Leveraging Deep Reinforcement Learning and Braid Representations to Explore Knot Theory

BYU SRC Conference

Dylan Skinner - Brigham Young University - 24 February 2024

Reinforcement Learning

What is Reinforcement Learning?

Introduction

- Agent is placed in an environment
- Agent interacts with the environment through a set of actions
- Agent chooses its actions to maximize a reward (goal)

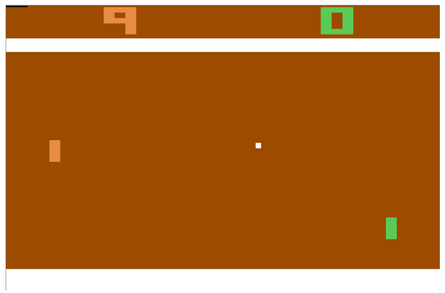

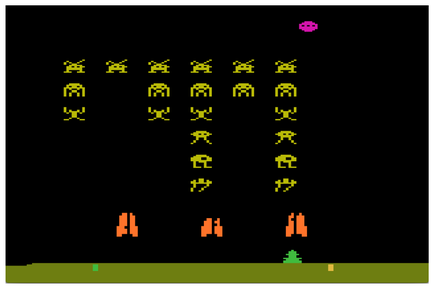

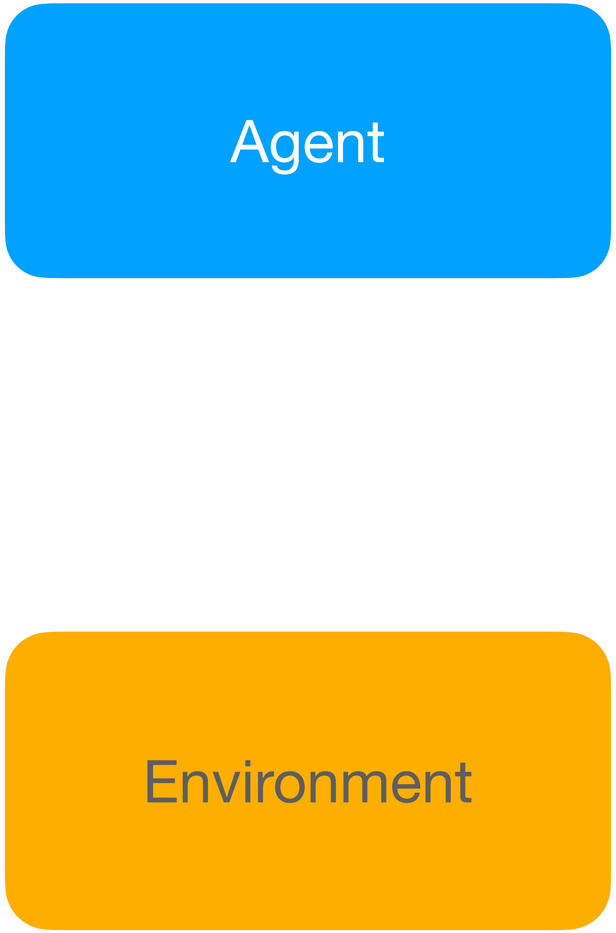
Agent
Environment

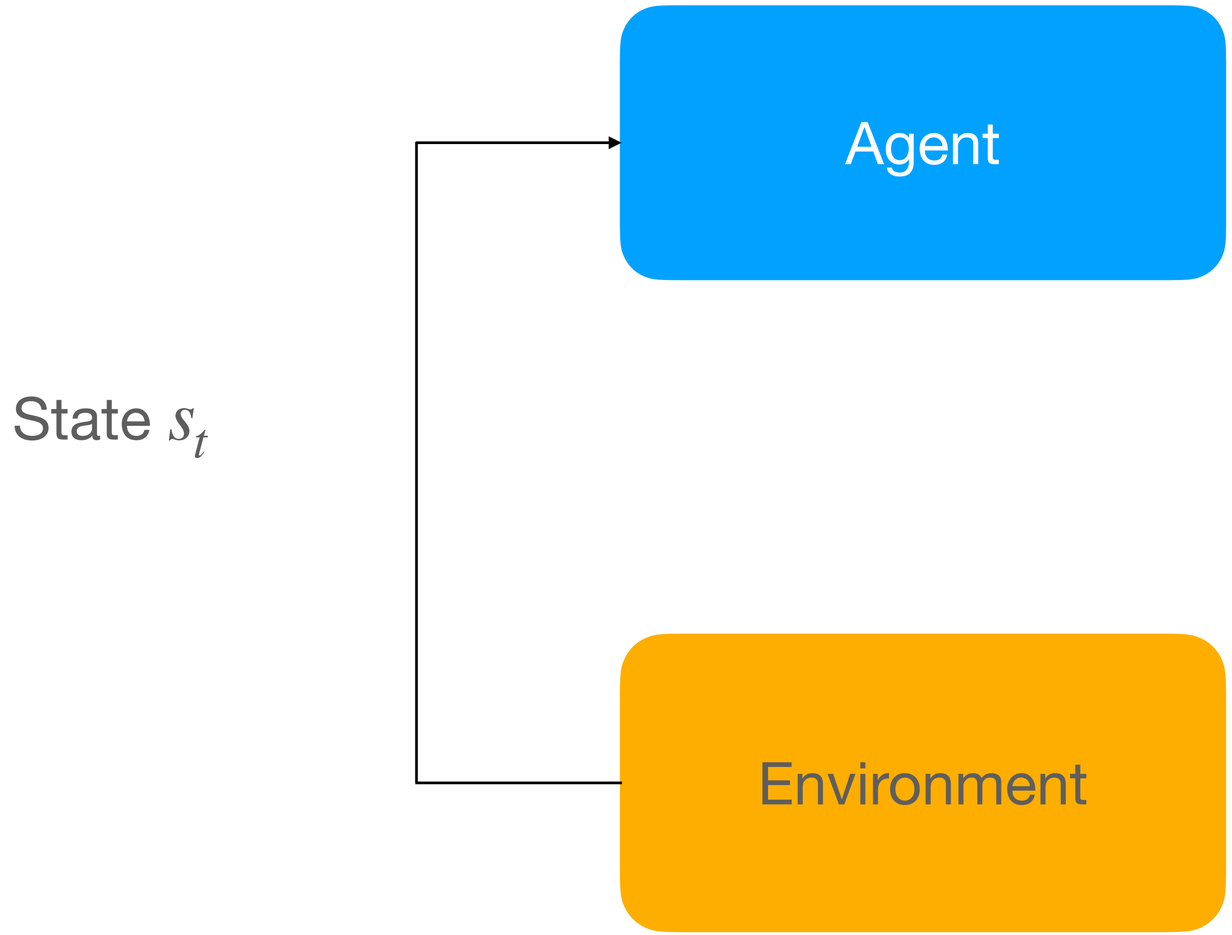
Agent
State s_t
Environment

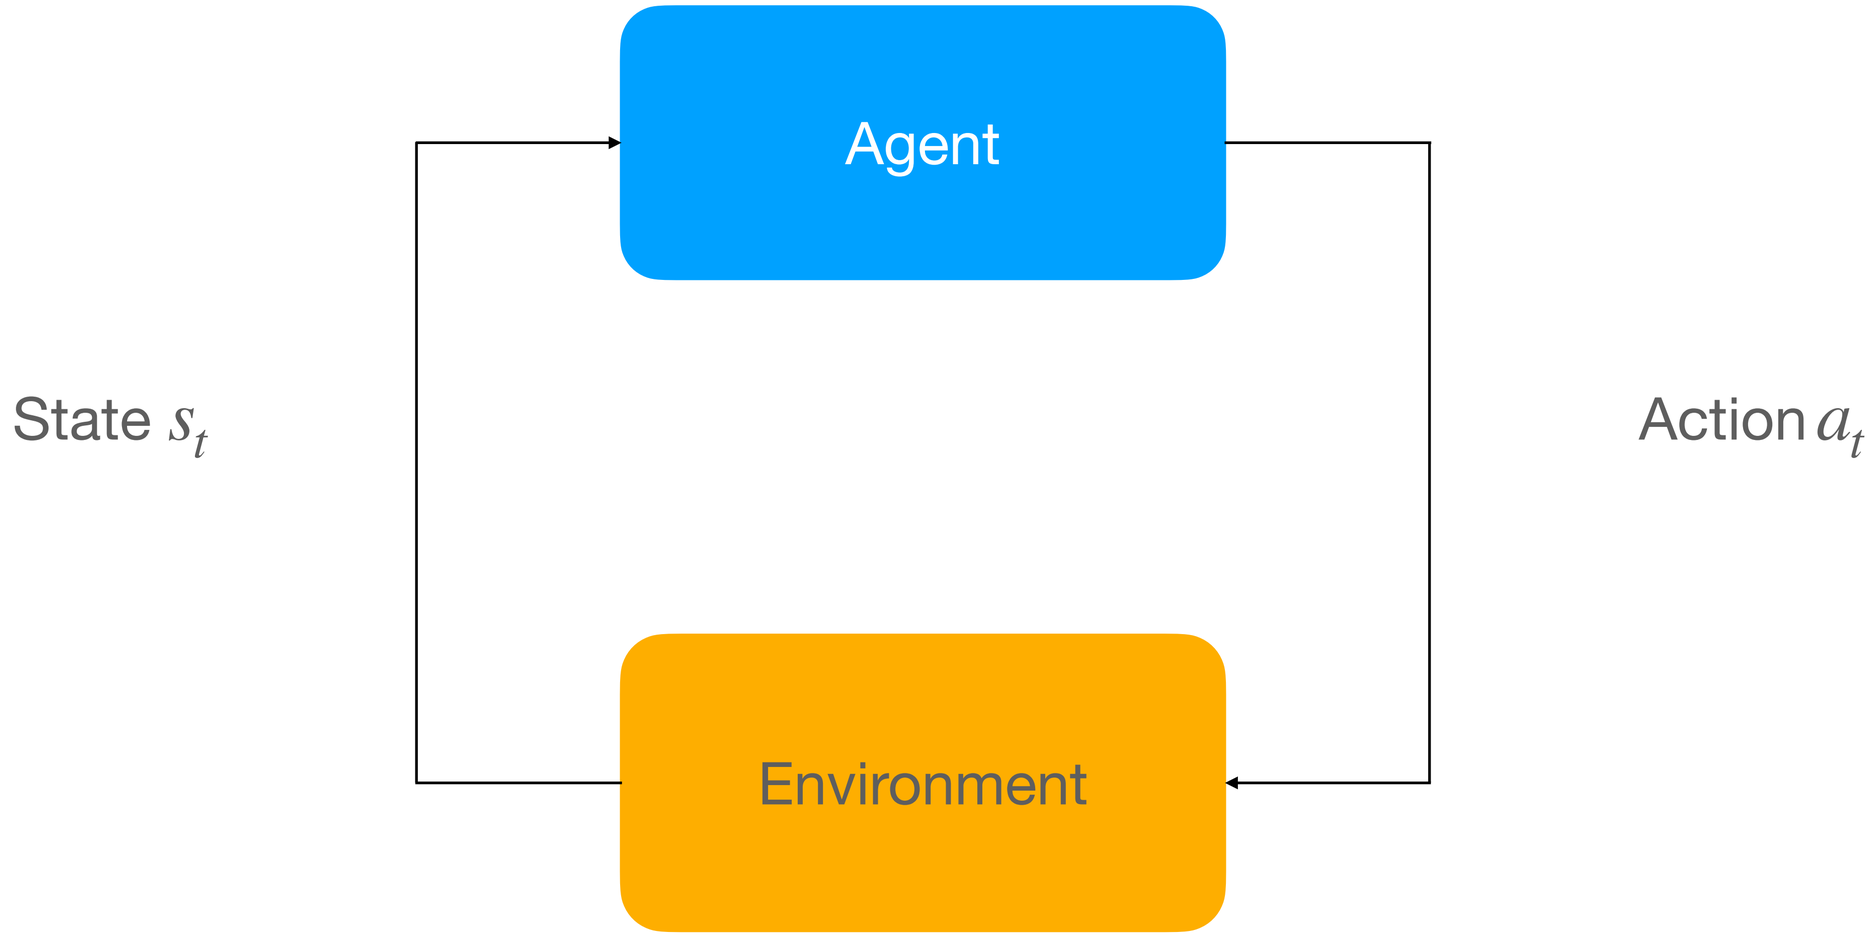
Agent
State s_t
Action a_t
Environment

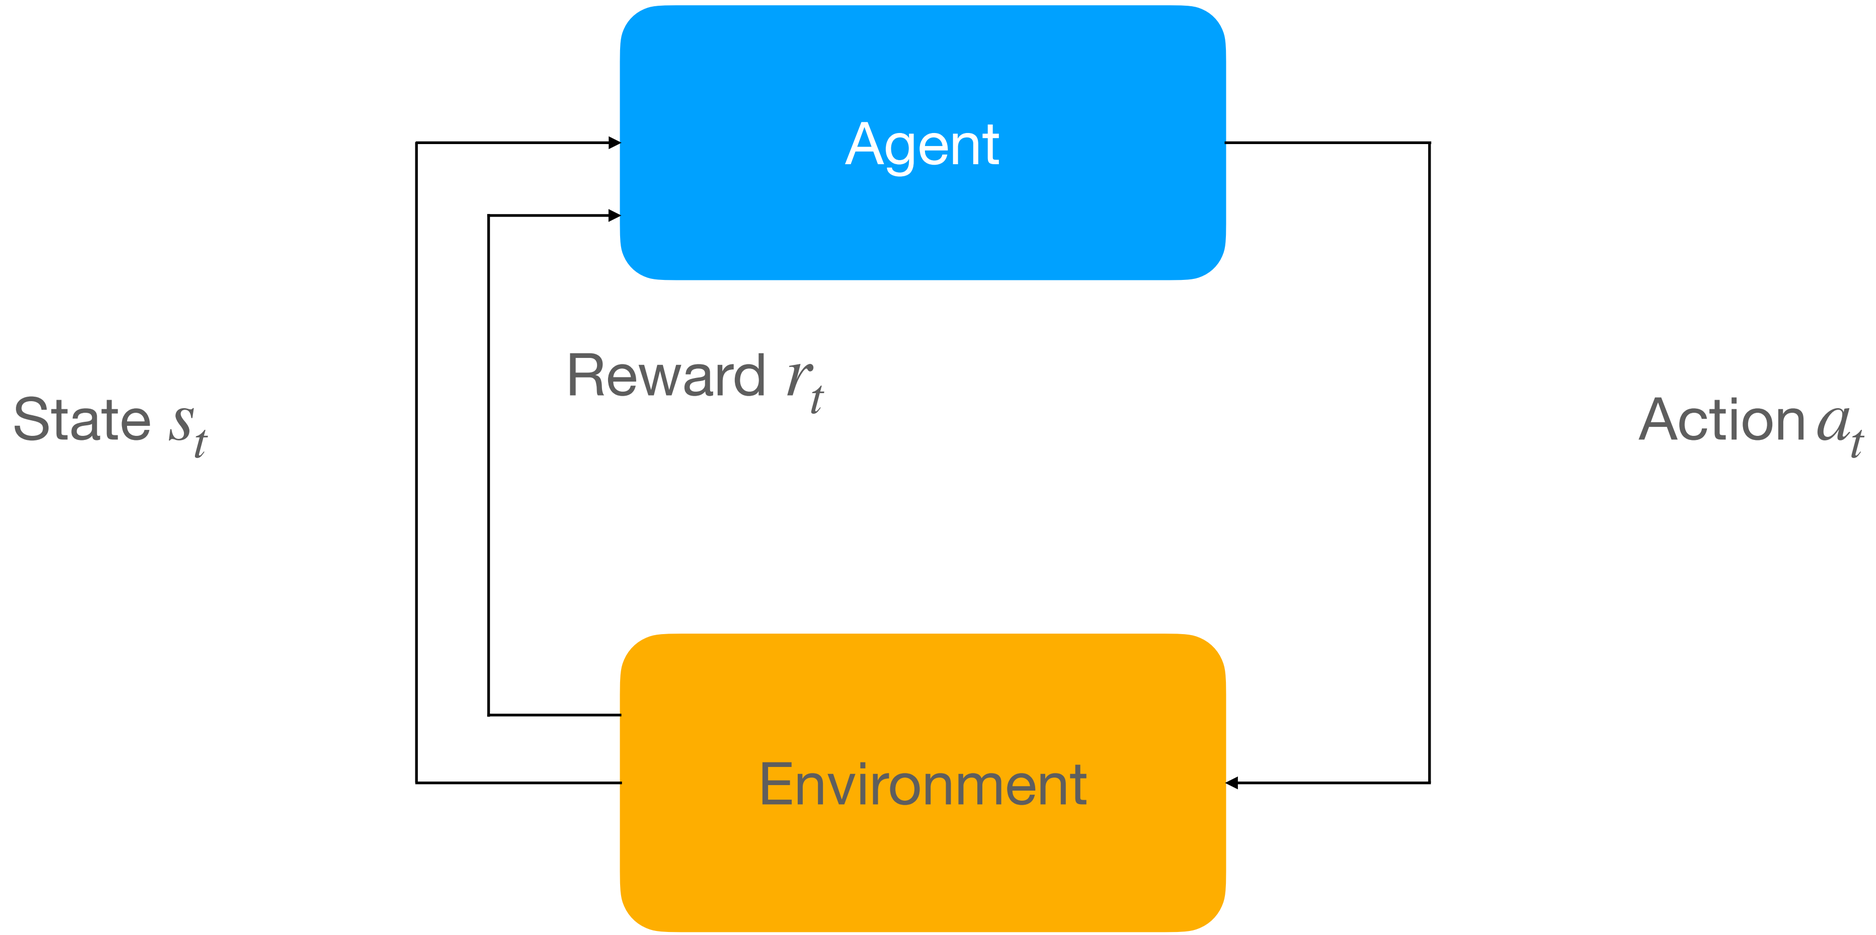
Agent
Reward r_t
State s_t
Action a_t
Environment

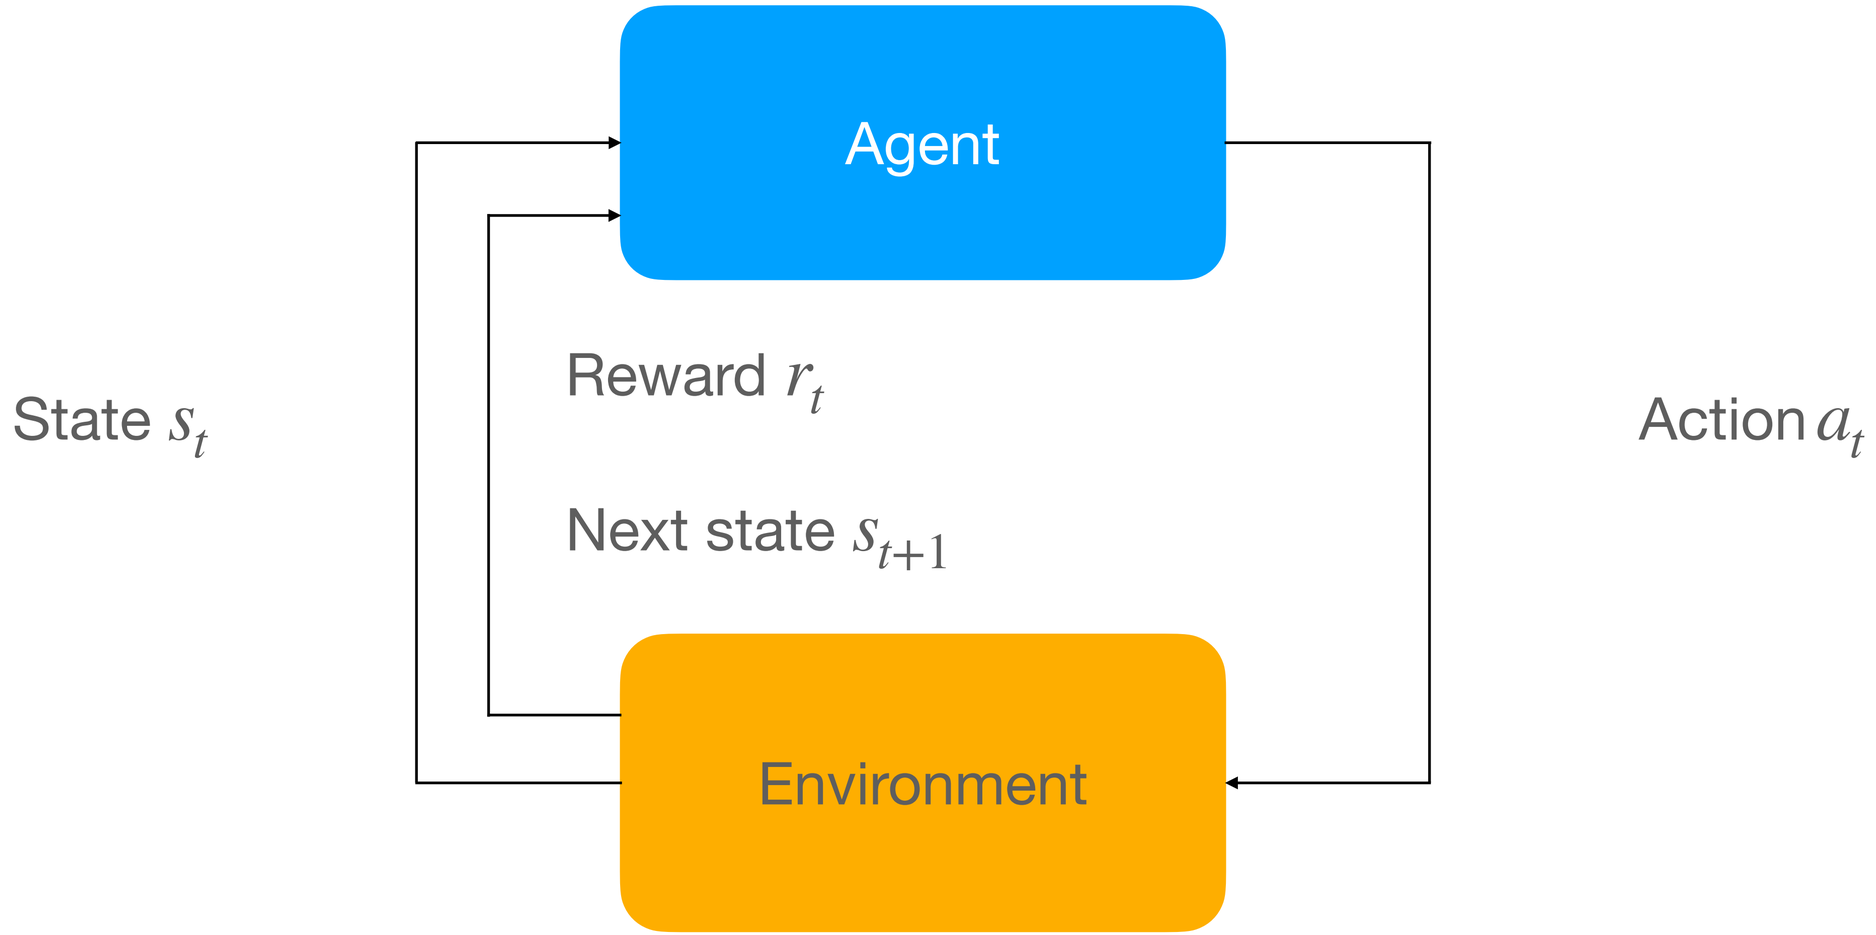
Agent
Reward r_t
State s_t
Action a_t
Next state s_{t+1}
Environment

Proximal Policy Optimization (PPO)

Introduction

- Algorithm developed by OpenAI in 2017
- Seeks a balance between ease of implementation, sample complexity, and ease of tuning
- Accomplished by computing update at each step to minimize cost function and deviate only slightly from current policy
- In order for this to work, the algorithm uses two separate policy networks

Proximal Policy Optimization (PPO)

Visualized

Knots, Seifert Surfaces, Slice Surfaces

Our Project

Knots

- Think of knots as a necklace whose clasps are together
- These knots live in dimension 3

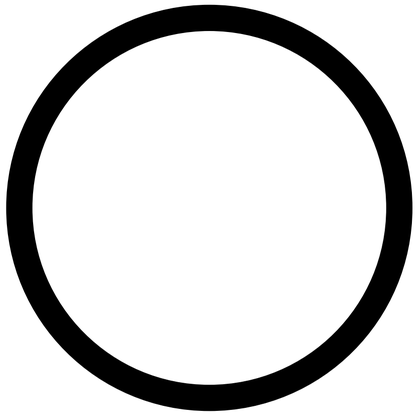

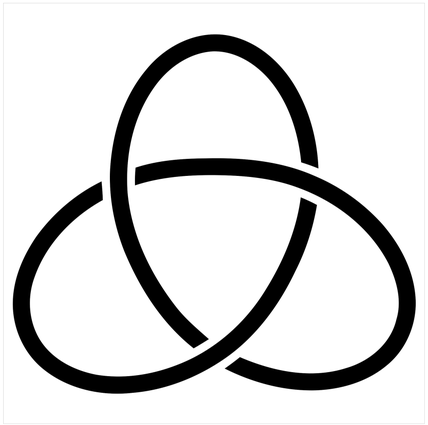

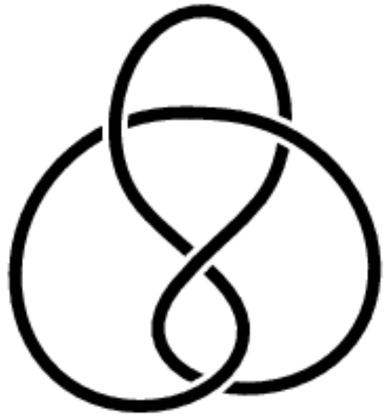

Our Project

Seifert Surfaces

- Seifert surfaces (3D) are orientable surfaces bounded by the knot

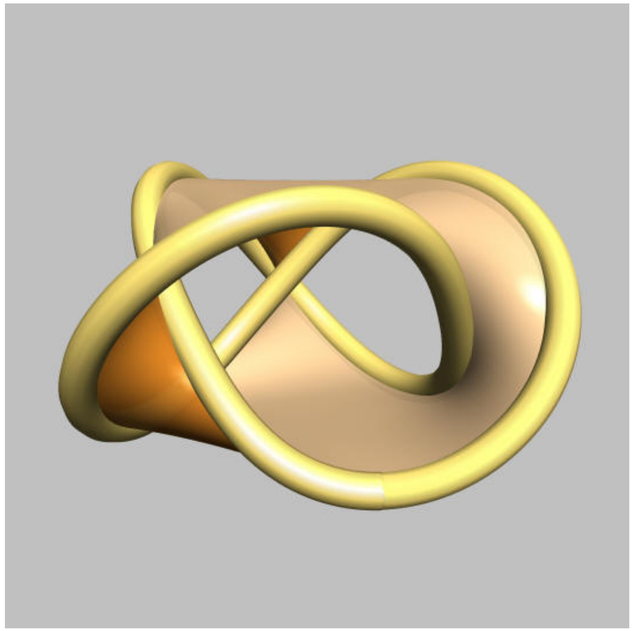

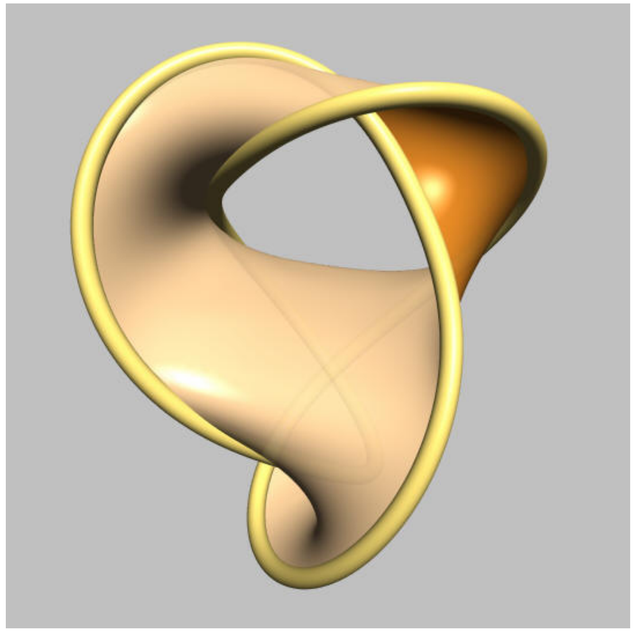

Our Project

Seifert Surfaces

- The unknot is the only knot that bounds a disk in S^3

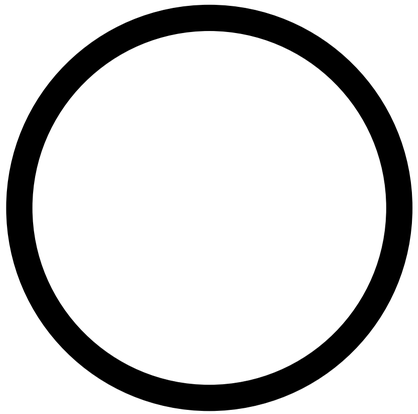

Our Project

Seifert Surfaces

- The unknot is the only knot that bounds a disk in S^3
- The trefoil and figure-eight knots both bound a punctured torus

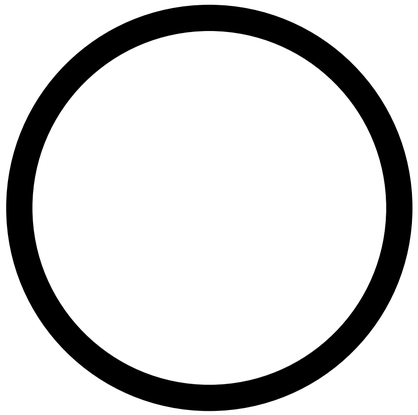

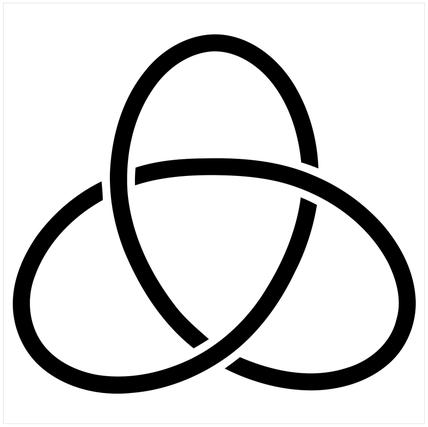

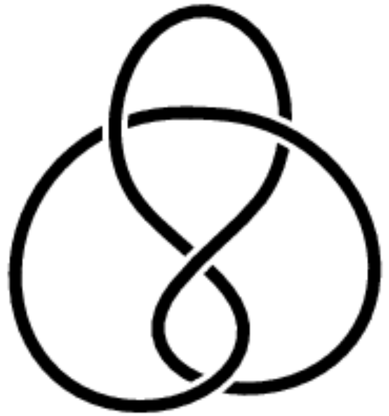

Our Project

Seifert Surfaces

- 6^1 knot also bounds a punctured torus in S^3

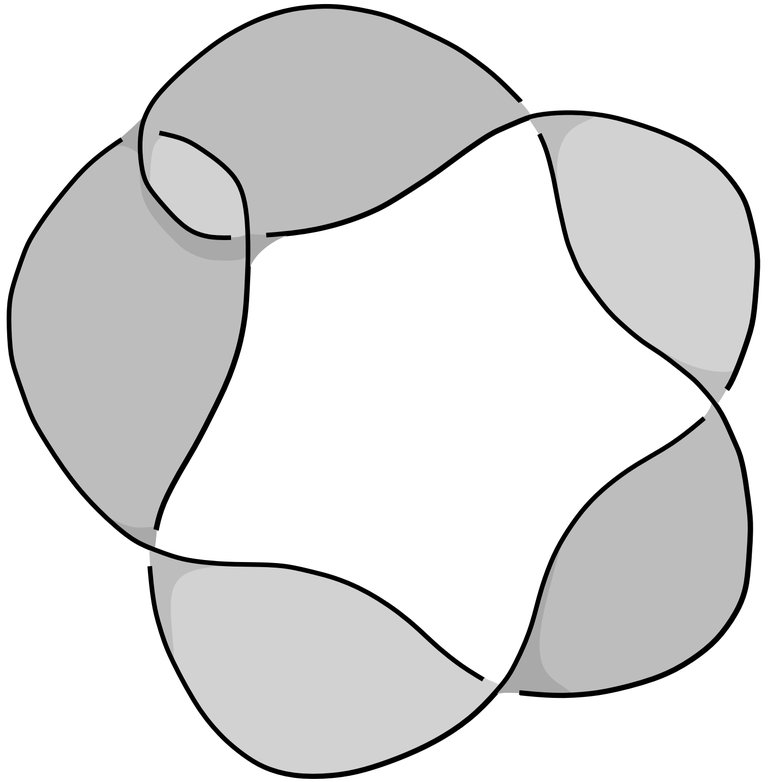

Our Project

Seifert Surfaces

- 6^1 knot also bounds a punctured torus in S^3
- Can we do better if we add an extra dimension?
- Yes!

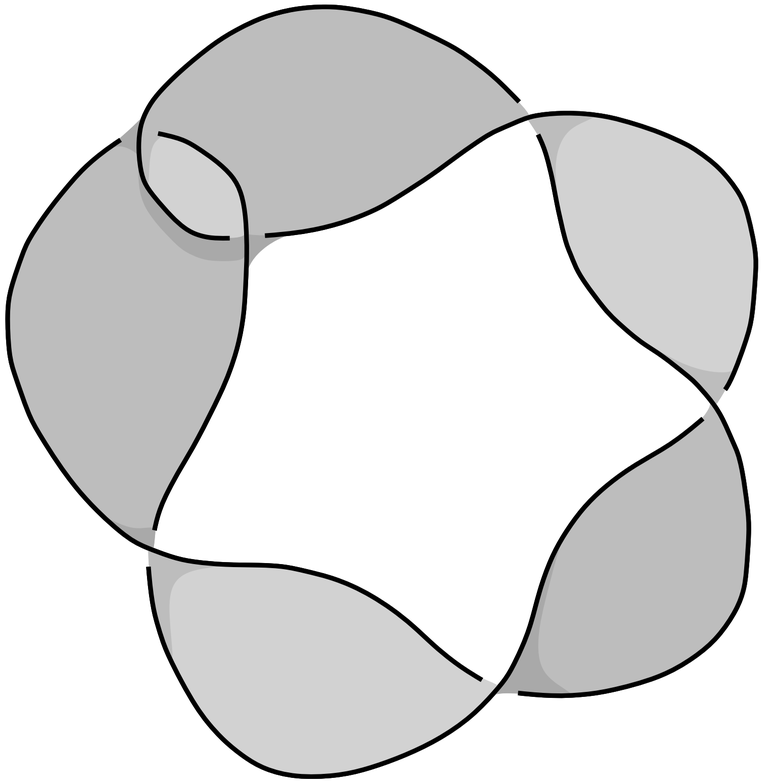

Our Project

Slice Surfaces

- Think of the knot on the surface S^3 of a ball of dimension 4
- Slice surfaces are equivalent to Seifert surfaces, just in dimension 4

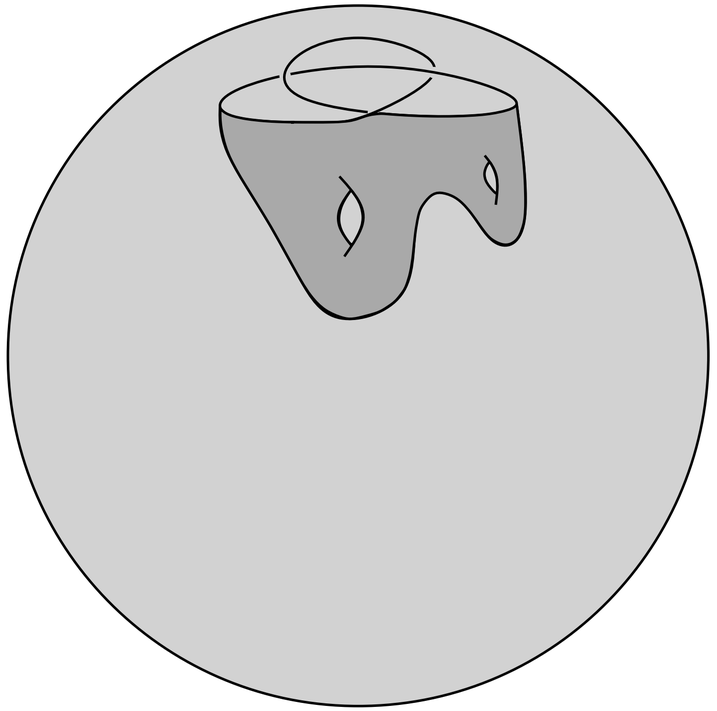

$S^3 = \partial B^4$

Our Project

Slice Surfaces

- One way we can look at slice surfaces is level sets

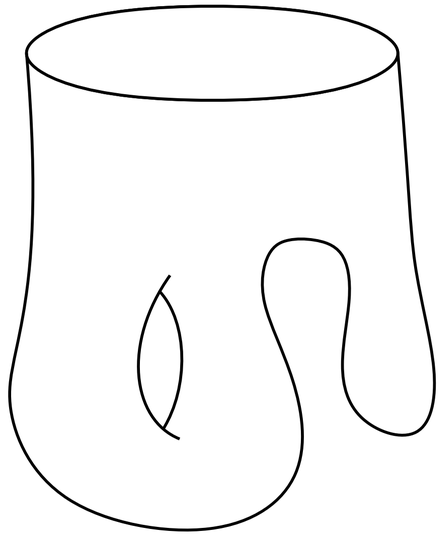

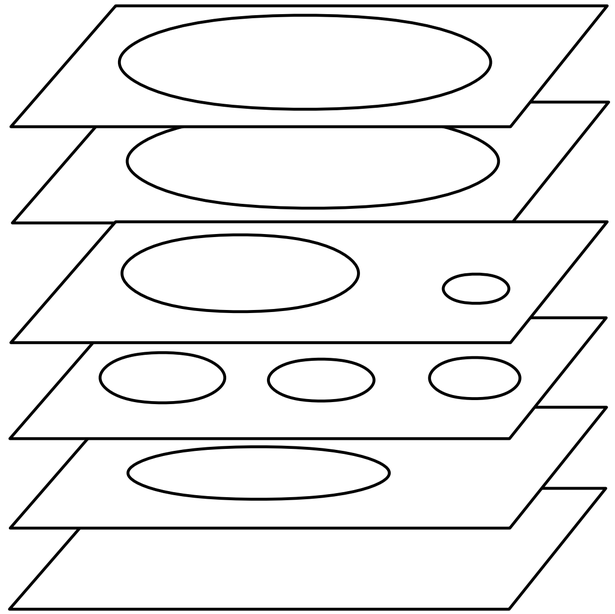

Our Project

Slice Surfaces

- Start with a strand
- Add some other strands
- Move the strands around (according to a set of moves)

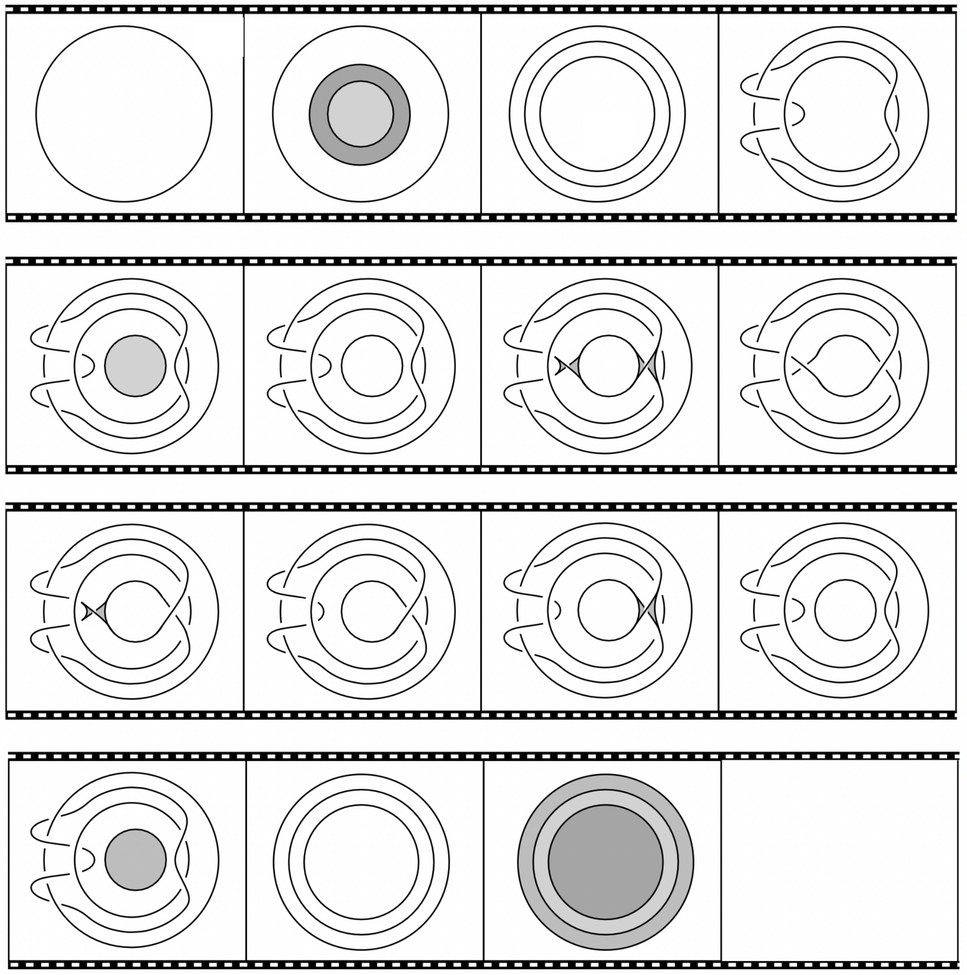

Braids

Our Project

Braids

- Every knot or link can be represented in braid form
- A braid is simply a set of n strings attached to a horizontal bar at the top and the bottom
- If you cut a knot at one point, you can attach it to two bars to make a braid

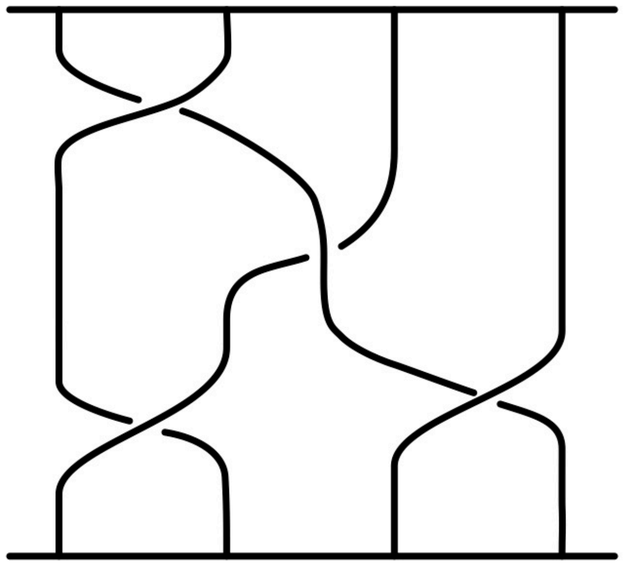

Our Project

Braids

- We represent braids through braid words
- If the nth strand crosses over the $n+1$th strand, represent as σ_n^{-1}
- If the nth strand crosses under the $n+1$th strand, represent as σ_n
- Braid word: $\sigma_1^{-1}\sigma_2\sigma_1^{-1}$

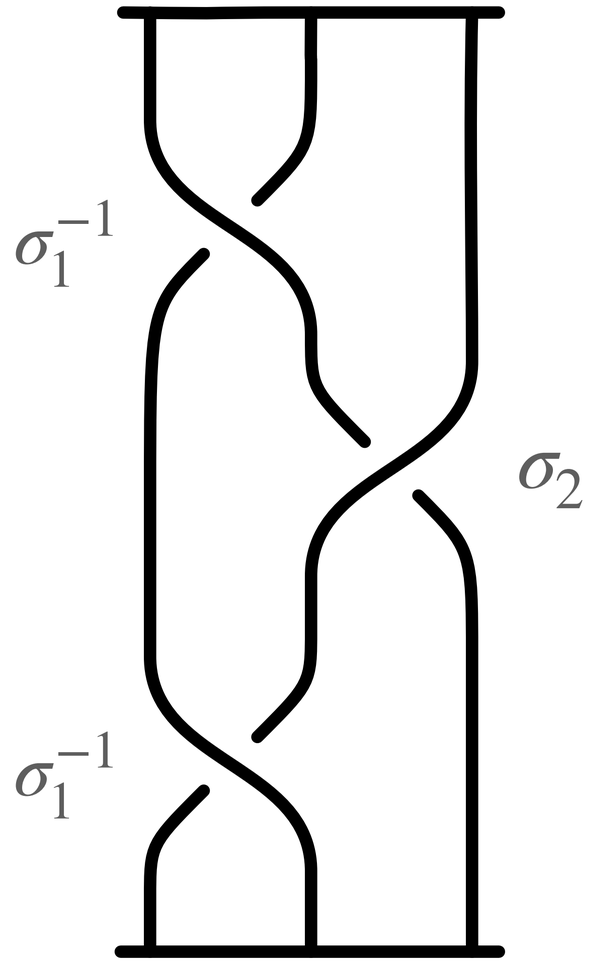

Our Project and Results

Our Project

Actions and Rewards

- Move 0: Remove a crossing
- Moves 1-10: Change the way the braid is presented
- Moves 11-12: Add crossings
- Inaction Penalty (Encourage our agent to find the answer as quick as possible)

- Reward: -1
- Reward: 0
- Reward: -1
- Reward: -0.05 (Hyperparameter)

Our Project

Rewards for Results

- If a move produces an unknotted, unlinked component, that component is removed and a reward of +1 is received
- If the agent builds a surface with the maximal Euler characteristic, the environment closes
- If the agent fails to build such a surface, a penalty of -350 is given

Our Project

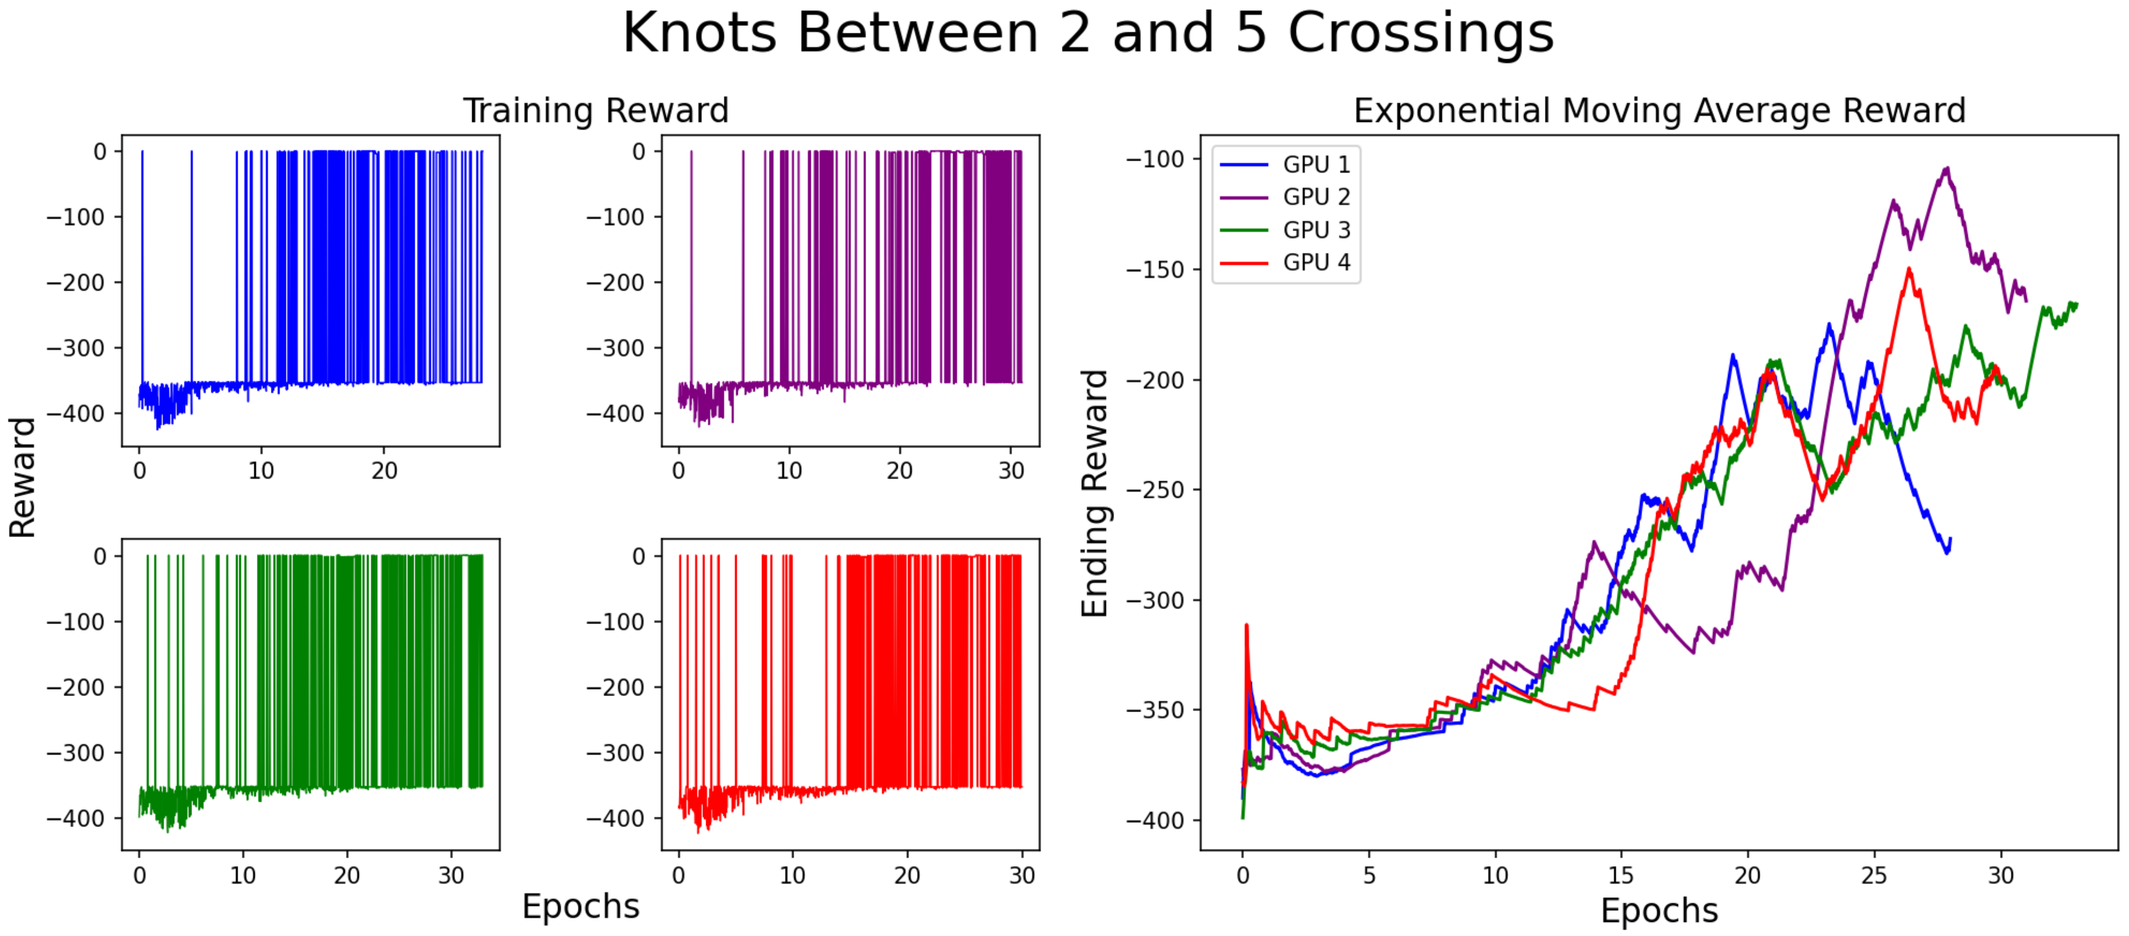
Knots Between 2 and 5 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

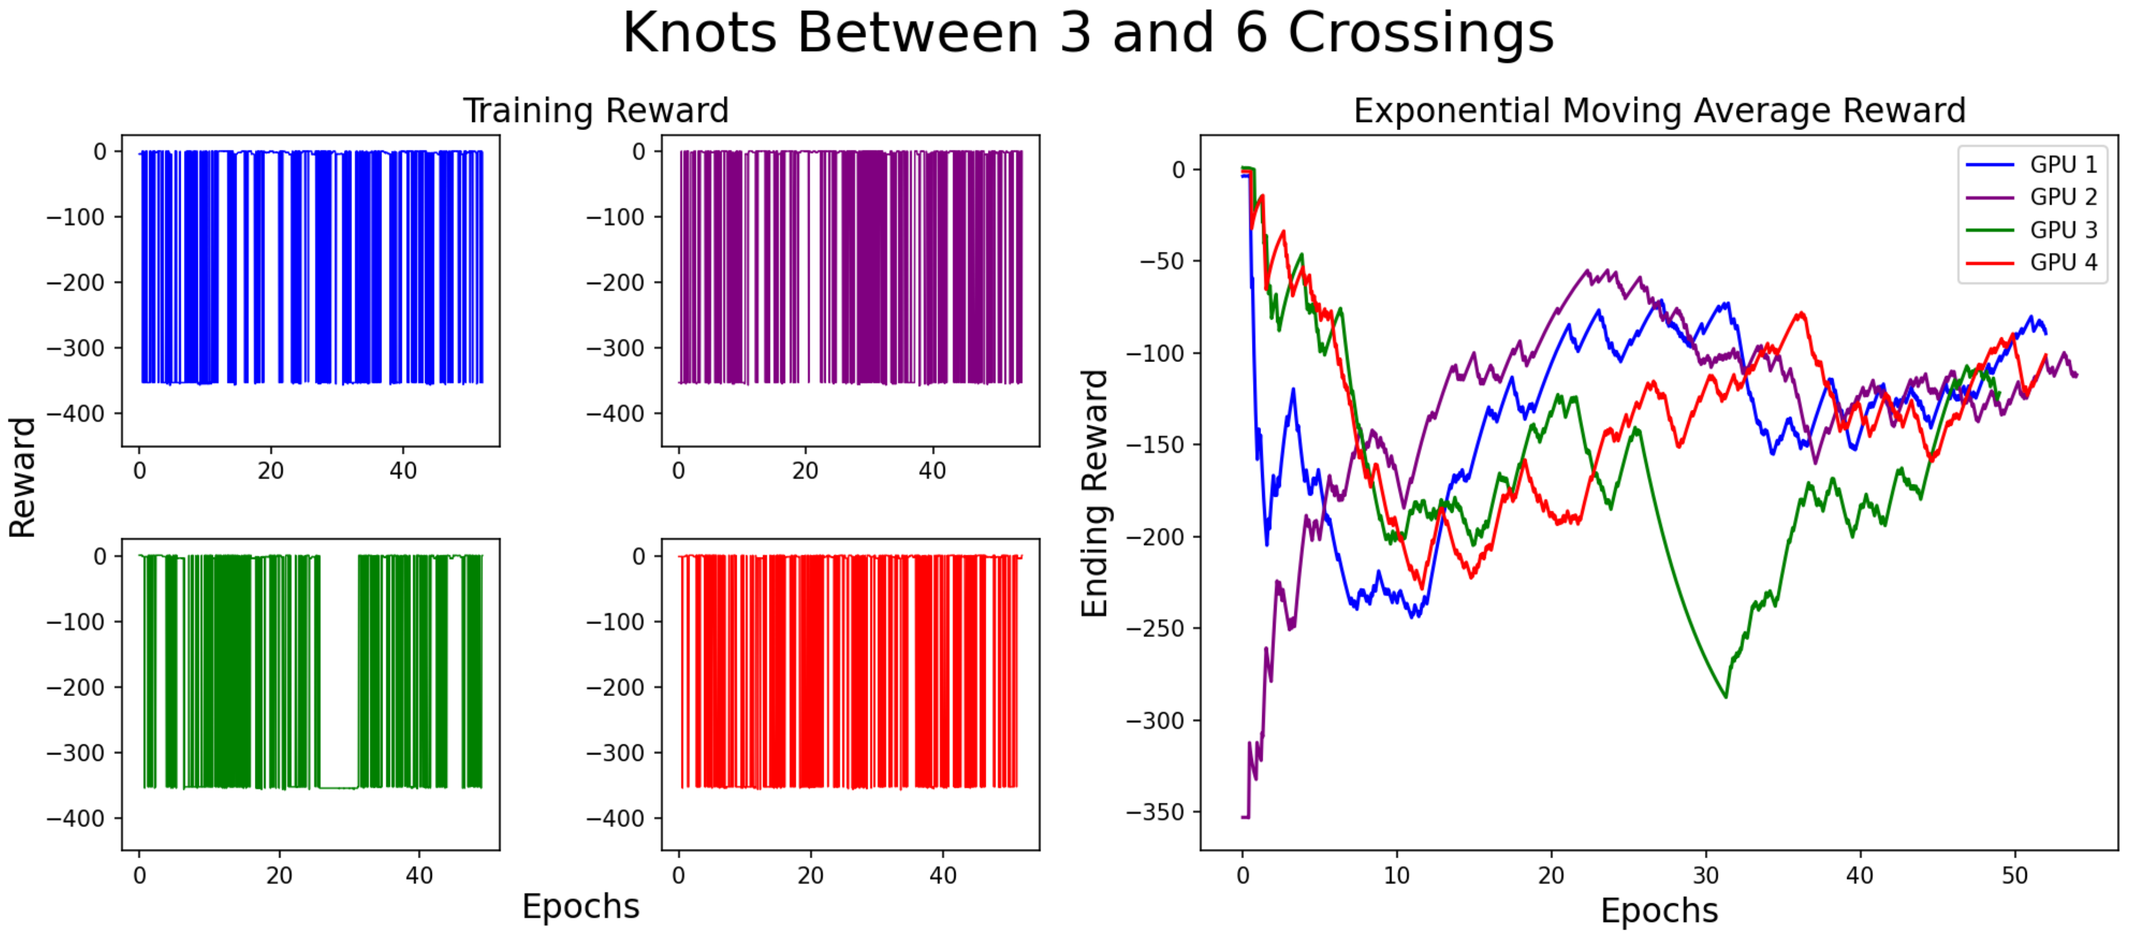
Knots Between 3 and 6 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

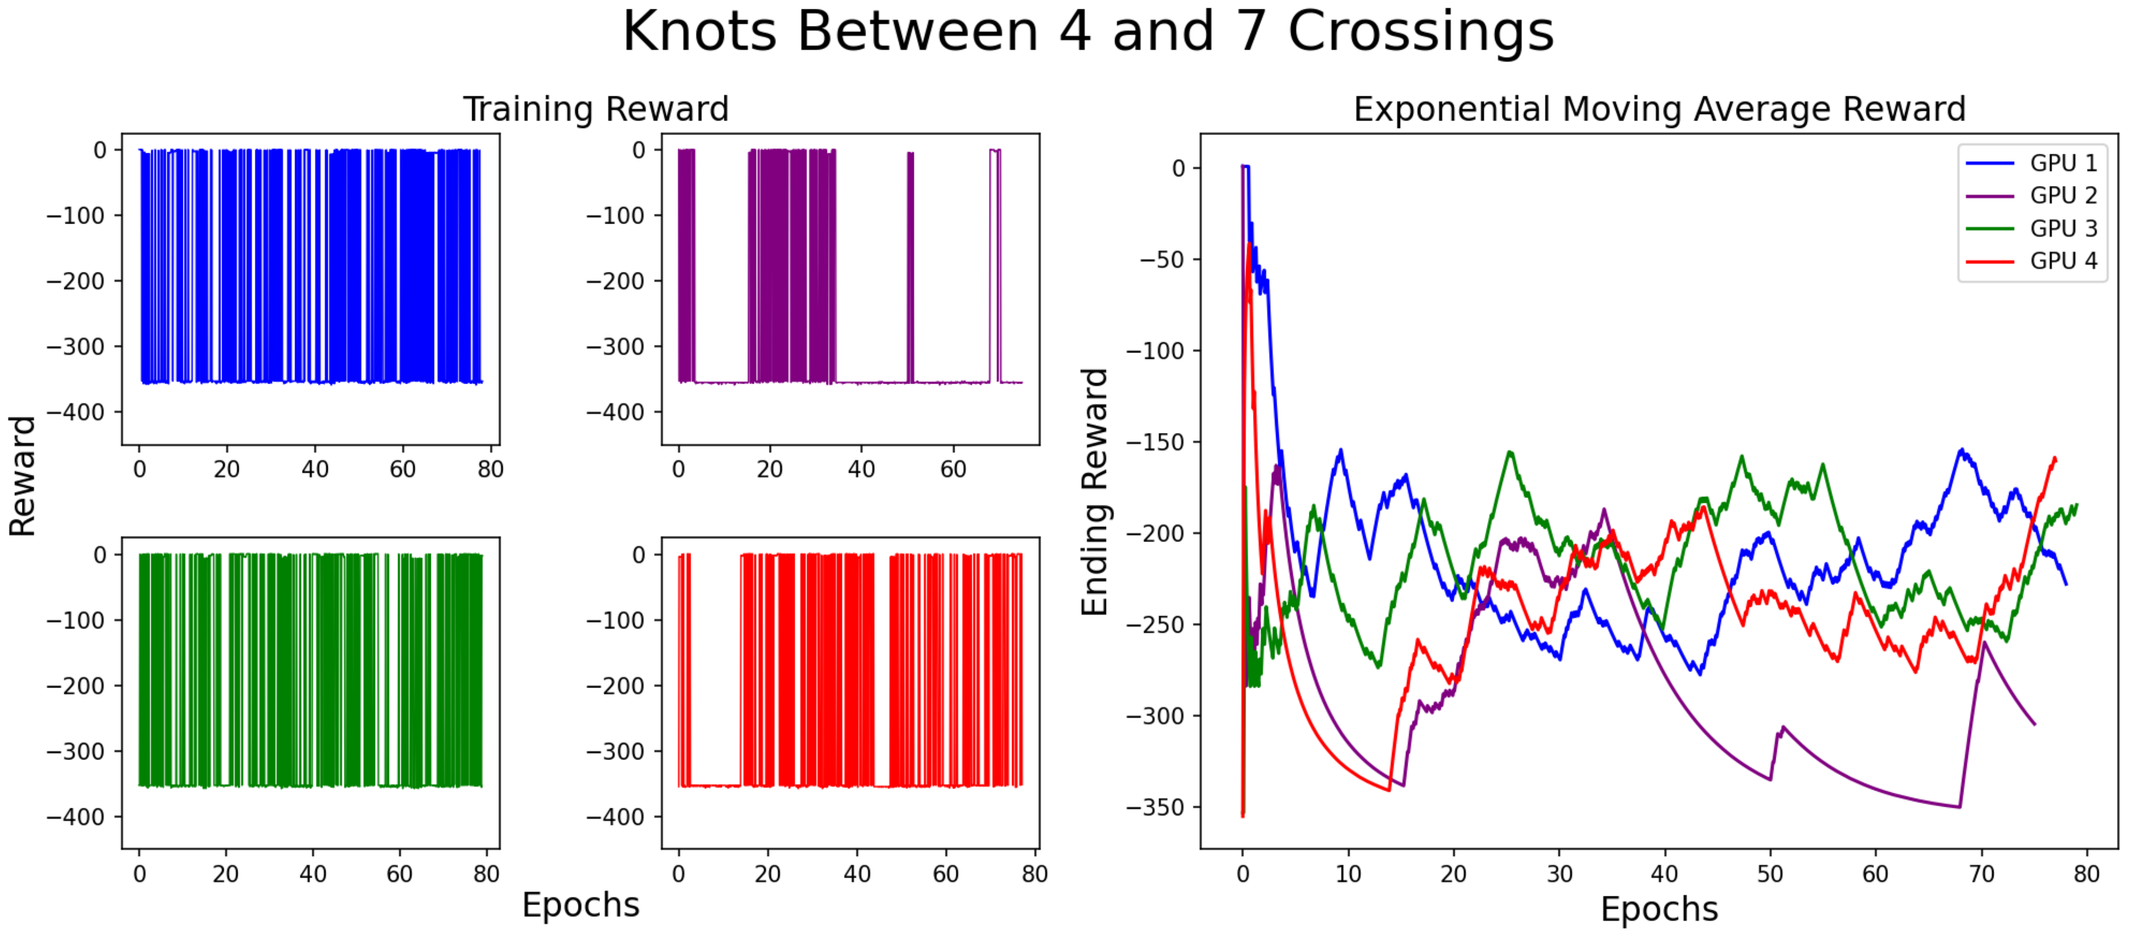
Knots Between 4 and 7 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

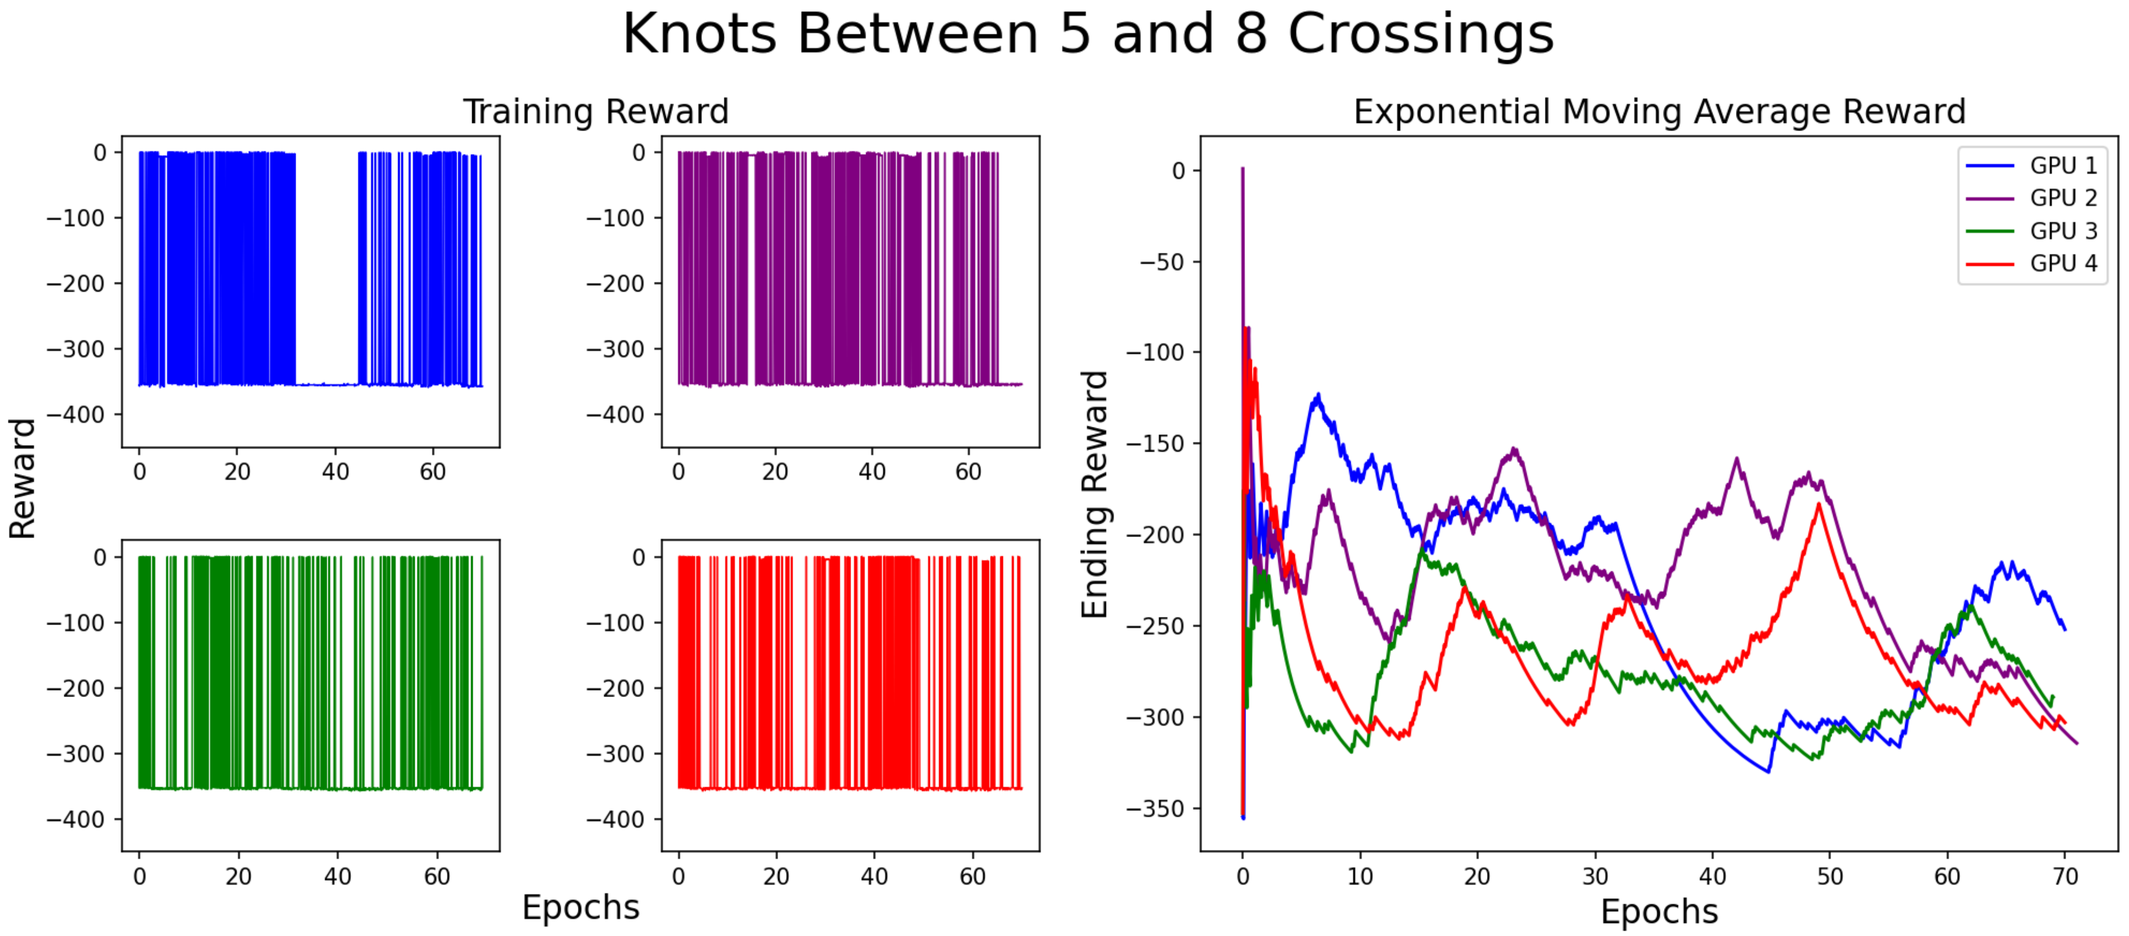
Knots Between 5 and 8 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

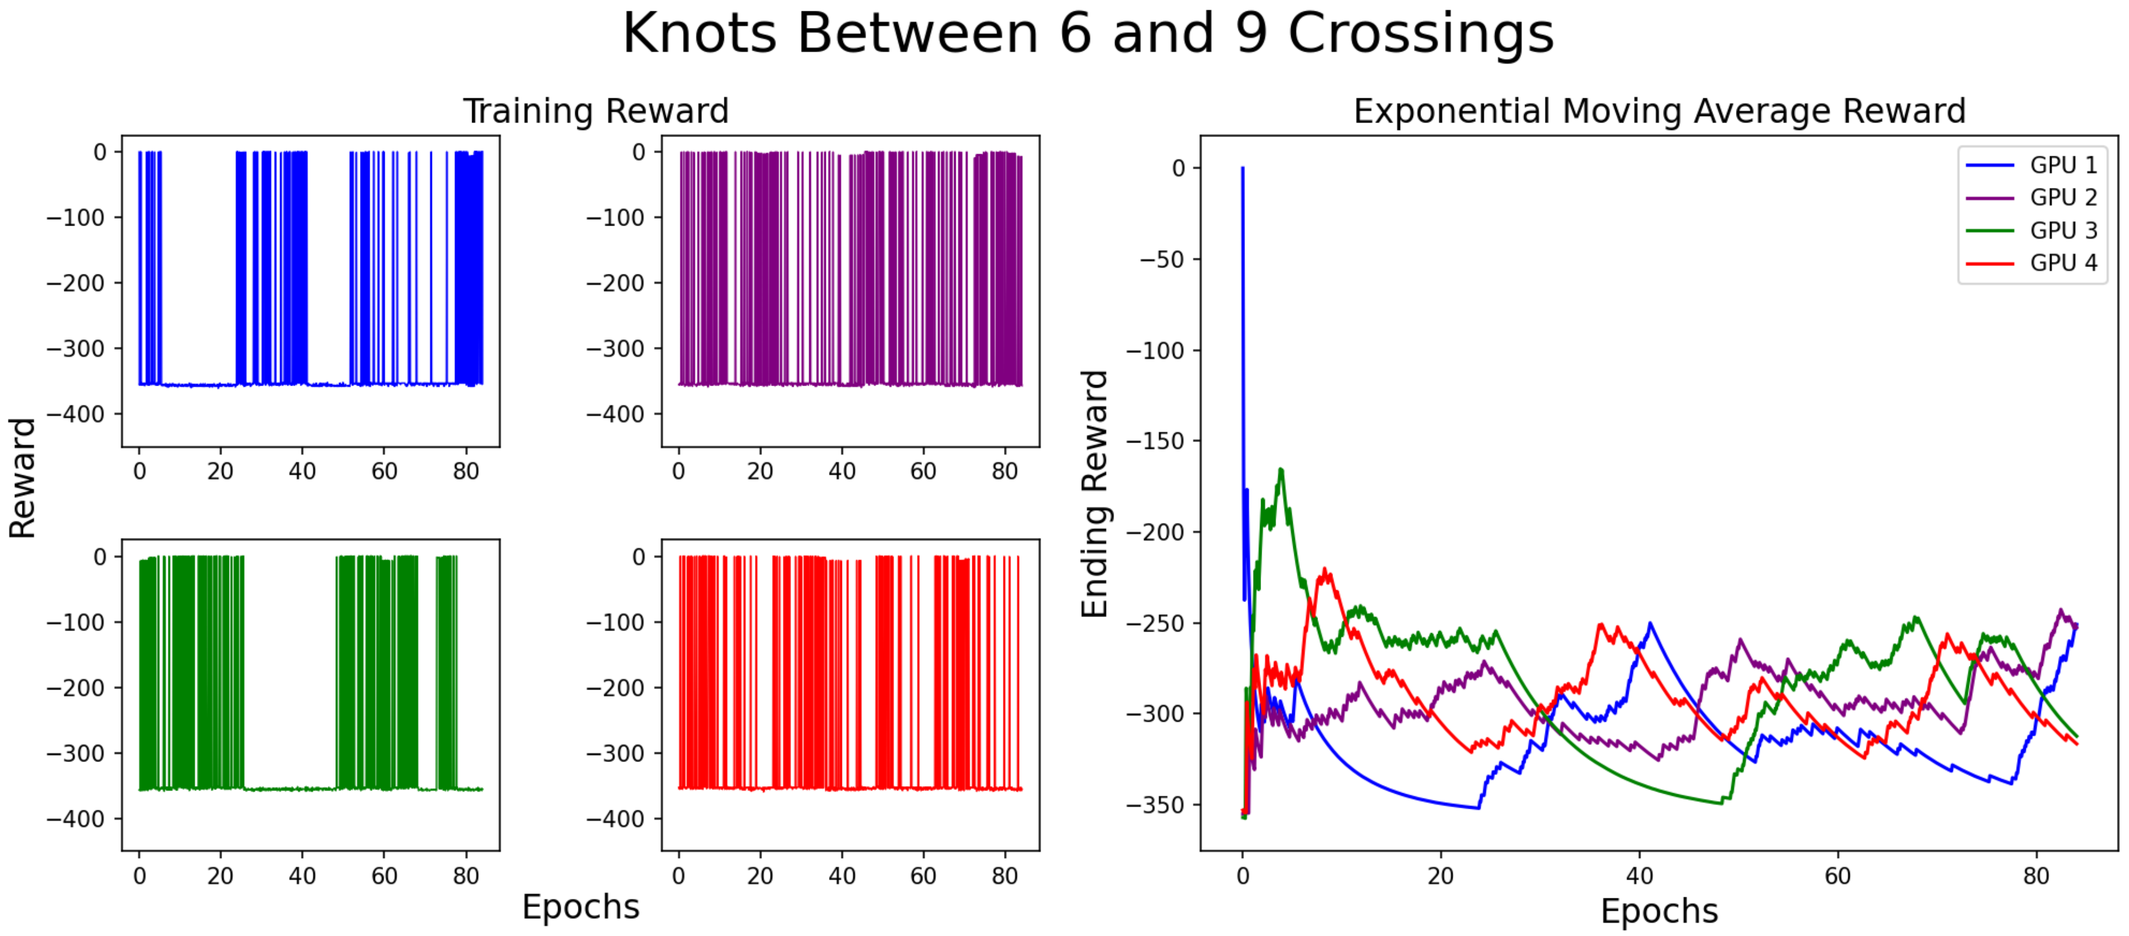
Knots Between 6 and 9 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

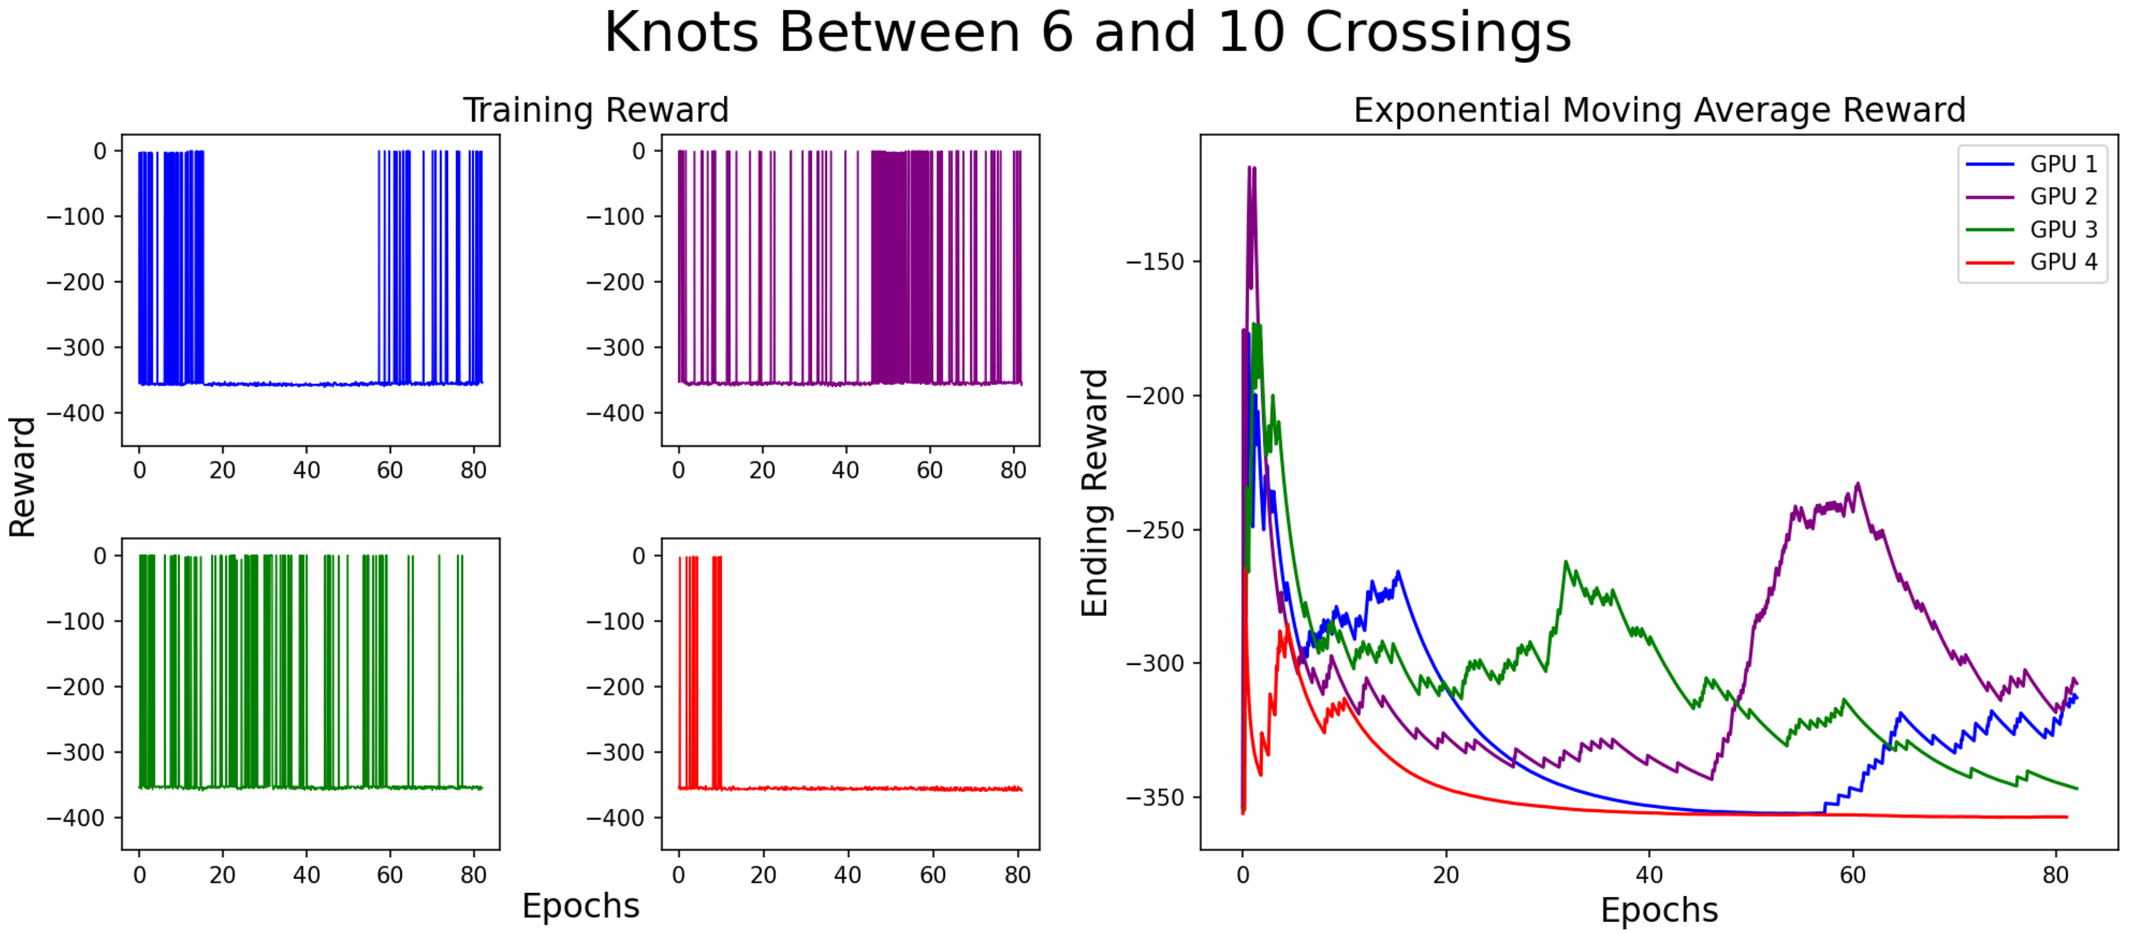
Knots Between 6 and 10 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

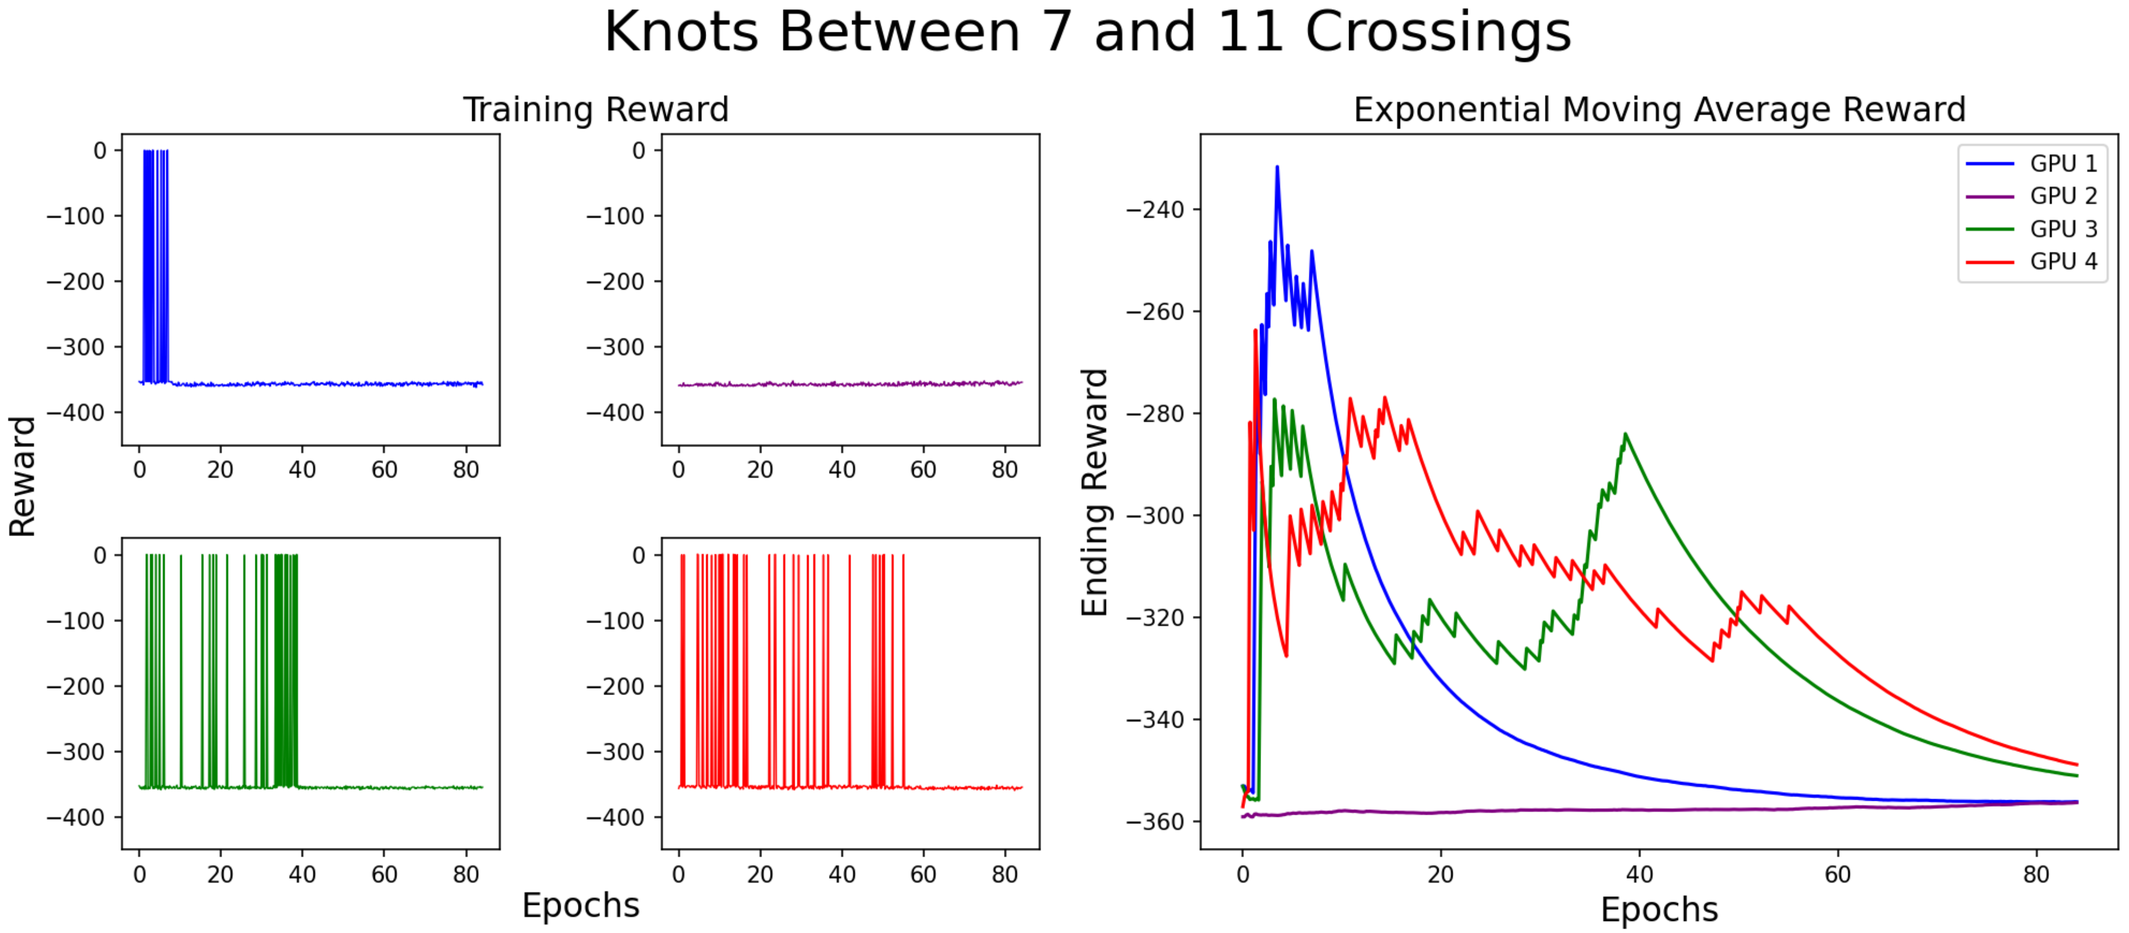
Knots Between 7 and 11 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

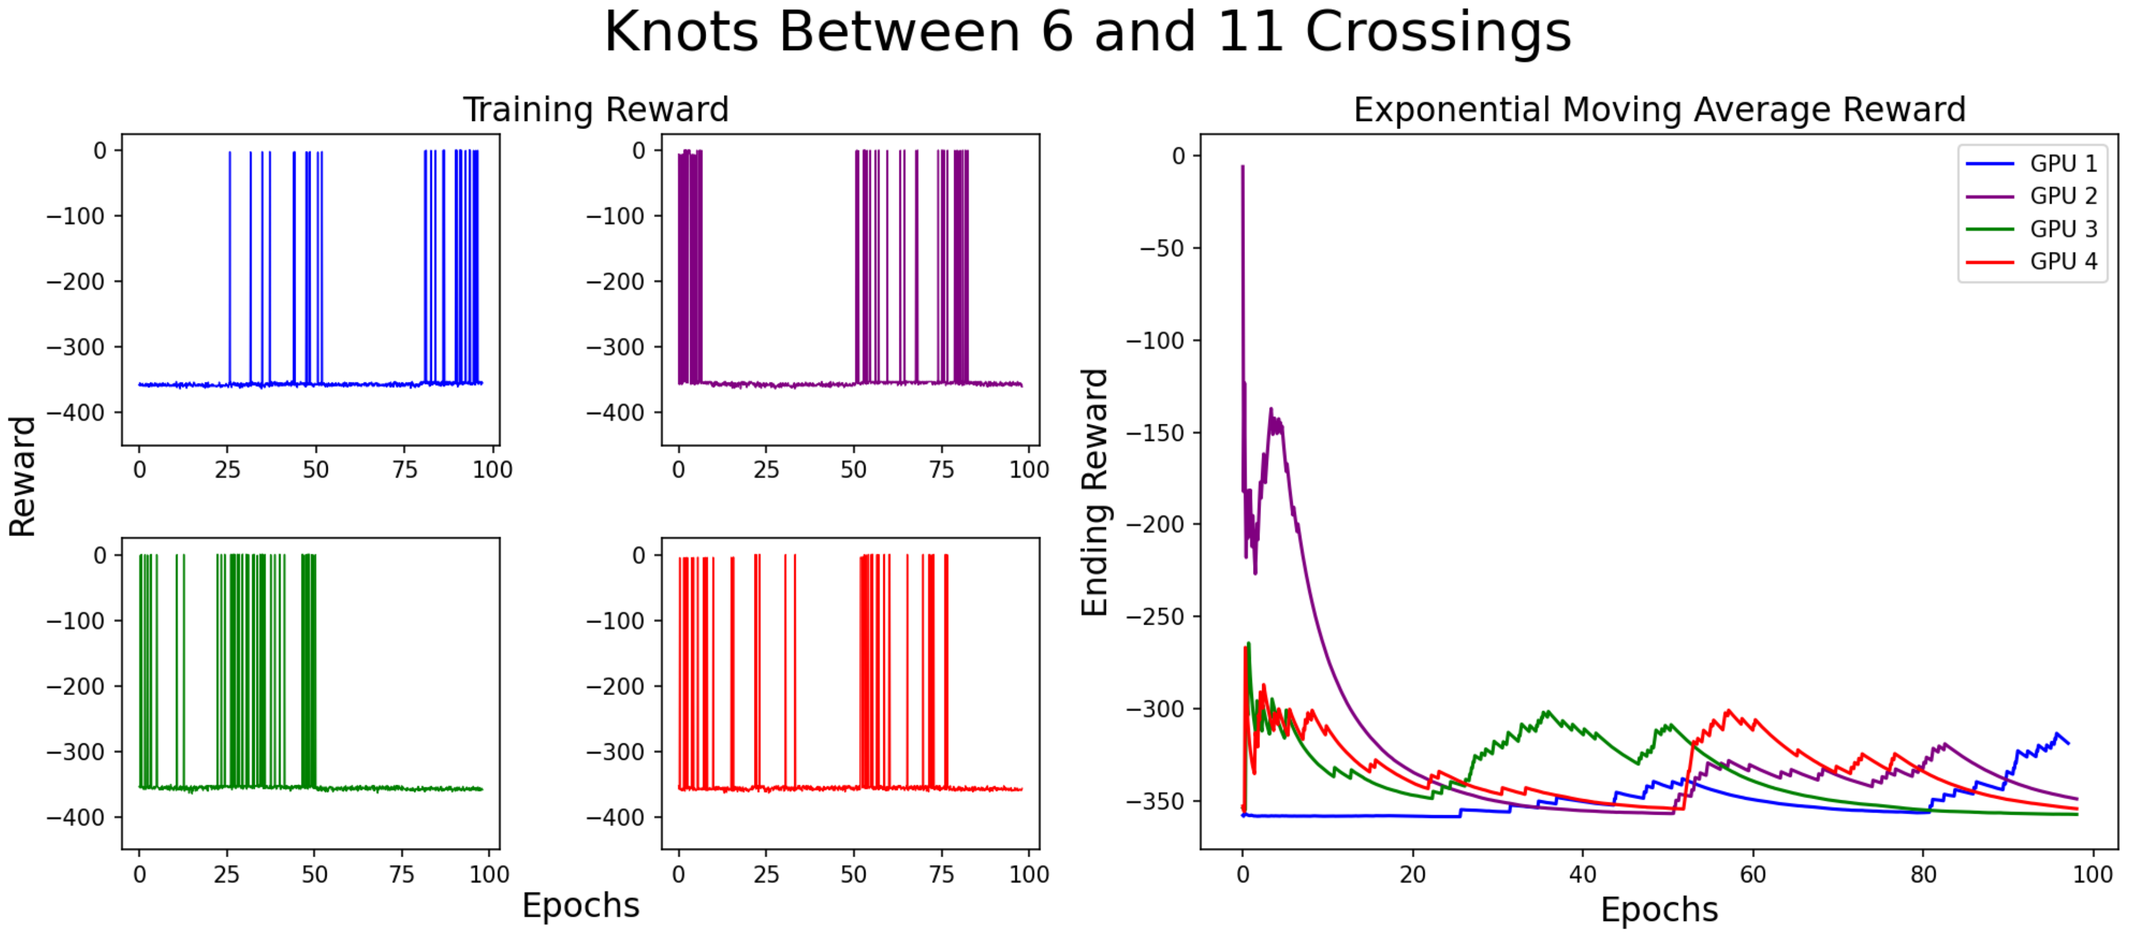
Knots Between 6 and 11 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

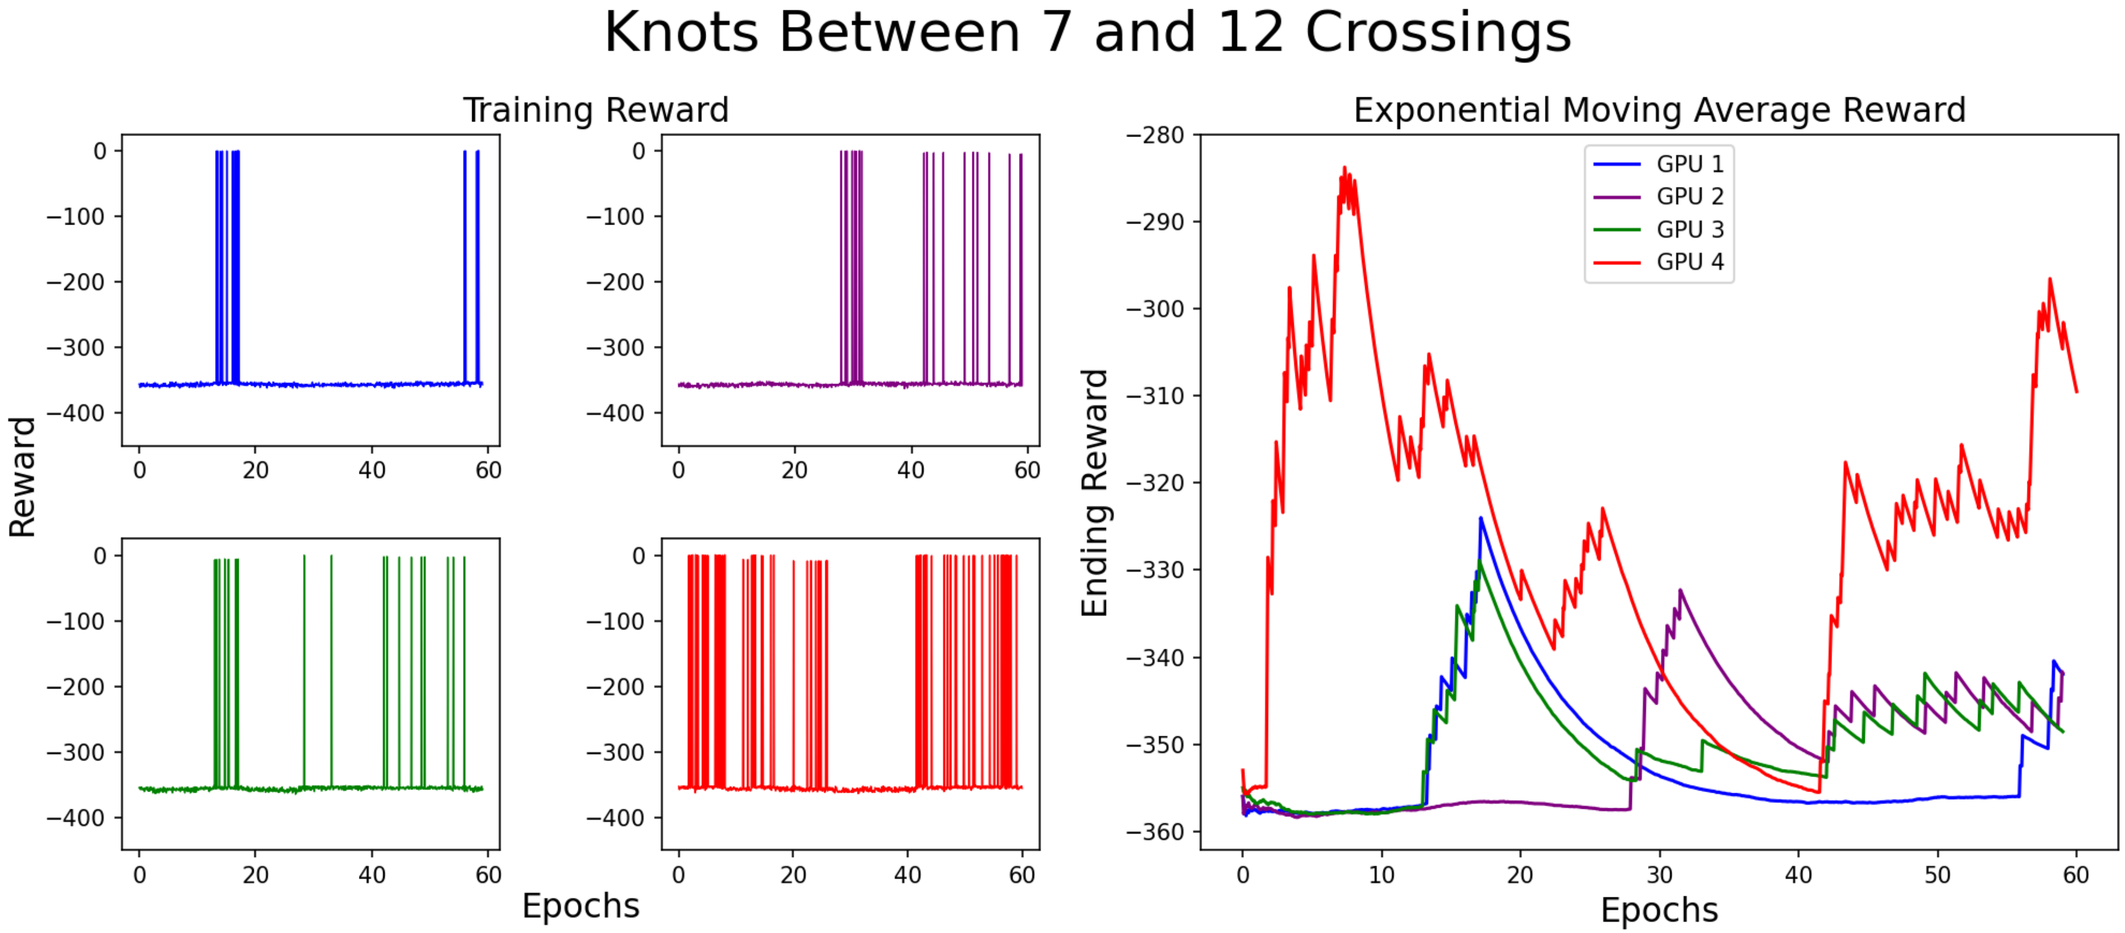
Knots Between 7 and 12 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

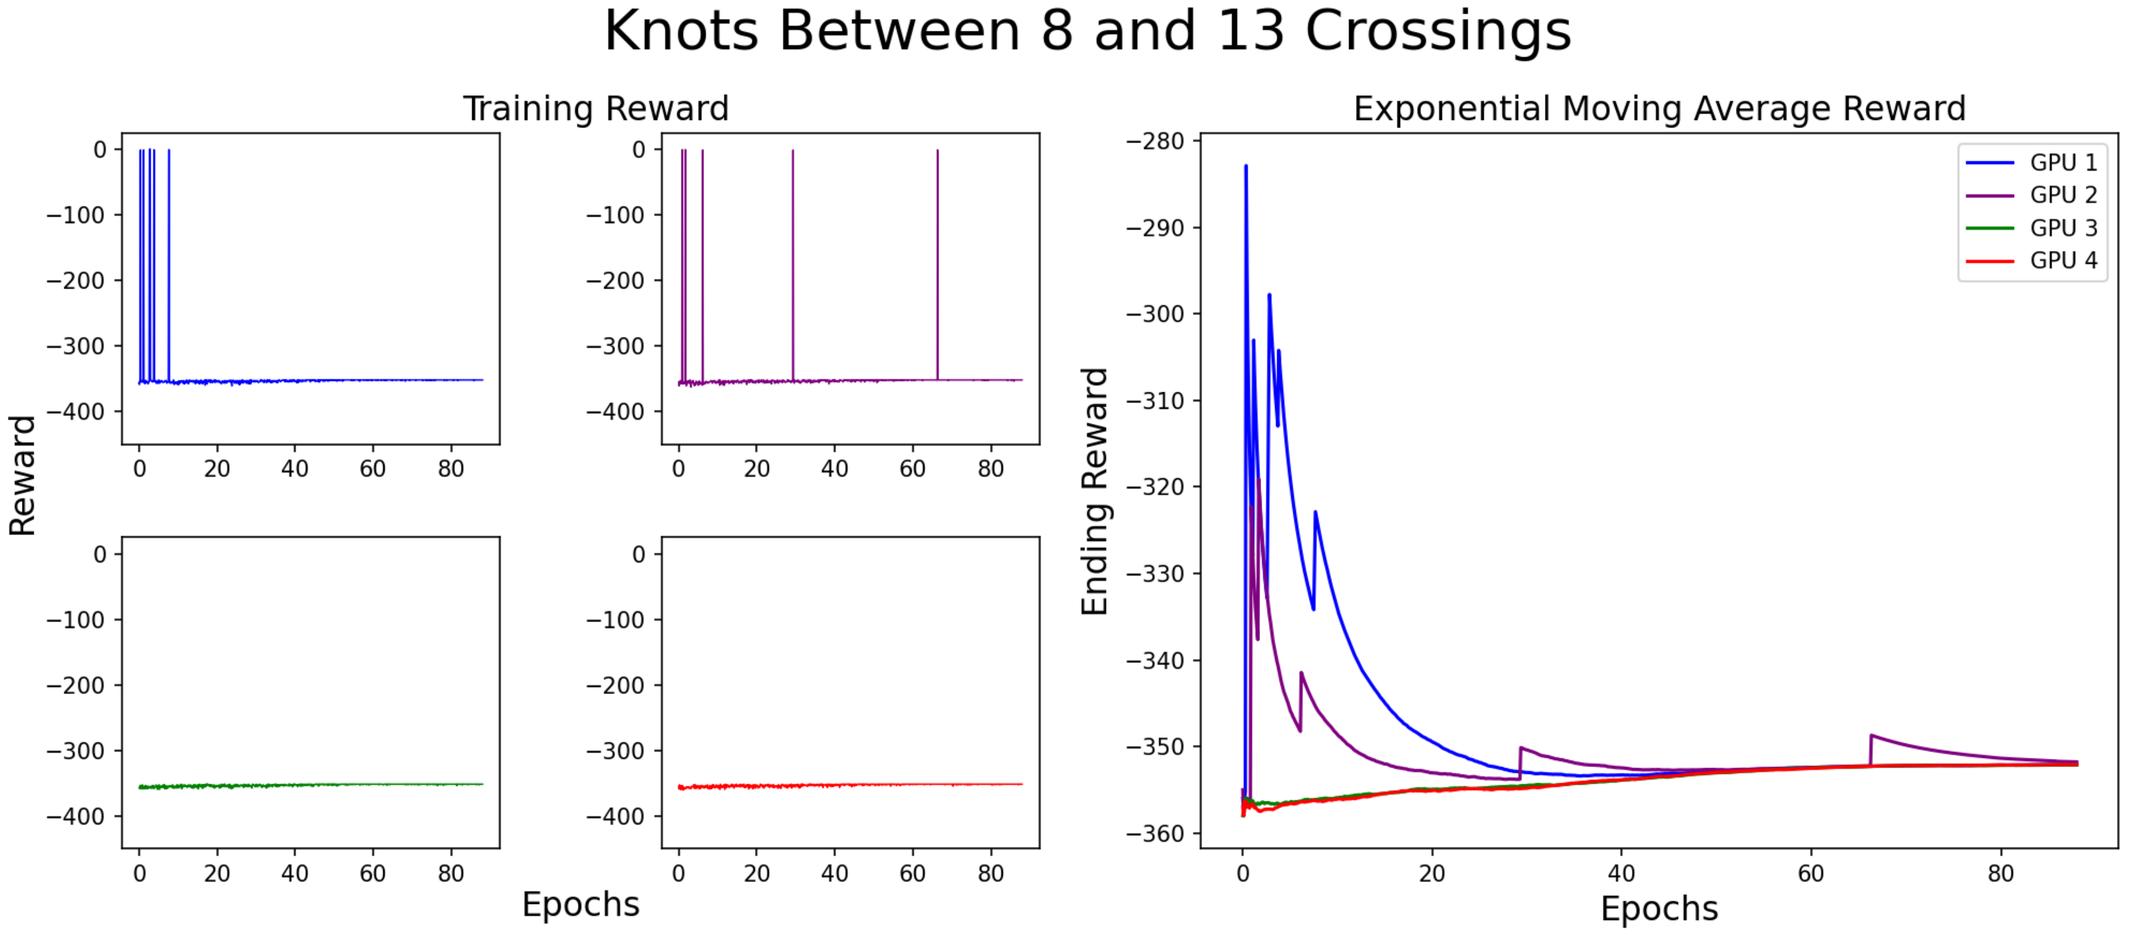
Knots Between 8 and 13 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Our Project

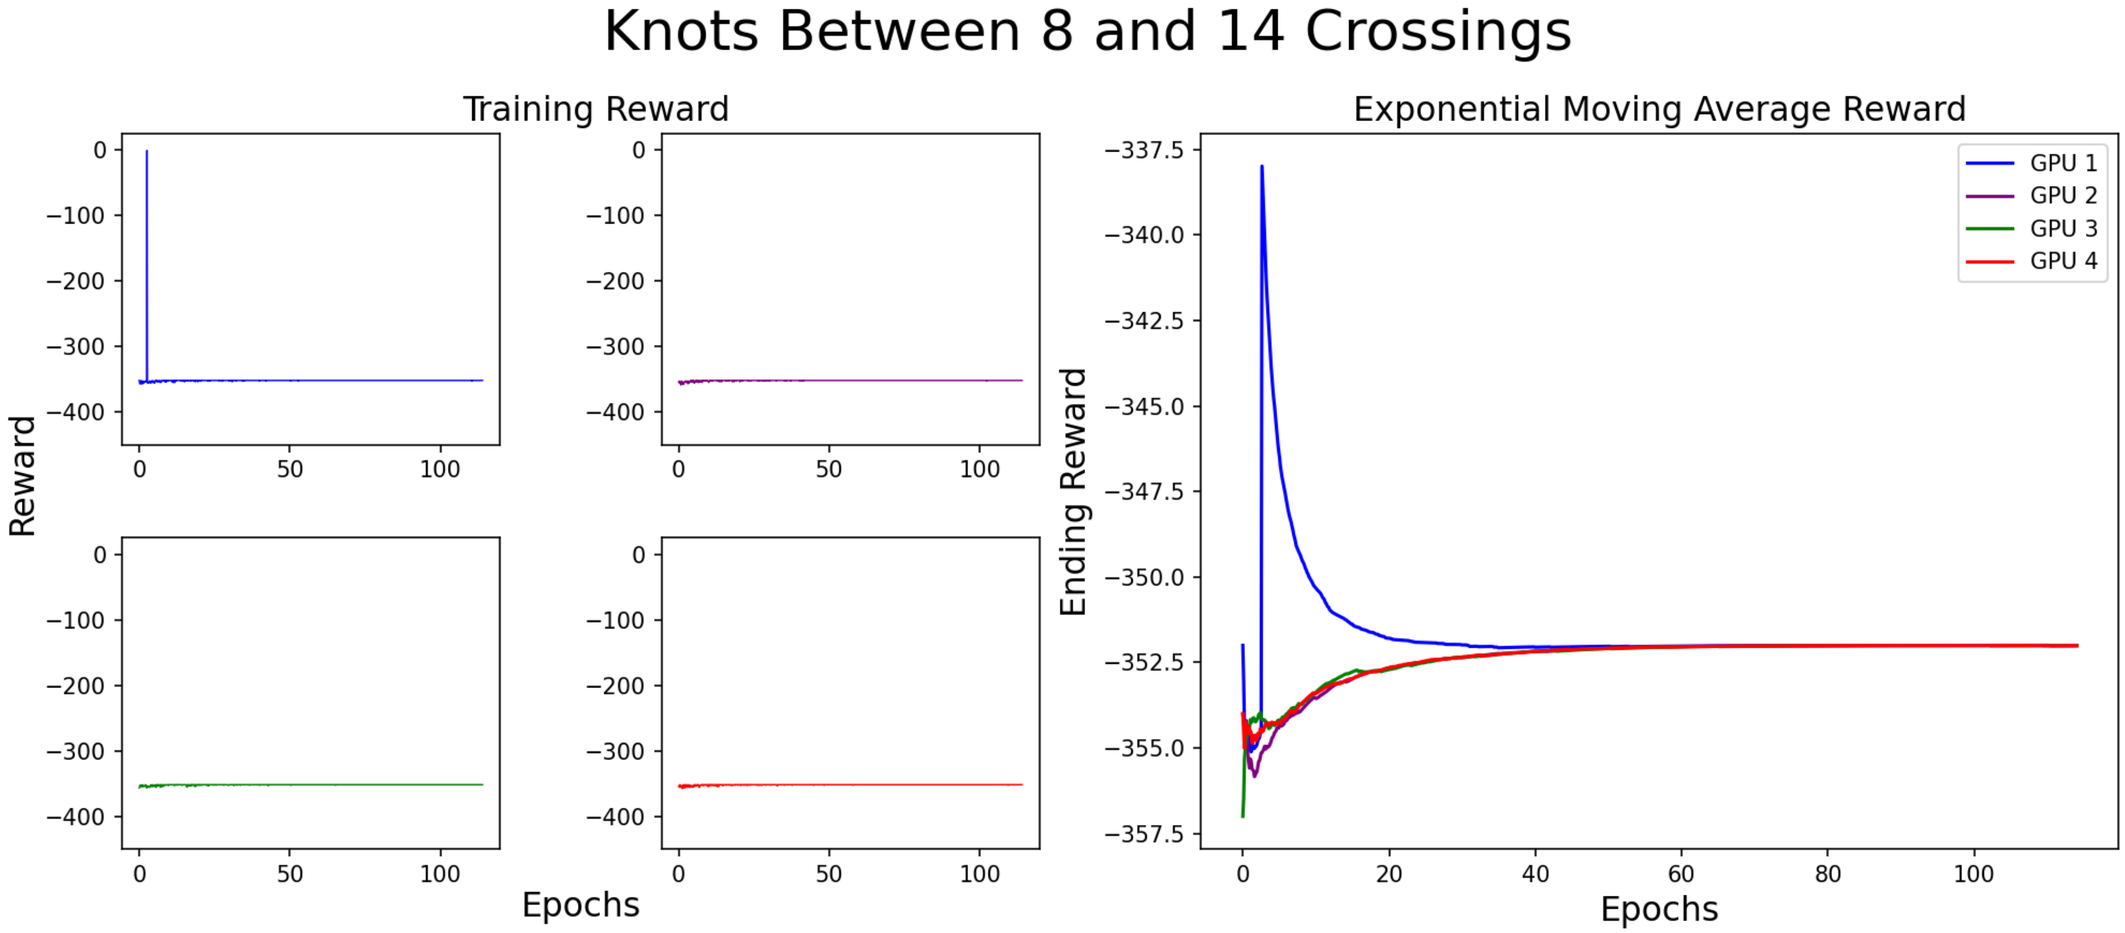
Knots Between 8 and 14 Crossings
Training Reward
Reward
Epochs
Exponential Moving Average Reward
Ending Reward
Epochs
GPU 1
GPU 2
GPU 3
GPU 4

Future Work

- Get above 13 crossings
- Increase network size
  - We ran into GPU memory issues when we made the network too large
- Maybe braids are not the best way to solve this problem
- Potentially change reward/penalty system