get at least two agents running successfully for radiation detection
- Explore DIAL approach

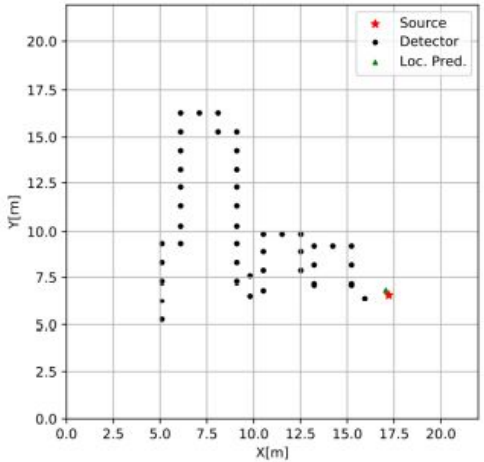

(a) Detector path.

Fig. 8: Graphical representation of a radiation detector choosing a path to get to its radiation destination. Proctor 2021

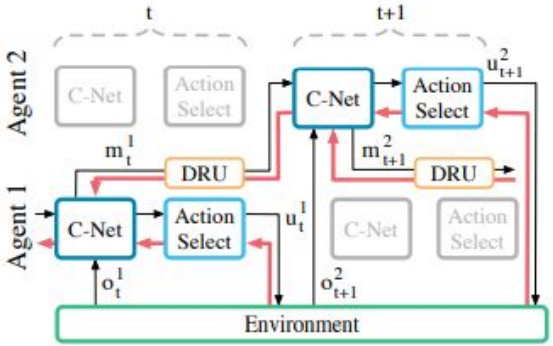

Fig. 9: Differentiable Inter-Agent Radiation Localization (DIAL) diagram, where two agents operate independently with a shared objective (Foerster et al 2016)

What are the next steps?

- Figure out how agents will be communicating with each other
  - What info will be exchanged?
  - What are their positions in relation to suspected source?
- Can DIAL and PPO work together?

Fig. 10: (bottom) DARPA radiation identification image of an urban environment, (right) radiation drone detector

What difference will your project make? For whom?

- Allows for homeland security to be able to safely and reliably detect radiation
  - Airports
  - Building security
  - Hazardous materials
- Benefits humanity by locating radiation sources without ACTUALLY putting humans at risk (using robots/drones, RIID, etc)

Fig. 11: RadSeeker holding Radiation Isotrope Identification Device (RIID)

Thank You!

Acknowledgements:

The REU Site is supported by the National Science Foundation under grant no 1758006. The research was supported by the Defense Threat Reduction Agency (DTRA) under grant no HDTRA1-18-1-0009.

I would like to thank Dr. Christof Teuscher, Philippe Proctor, Adrian Jimenez, and the Defense Threat Reduction Agency (DTRA) for your wisdom, patience, and guidance through such complex material.