Reinforcement Learning from a GA population

Daniel Kudenko: 2018 Summer Project 1

Reinforcement Learning

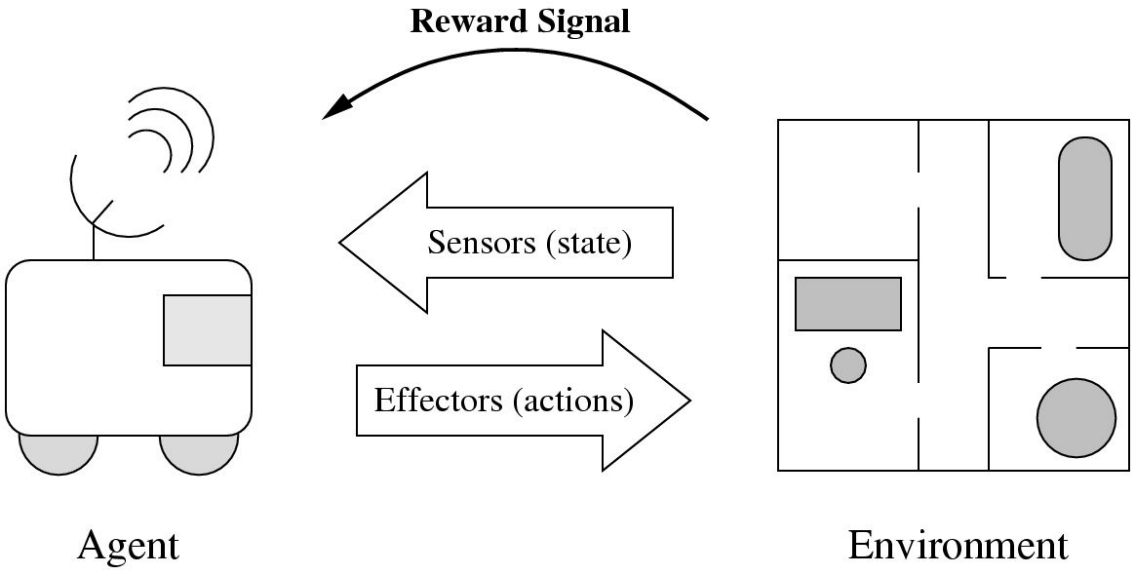

- Problem: Does not scale well to complex domains.
- Solution: learning from expert demonstrations.
- But: There may only be a group of sub-optimal experts. Where to get these experts?

Learning from Multiple Demonstrations

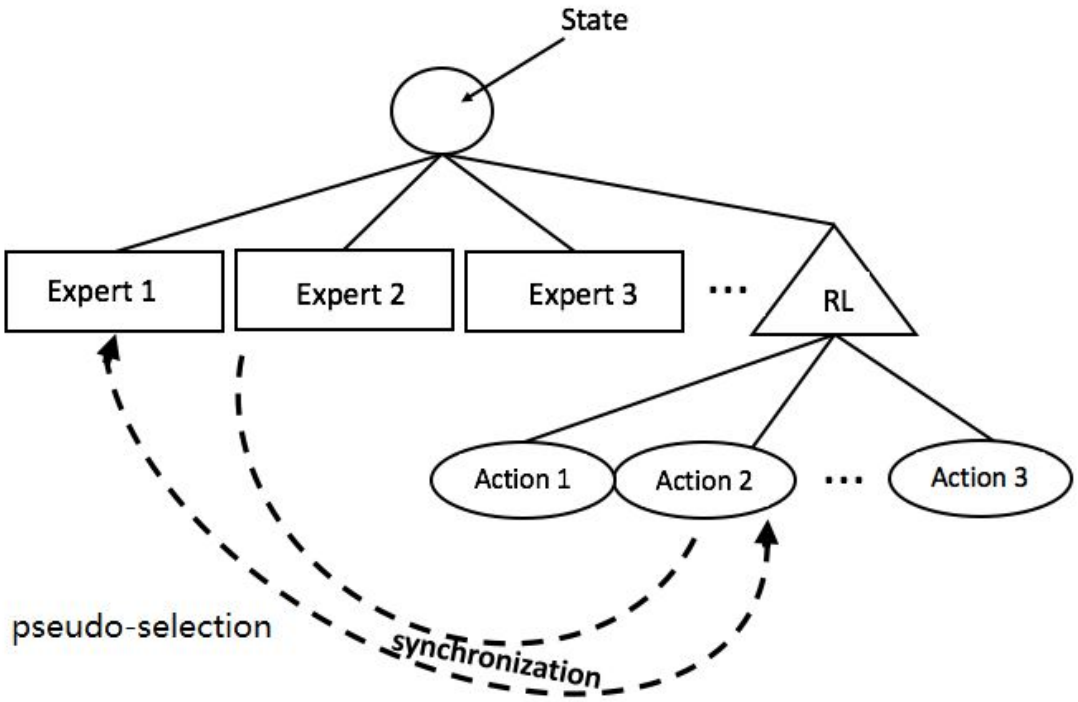

- Two-level Q-learning [Li & Kudenko 18]
- Learn trust value for each expert in each state.

Genetic Algorithms

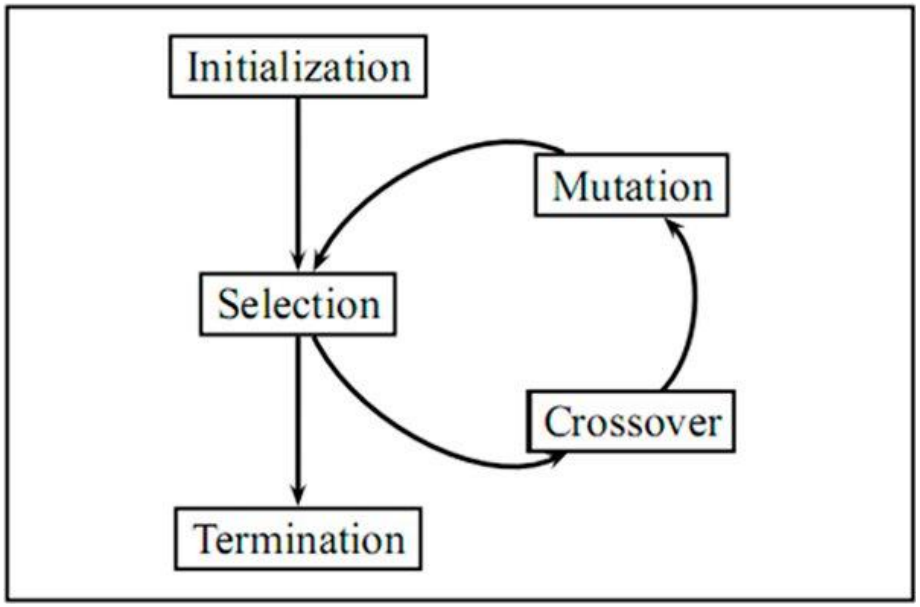

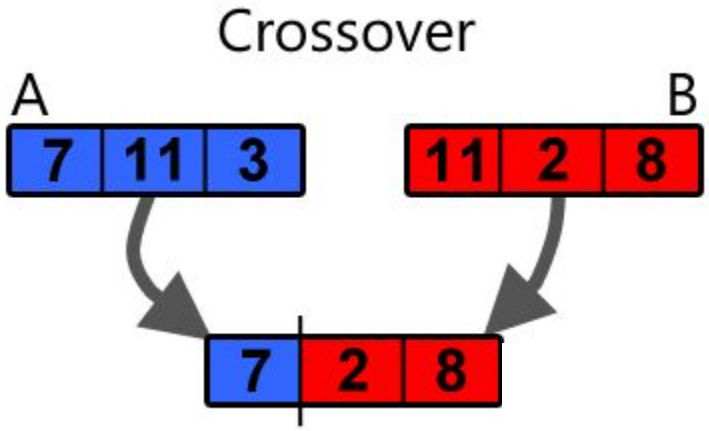

- Modelled on evolution.
- Keeping track of a population of AI agents.

Project Idea

- Apply Genetic Algorithms to a game, e.g. Asteroids.
- Use the GA population to provide demonstrations.
- Feed these into TLQL.
- Compare performance to:
  - RL agent without expert demonstrations.
  - RL agent with hand-coded experts.

Desired Student Skills

- Confident programmer.
- Enthusiastic about blue-sky research.