Tiny Treasure Hunter

Teaching the VEX AIM Robot how to find objects

Author: Santiago Arambulo Patino (andrew ID: sarambul)
Course: Cognitive Robotics 15-694 @ CMU

Problem: The VEX Aim robot does not know how to search for things, even when controlled by an LLM

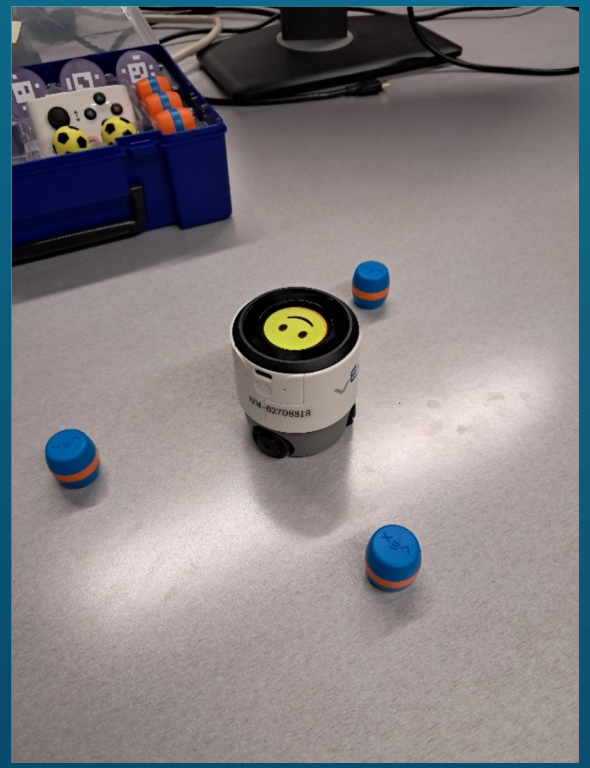

LLM Prompt: How many barrels are around you? Look in every direction to find the barrels. They might be behind you.

Outcome: The robot only identified the barrel in front of it. It did not look around before answering

Author: Santiago Arambulo Patino | Andrew ID: sarambul

Approach: Use Reinforcement Learning to develop a search strategy

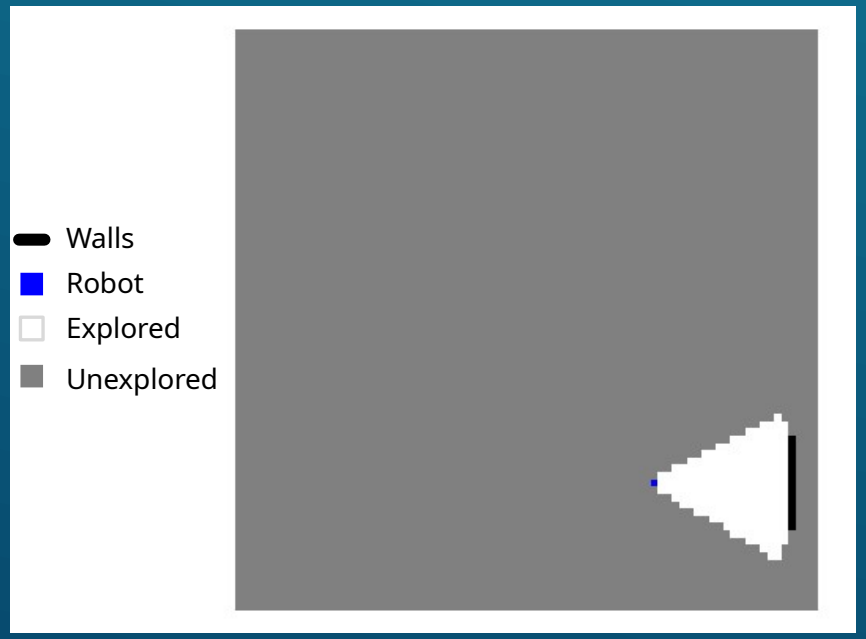

Training environment

- Observation space: 2D Grid World encoding robot and object positions
- Action Space: Turn and move up, left, right, down
- Rewards: Number of grid cells revealed at each step (based on robot's field of view) OR penalty for invalid actions (bumping walls)
- Random initial position and orientation

Author: Santiago Arambulo Patino | Andrew ID: sarambul

Interesting aspects of solution: Learned action distribution is almost random

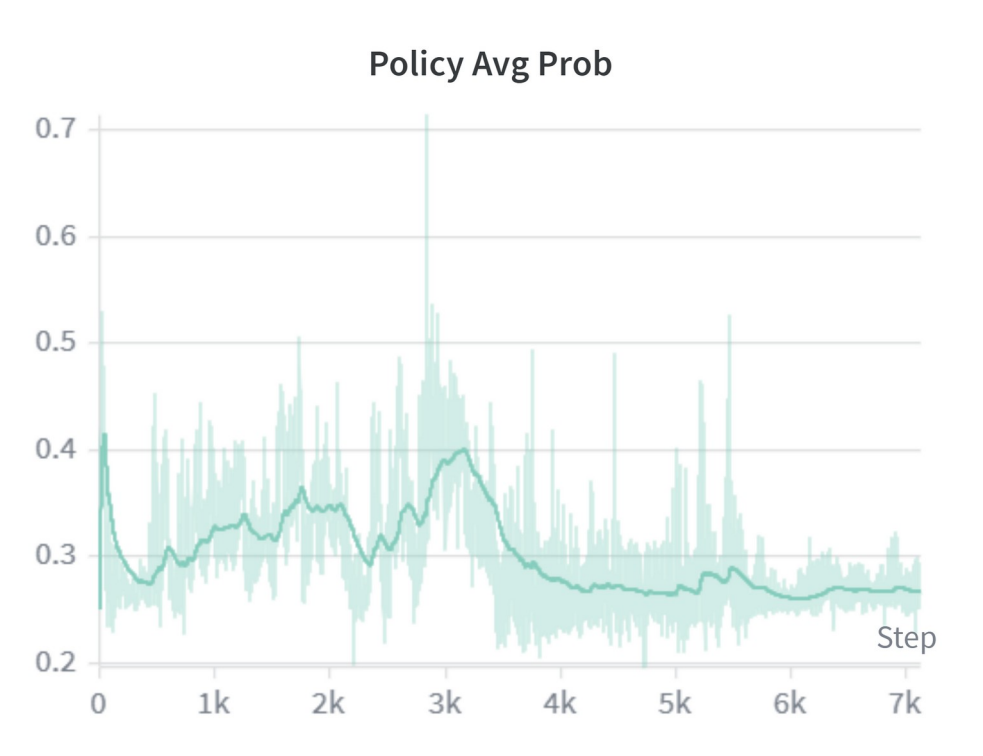

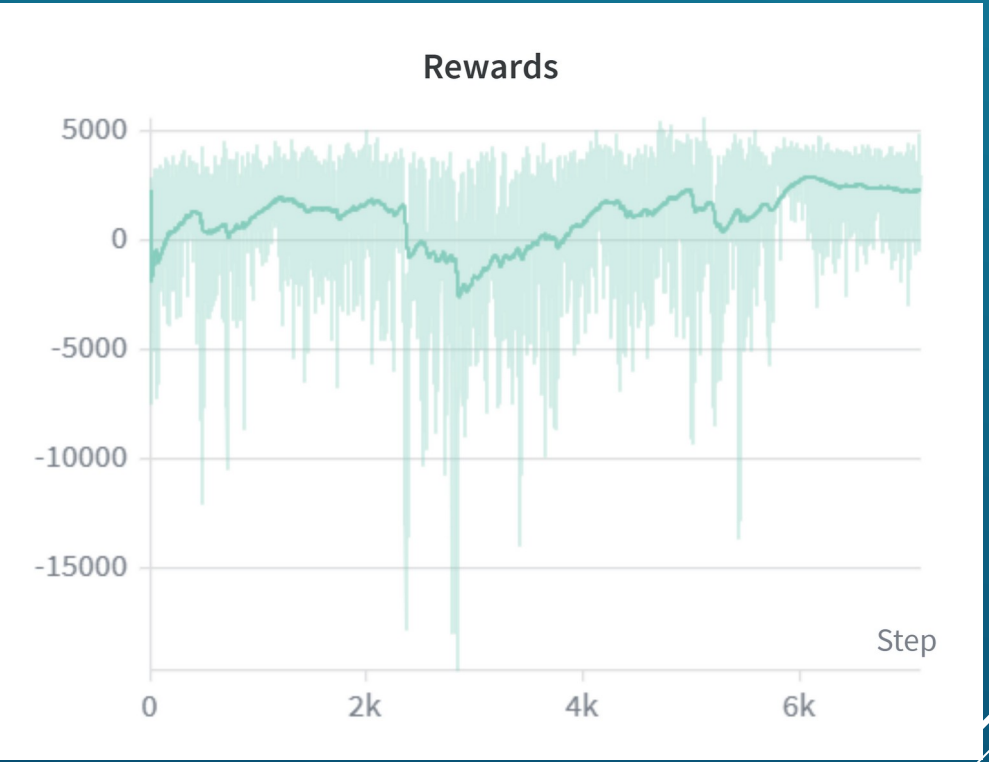

Author: Santiago Arambulo Patino | Andrew ID: sarambul

Demo

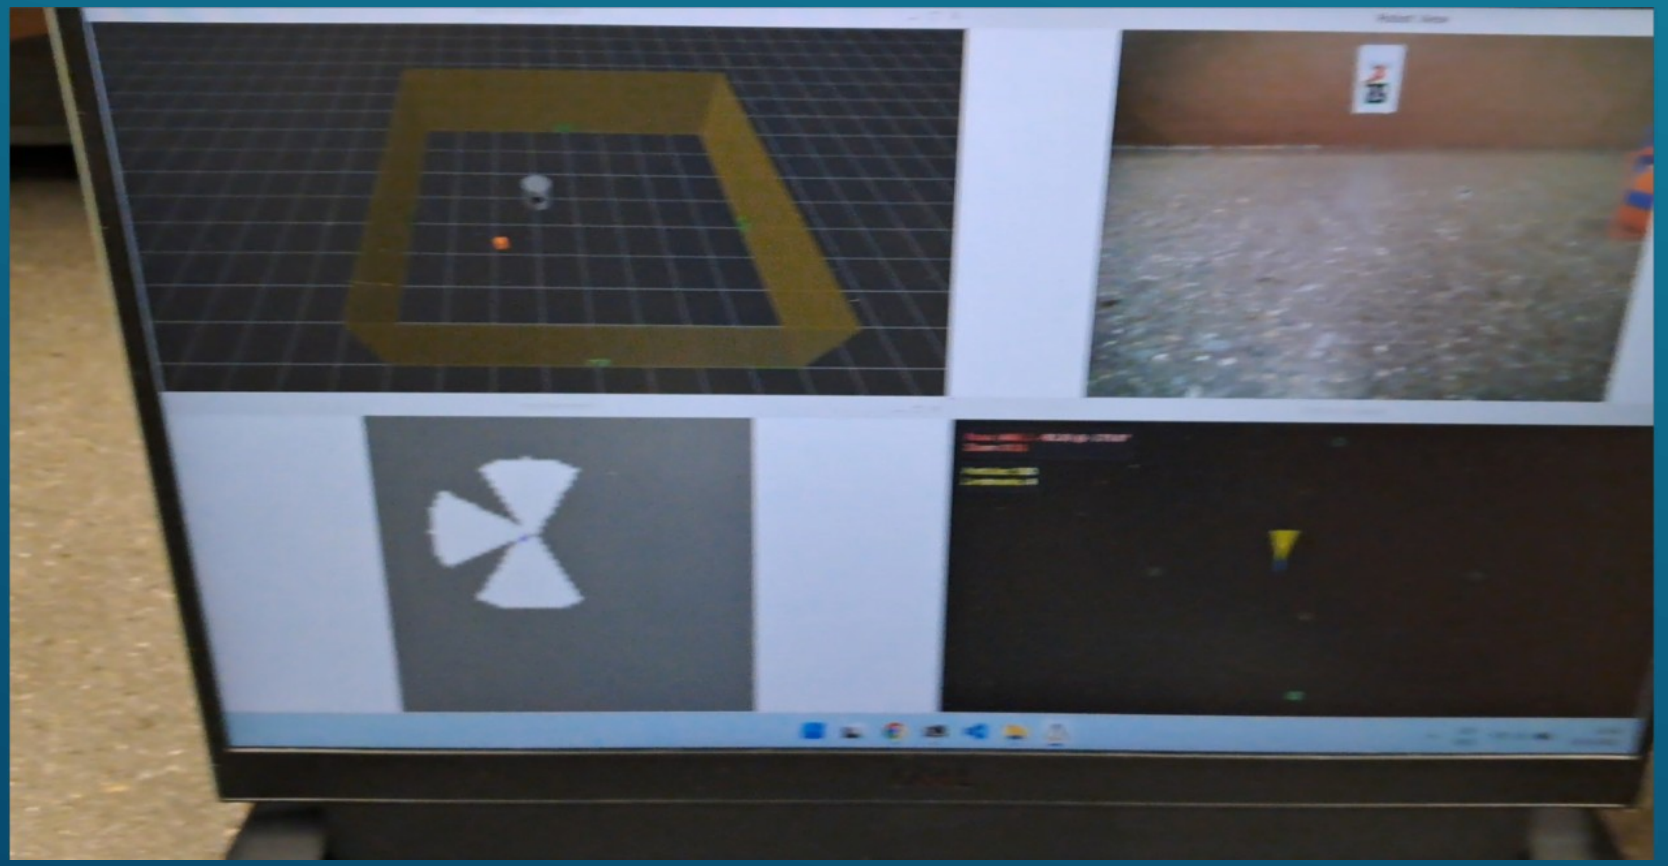

Link to video

Author: Santiago Arambulo Patino | Andrew ID: sarambul

Results

- The robot finds barrels in the empty room, but needs to perform many actions to do so
- The policy does not avoid walls as much as desired

Relevant things outside the project's scope:

- The robot needs to get close to a barrel to identify it
- Adding a barrel to the world maps required around 1-2 seconds (delays the search)

Author: Santiago Arambulo Patino | Andrew ID: sarambul

Future work: Occlusion, Smaller turns, and SLAM

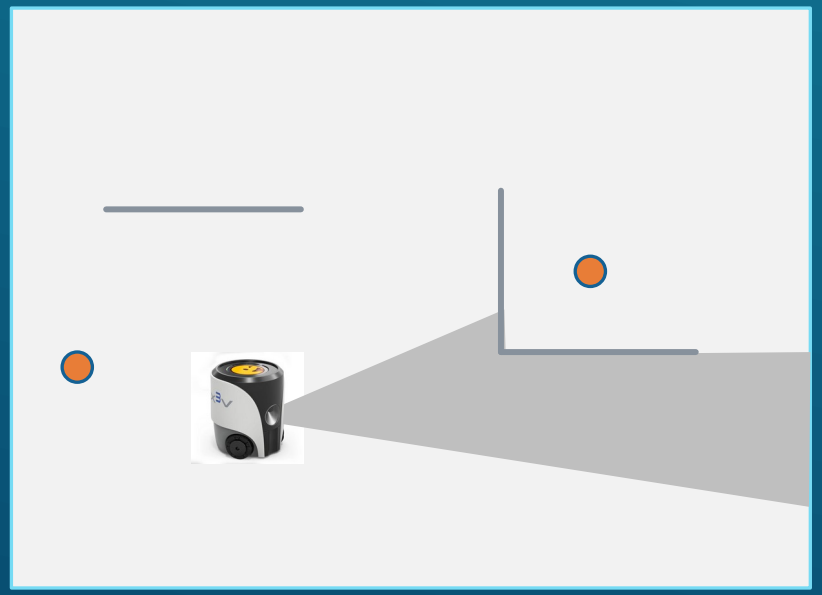

- The current environment has built-in support for rooms with occlusion, but the agent was not trained on such rooms
- The robot's actions could include turning and moving diagonally to explore the rooms more easily
- The world map is predefined and has landmarks, but we could use SLAM to build the world map.

Author: Santiago Arambulo Patino | Andrew ID: sarambul