Beamline Tuning with Reinforcement Learning

David Wang, Paul Jung, Olivier Shelbaya, Spencer Kiy, Wojtek Fedorko, Rick Baartman, Oliver Kester

OLIS Beamline

- Starting point for AI-tuning
  - Low current
  - Non-radioactive
- Manual tuning by operators takes many hours, taking away from research beam time
- Goal for reinforcement learning agent:
  - Optimize beam transmission
  - Offset unknown misalignments
  - Improve speed and accuracy of tuning

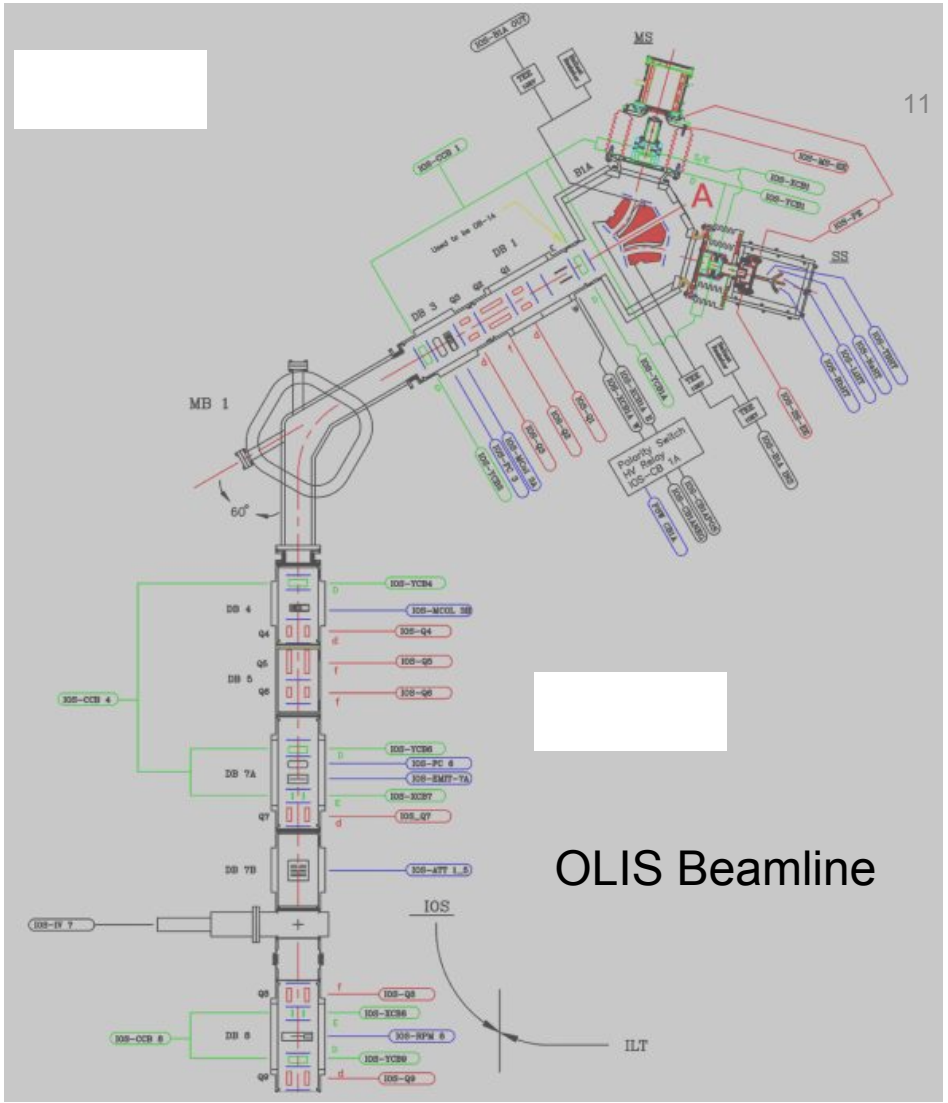

Source: Beam Physics Note TRI-BN-20-13R, Olivier Shelbaya

Source: https://towardsdatascience.com/getting-started-with-reinforcement-q-learning-77499b1766b6

Reinforcement learning

- Challenges of beamline environment:
  - Partially observed (only a few measurable spots)
  - Continuous and large action spaces
- Proposed Algorithm: Recurrent Deep Deterministic Policy Gradients (RDPG)
  - Actor-critic algorithm utilizing actor and critic networks to optimize agent learning
  - Long Short-term Memory (LSTM) networks to operate in partially observed environment

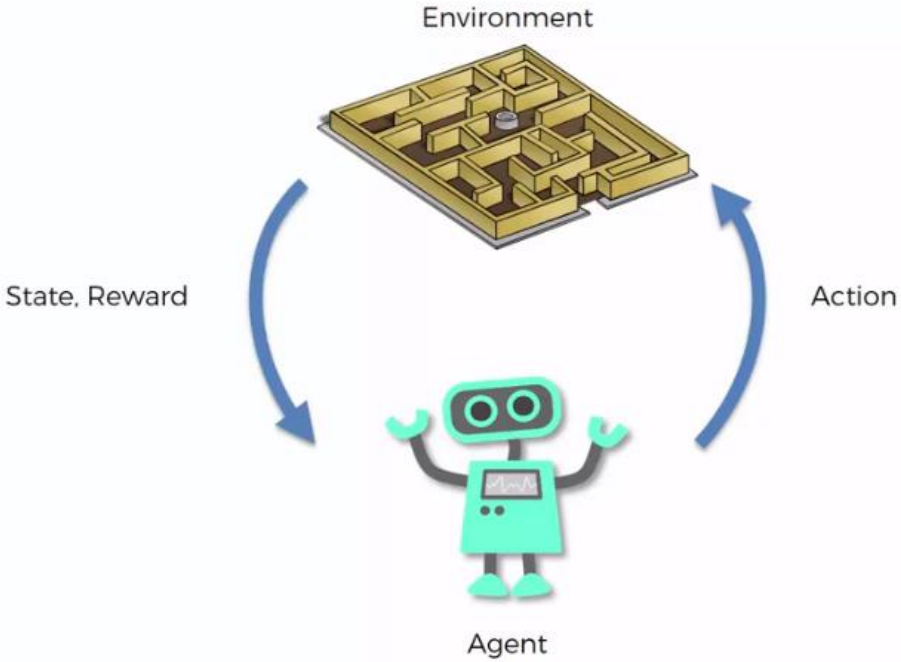

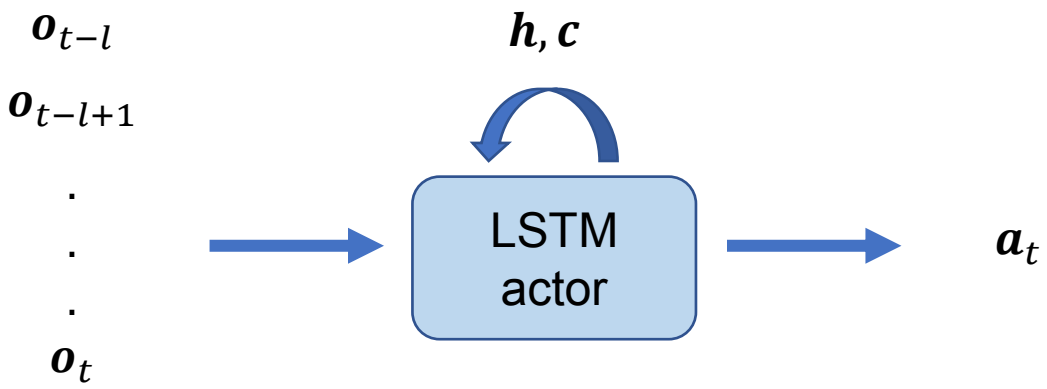

o_t : observation
for example, current measured at 2 faraday cups
a_t : predicted action
for example, angles to rotate each steerer
l : memory length of LSTM actor
h, c : hidden states of LSTM actor

Current Progress and Next Steps

- Beamline simulation
  - Approximate as a Gaussian particle distribution
  - Analyze in only 1 dimension
  - Use centroid (solid line) and envelope (dotted line) to determine transmission
- Current model trains well on simulation
  - Using realistic observations but artificial reward function
- Plans:
  - More realistic simulations of measurement and reward
  - Develop strategy and tools for real beamline tuning
  - Extend to ISAC and other beamlines