Arthur Juliani

Machine Learning Engineer
Unity Technologies
@awjuliani

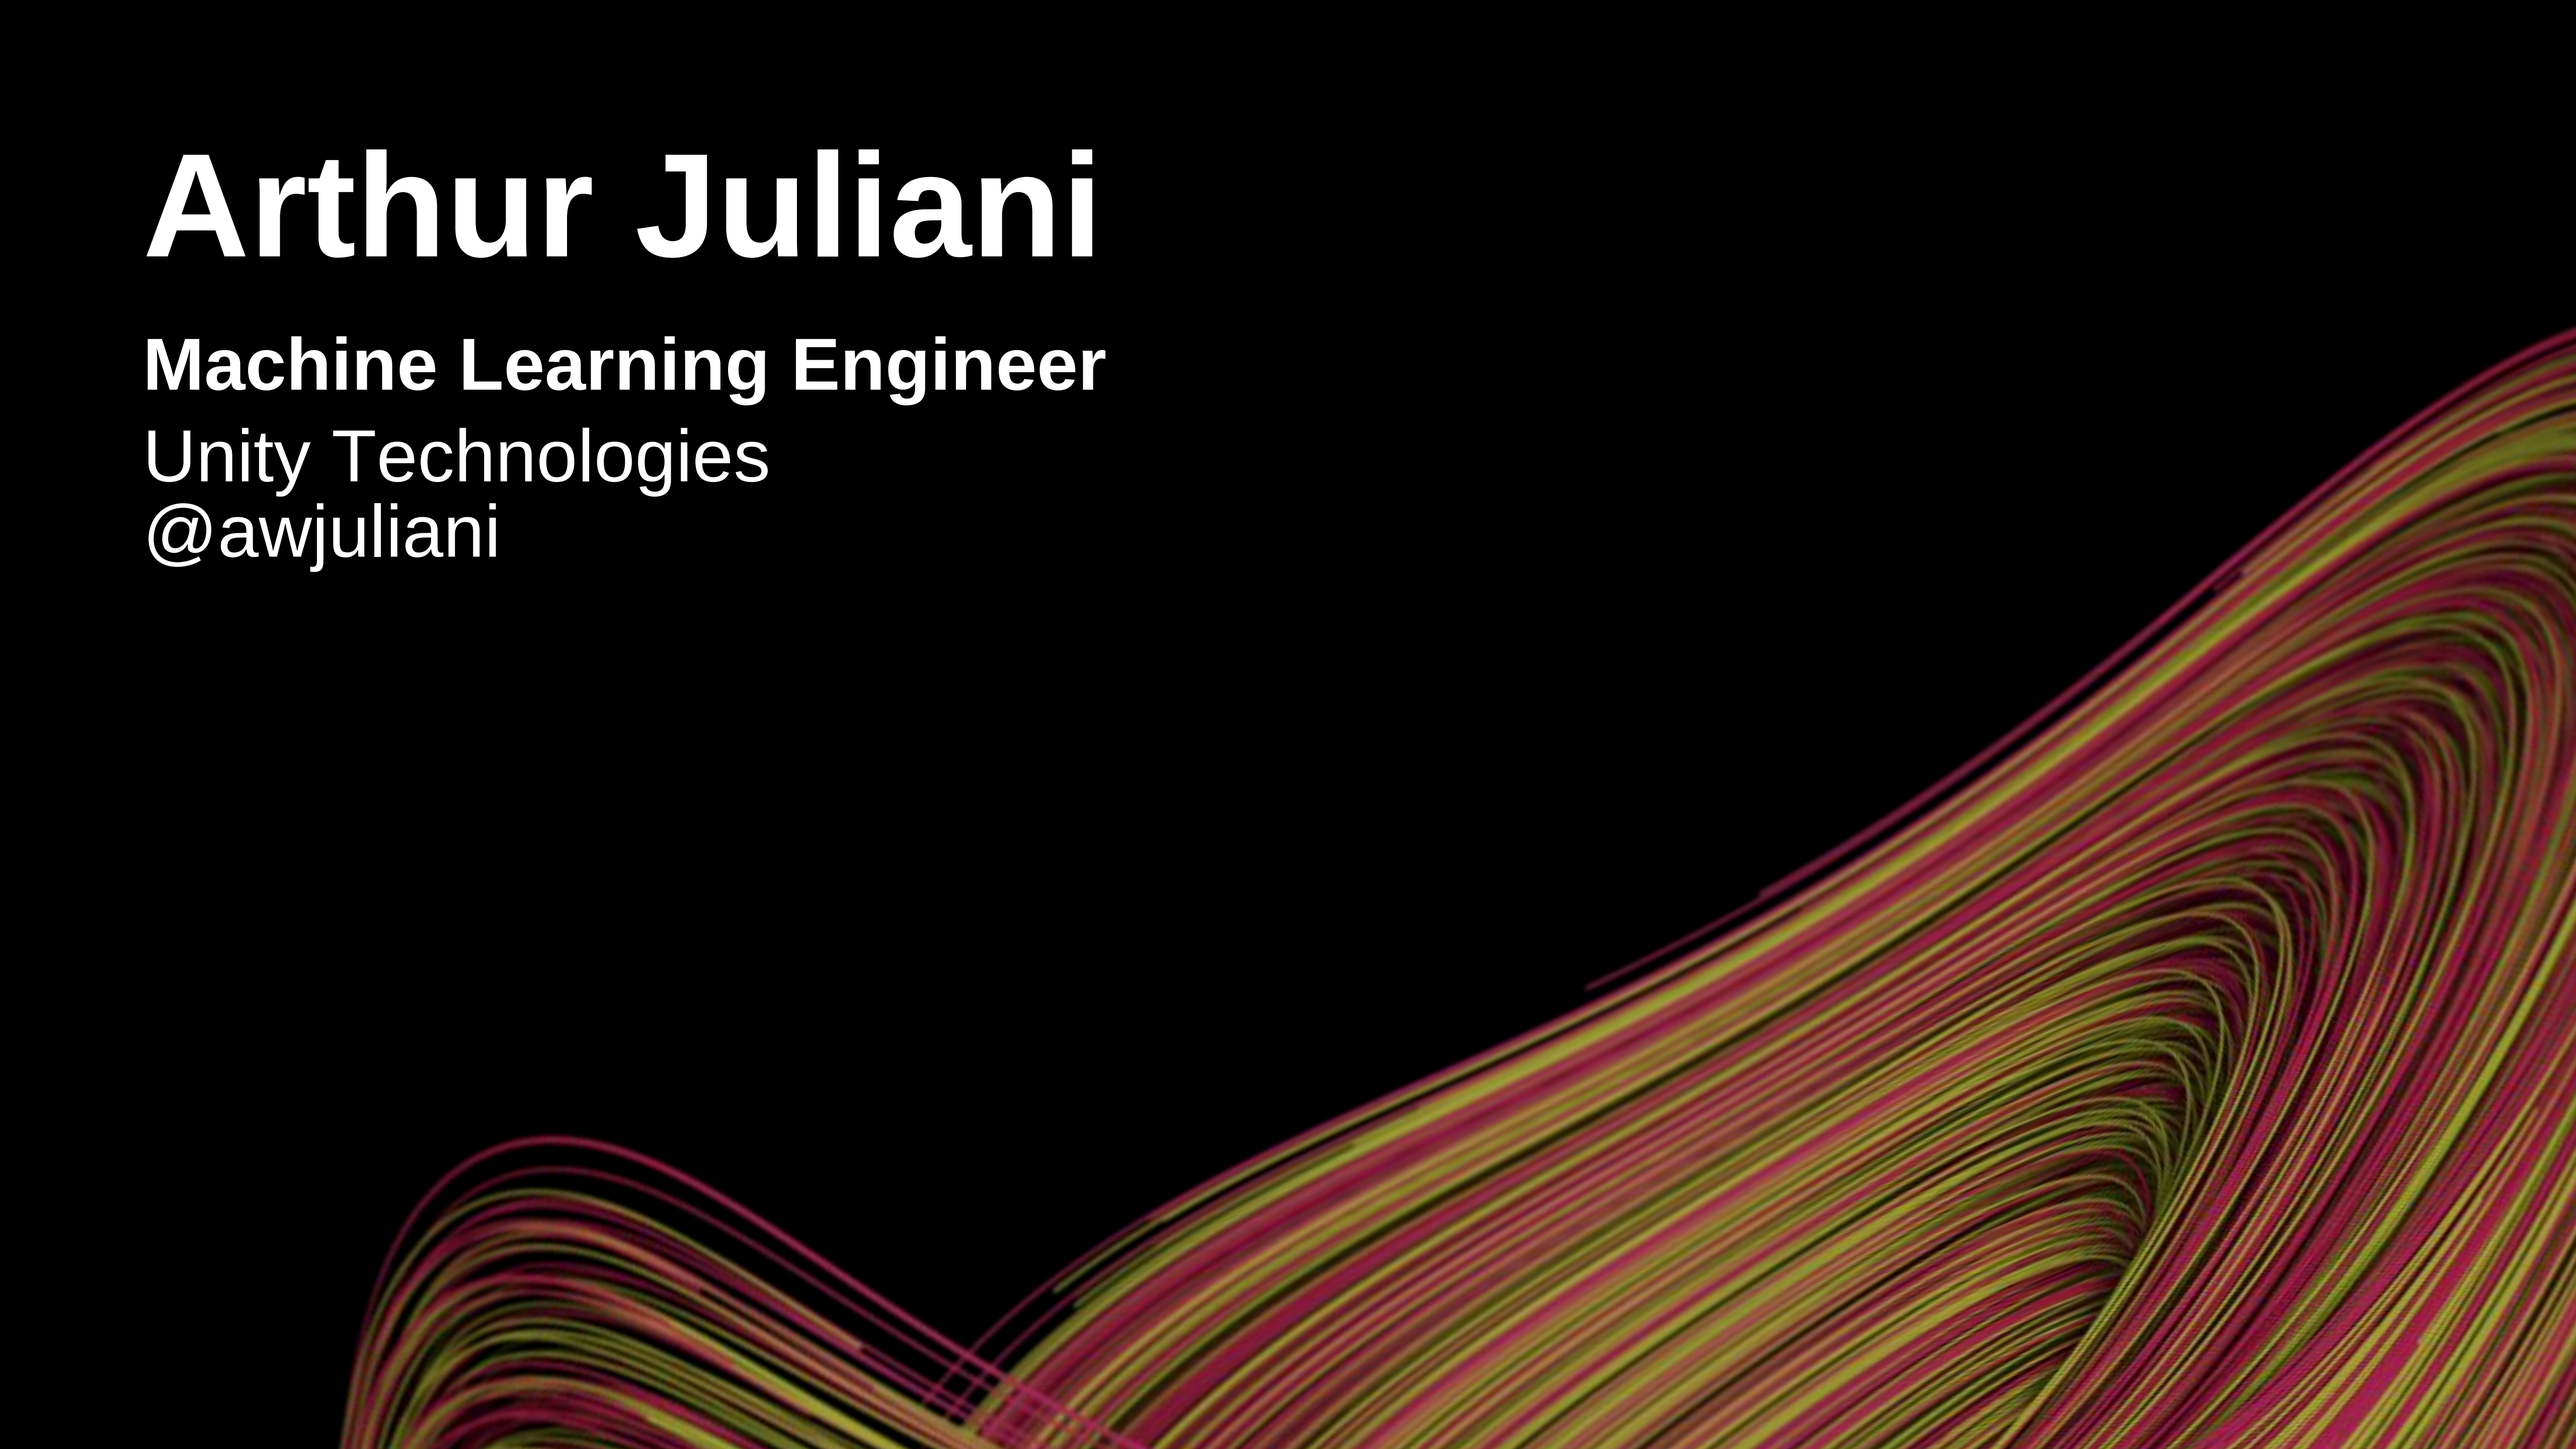

Unity ML-Agents: A flexible platform for Deep RL research

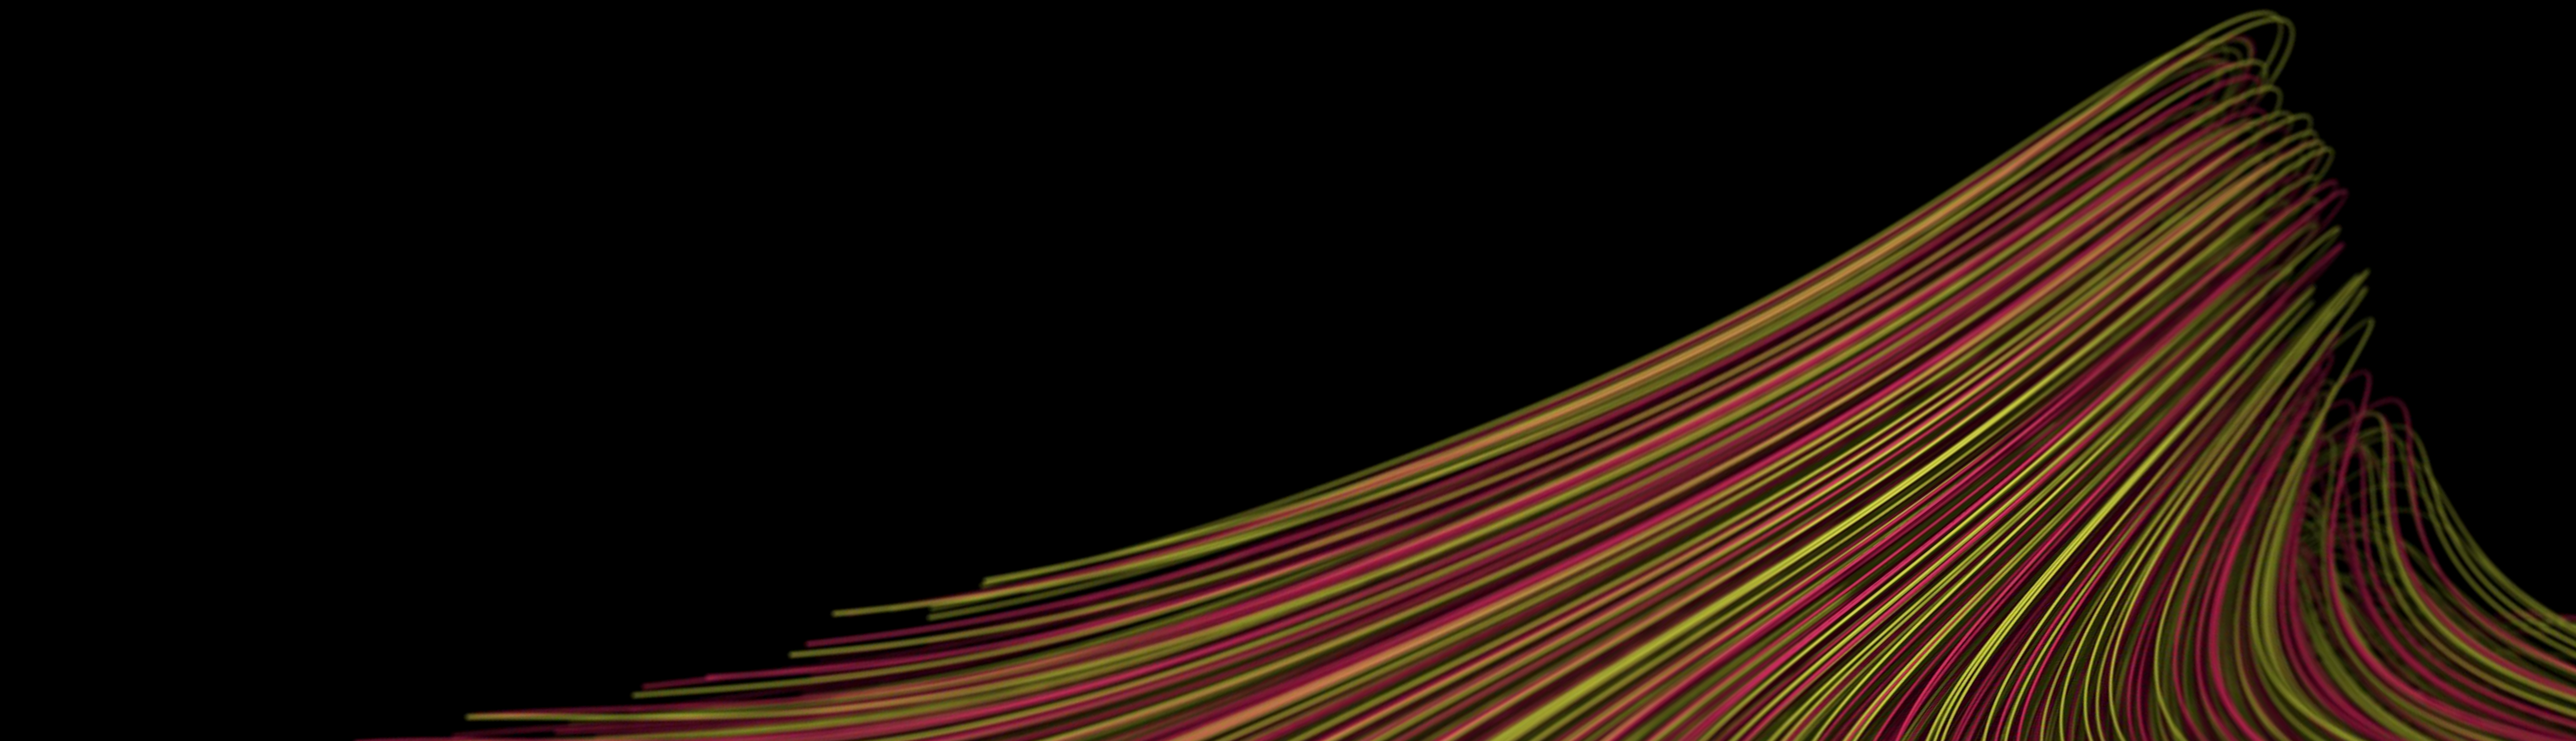

About Unity

“Creation Engine”

- Games
- AR/VR
- Cinematics
- Simulations
- 40+ Platforms
- Free for personal use

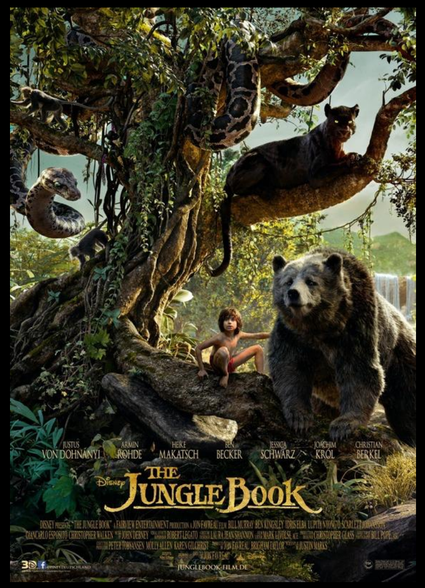

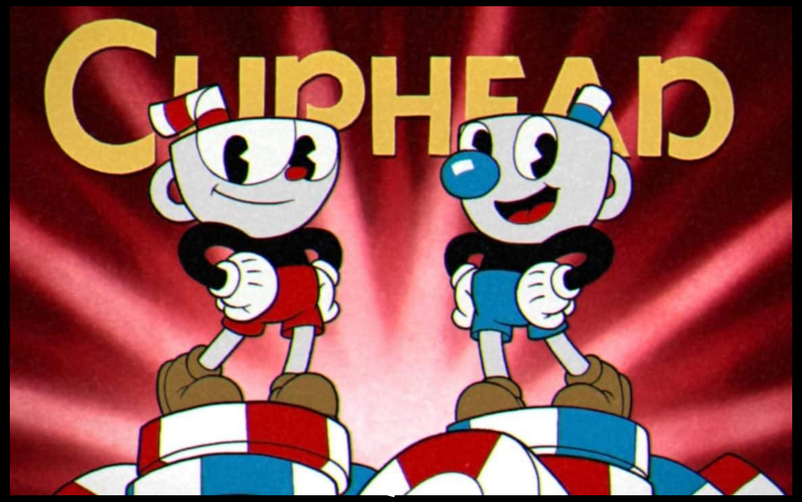

Research Environments

Visual Complexity

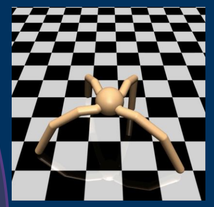

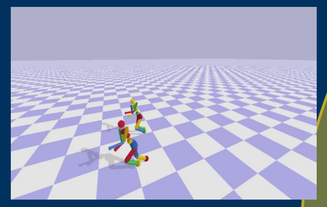

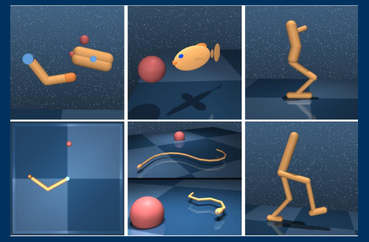

Physical Complexity

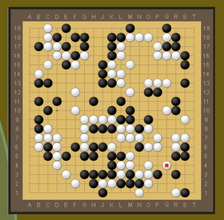

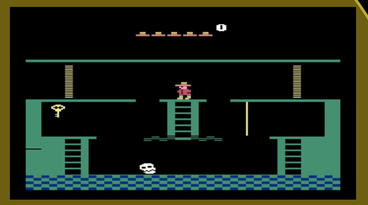

Cognitive Complexity

The Unity Ecosystem

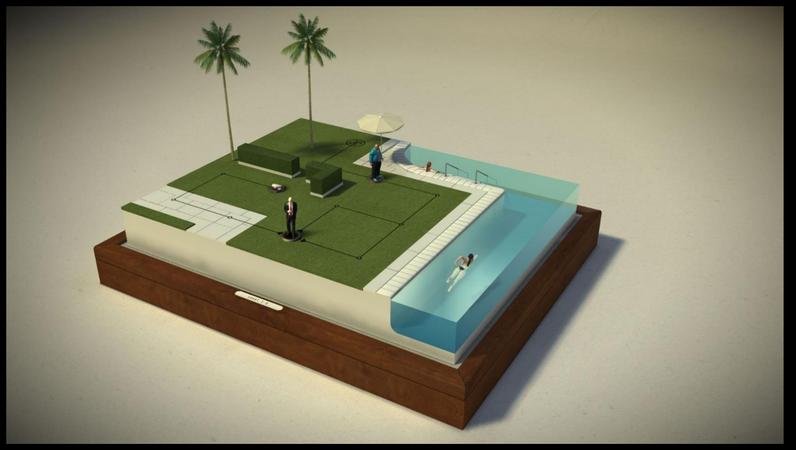

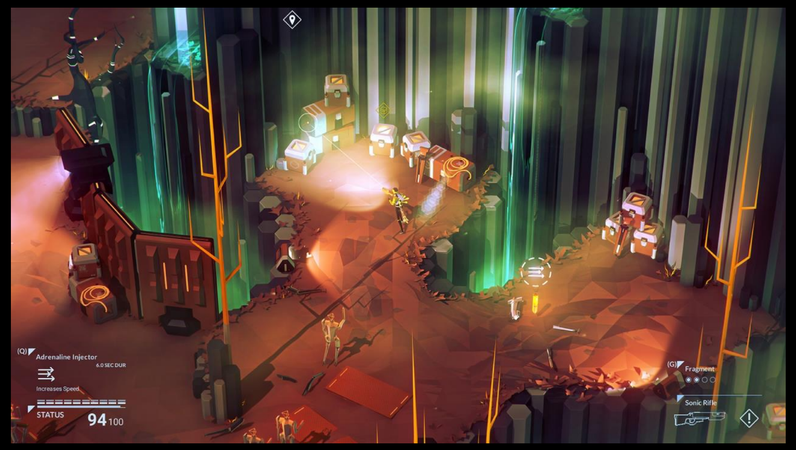

unity

Machine Learning Agents

Unity ML Agents Workflow

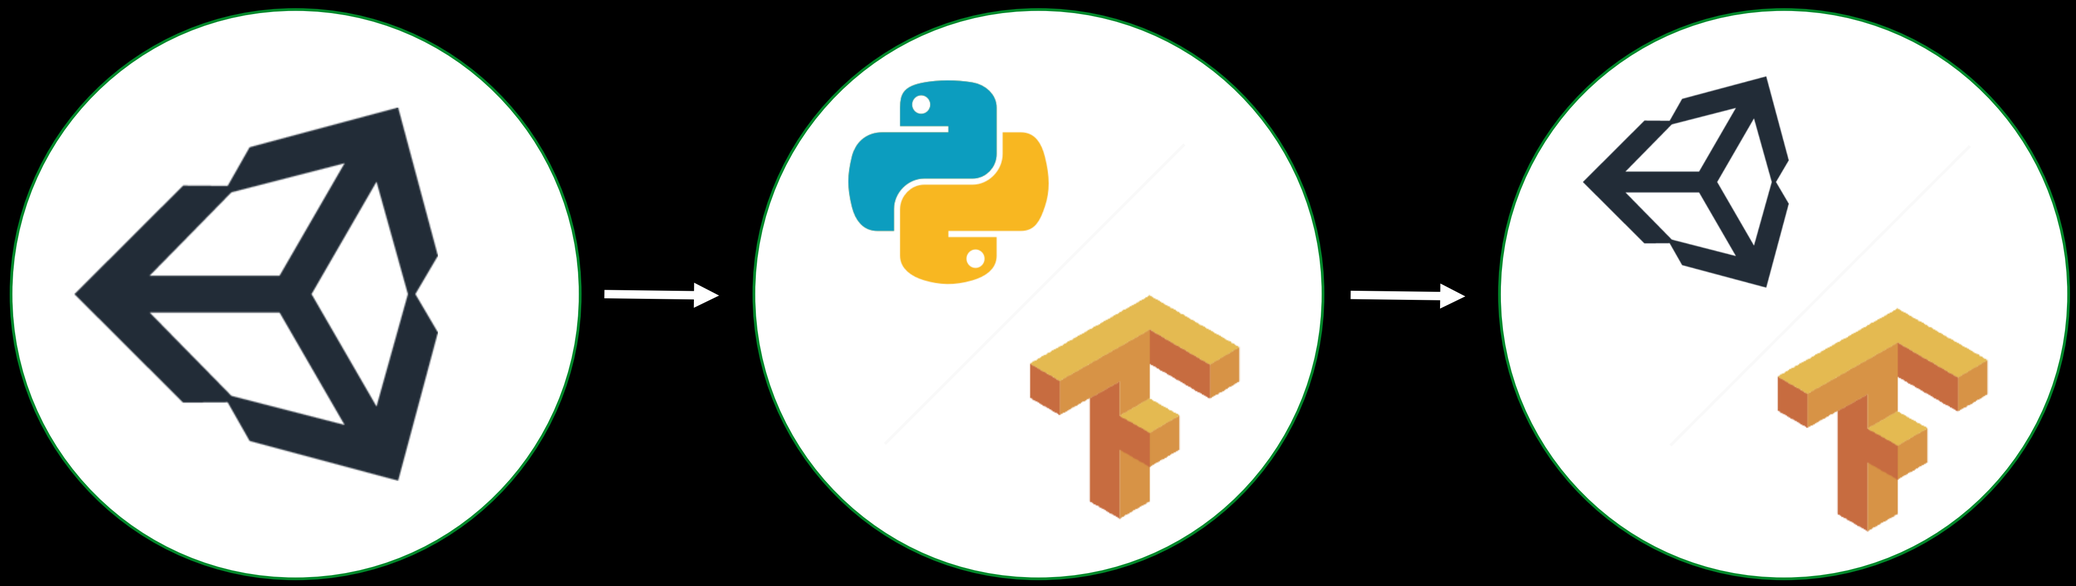

Set Up Environment

Train Agents

Embed Agents

Unity ML Agents Workflow

Set Up Environment

Train Agents

Embed Agents

Create Environment (Unity)

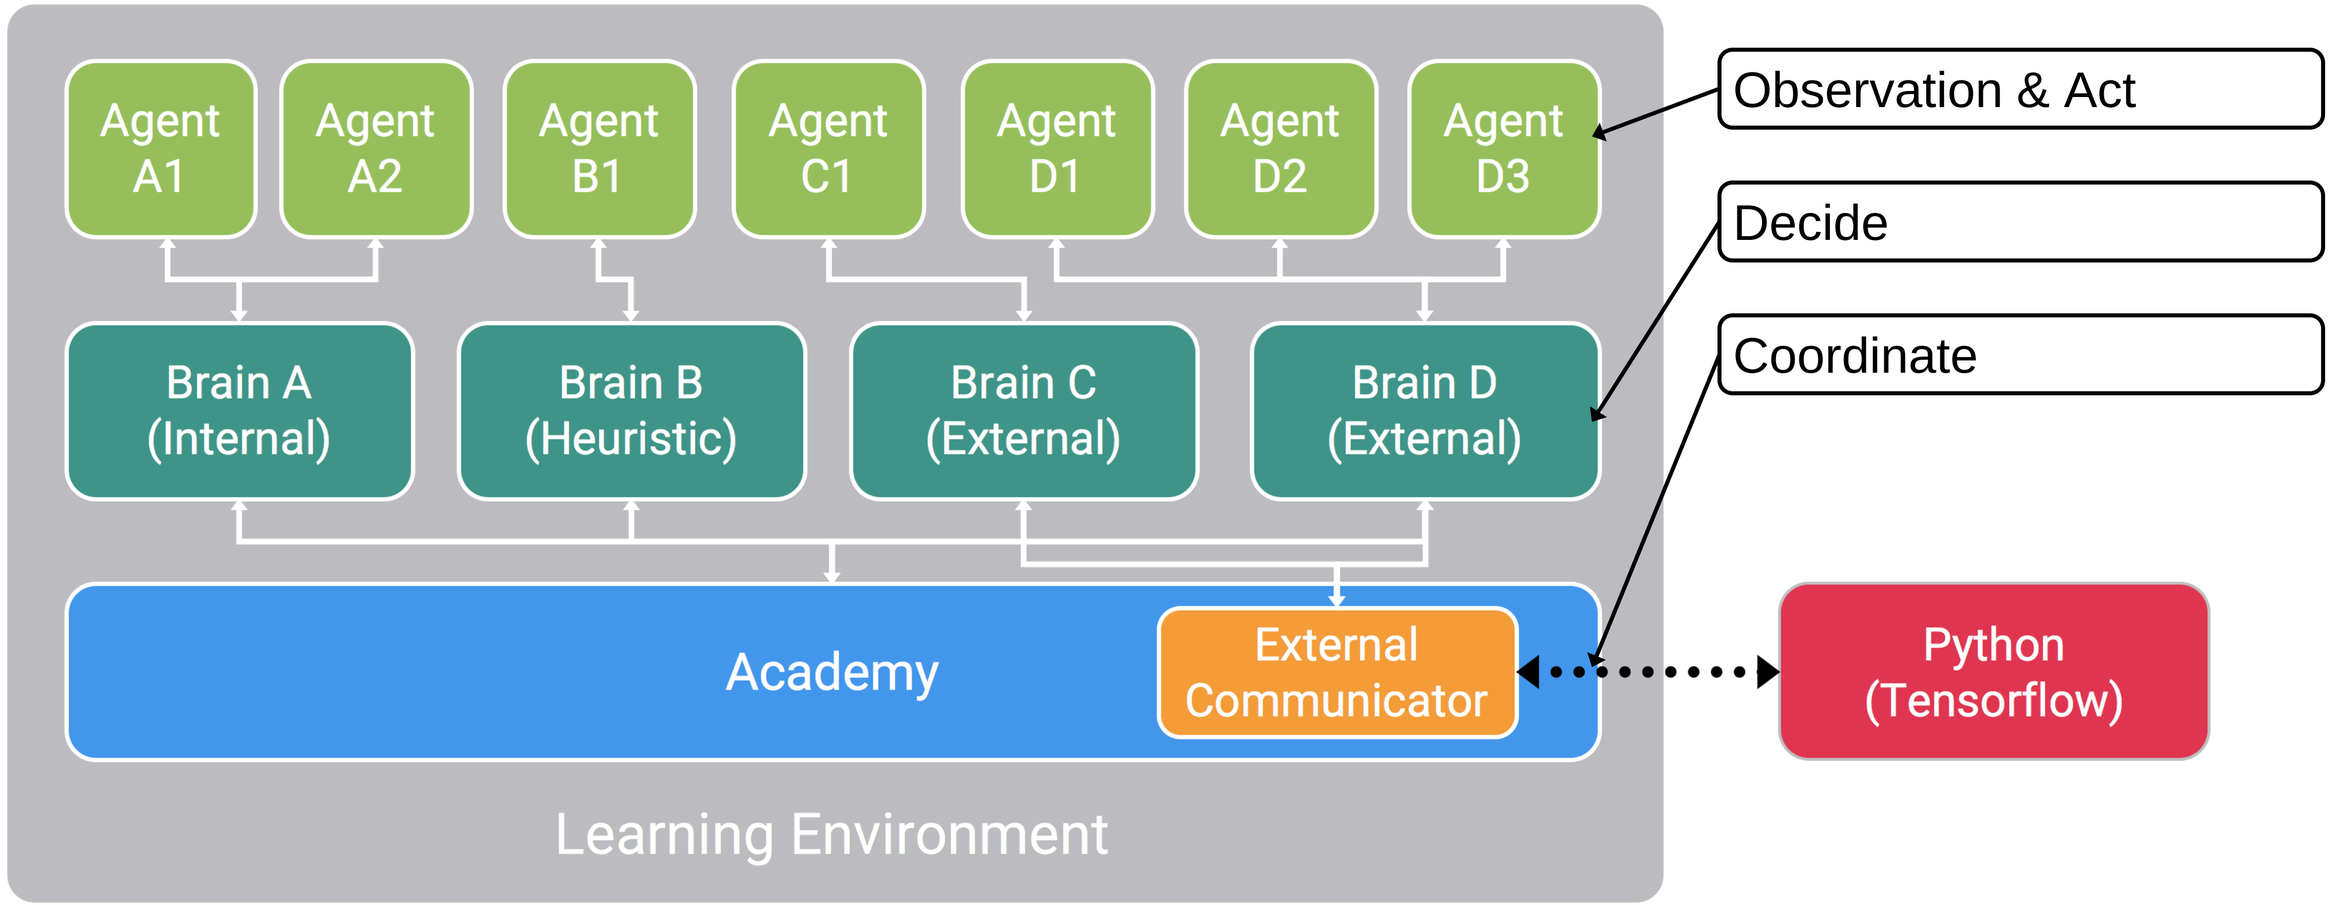

Unity ML Agents Workflow

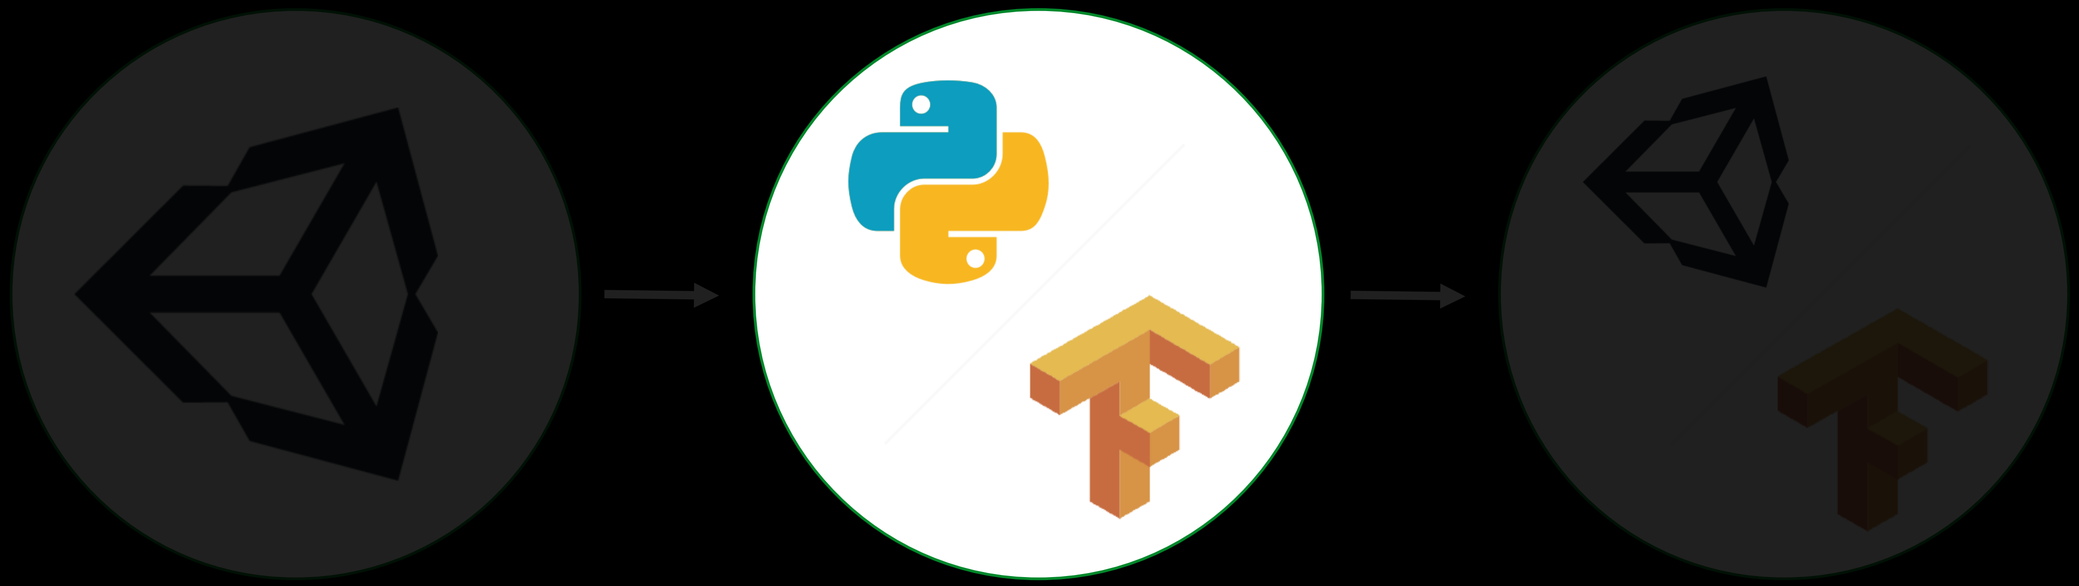

Set Up Game for Training

Train Agents

Embed Agents

Agent Training Process

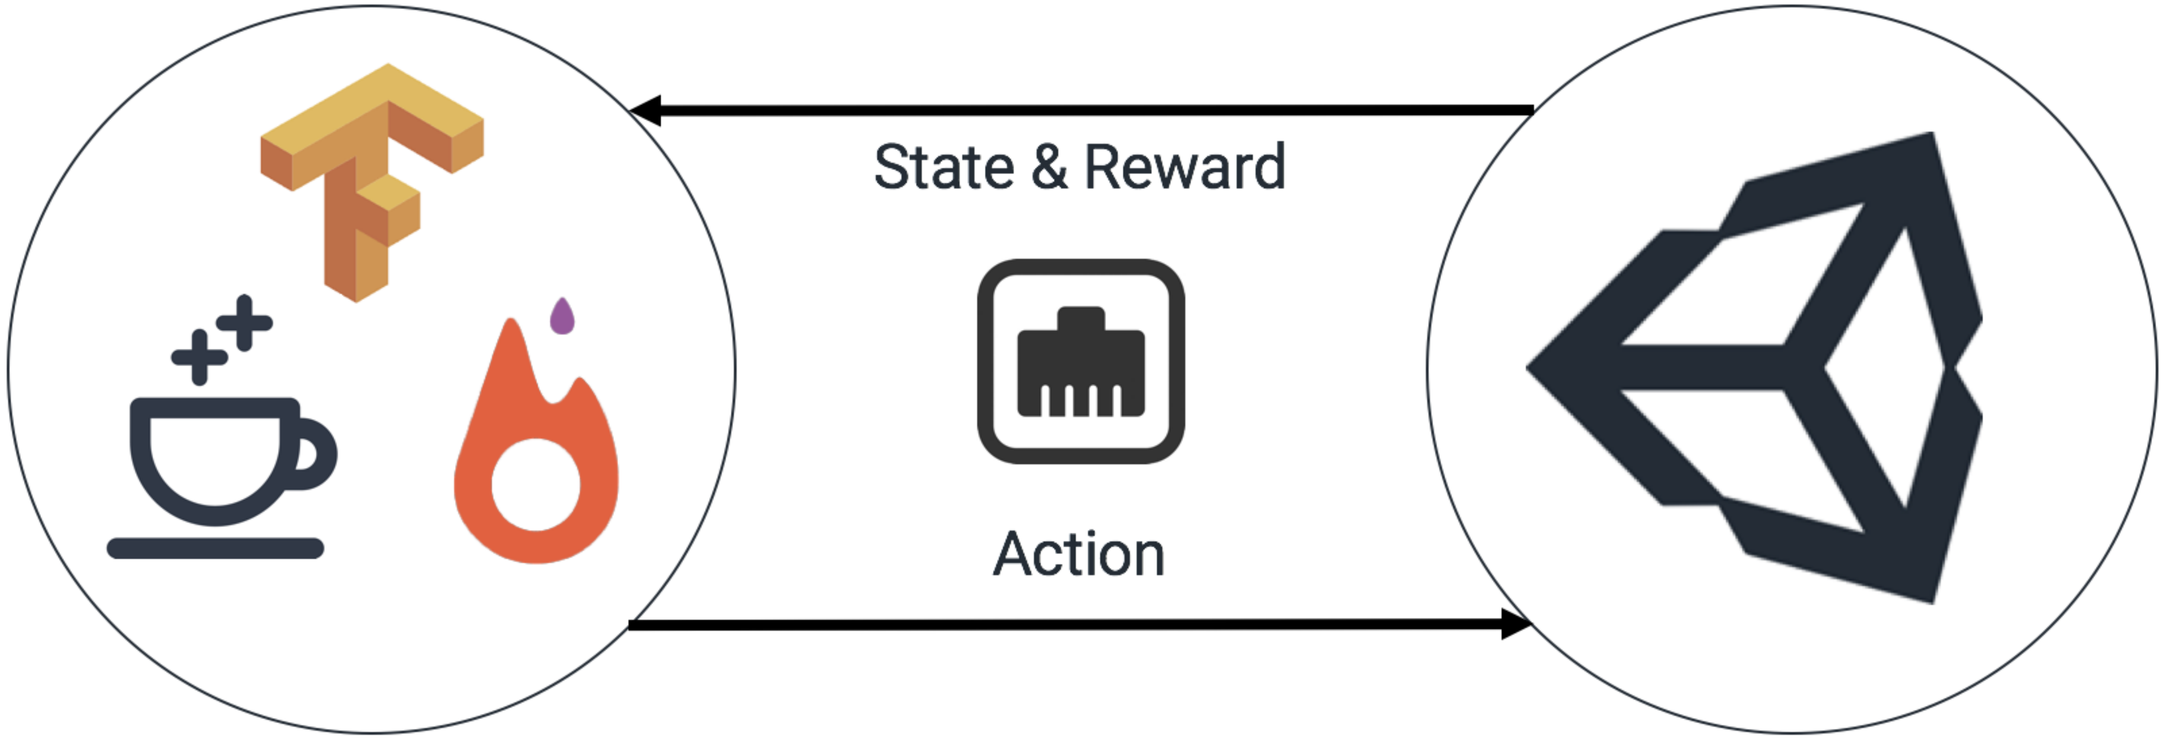

Training Methods

1. Reinforcement Learning

2. Imitation Learning

Reinforcement Learning Process

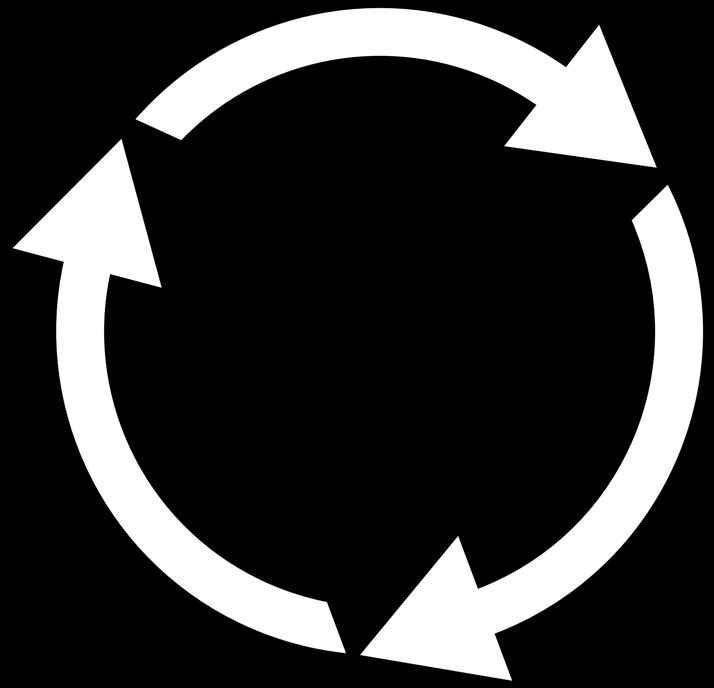

Explore → Exploit

Example: Chicken Crossing the Road

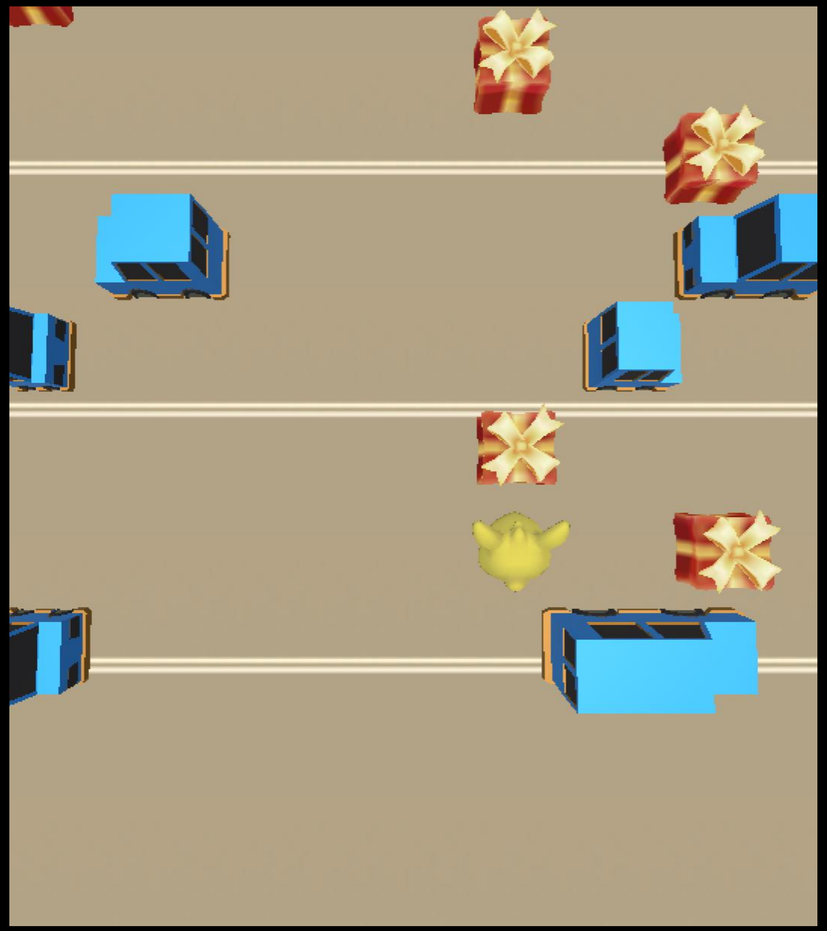

- Observe: Pixels in frame
- Actions: 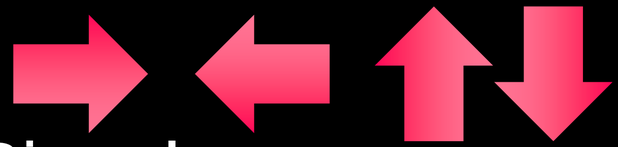
- Reward Signal
  - Negative for being hit
  - Positive for gift pickup

Imitation Learning Process

Demonstrate to the machine how it’s done

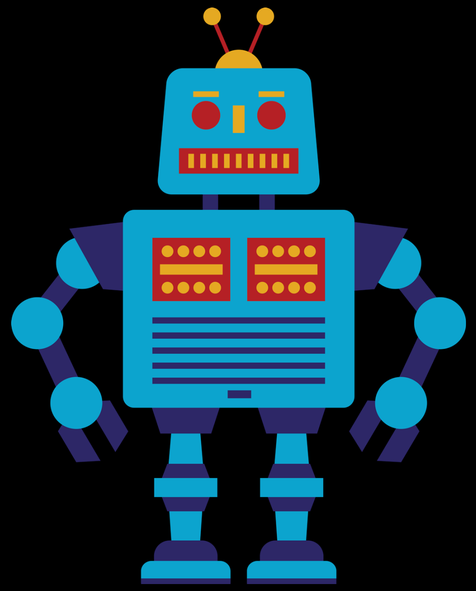

Policy is created by imitating the human

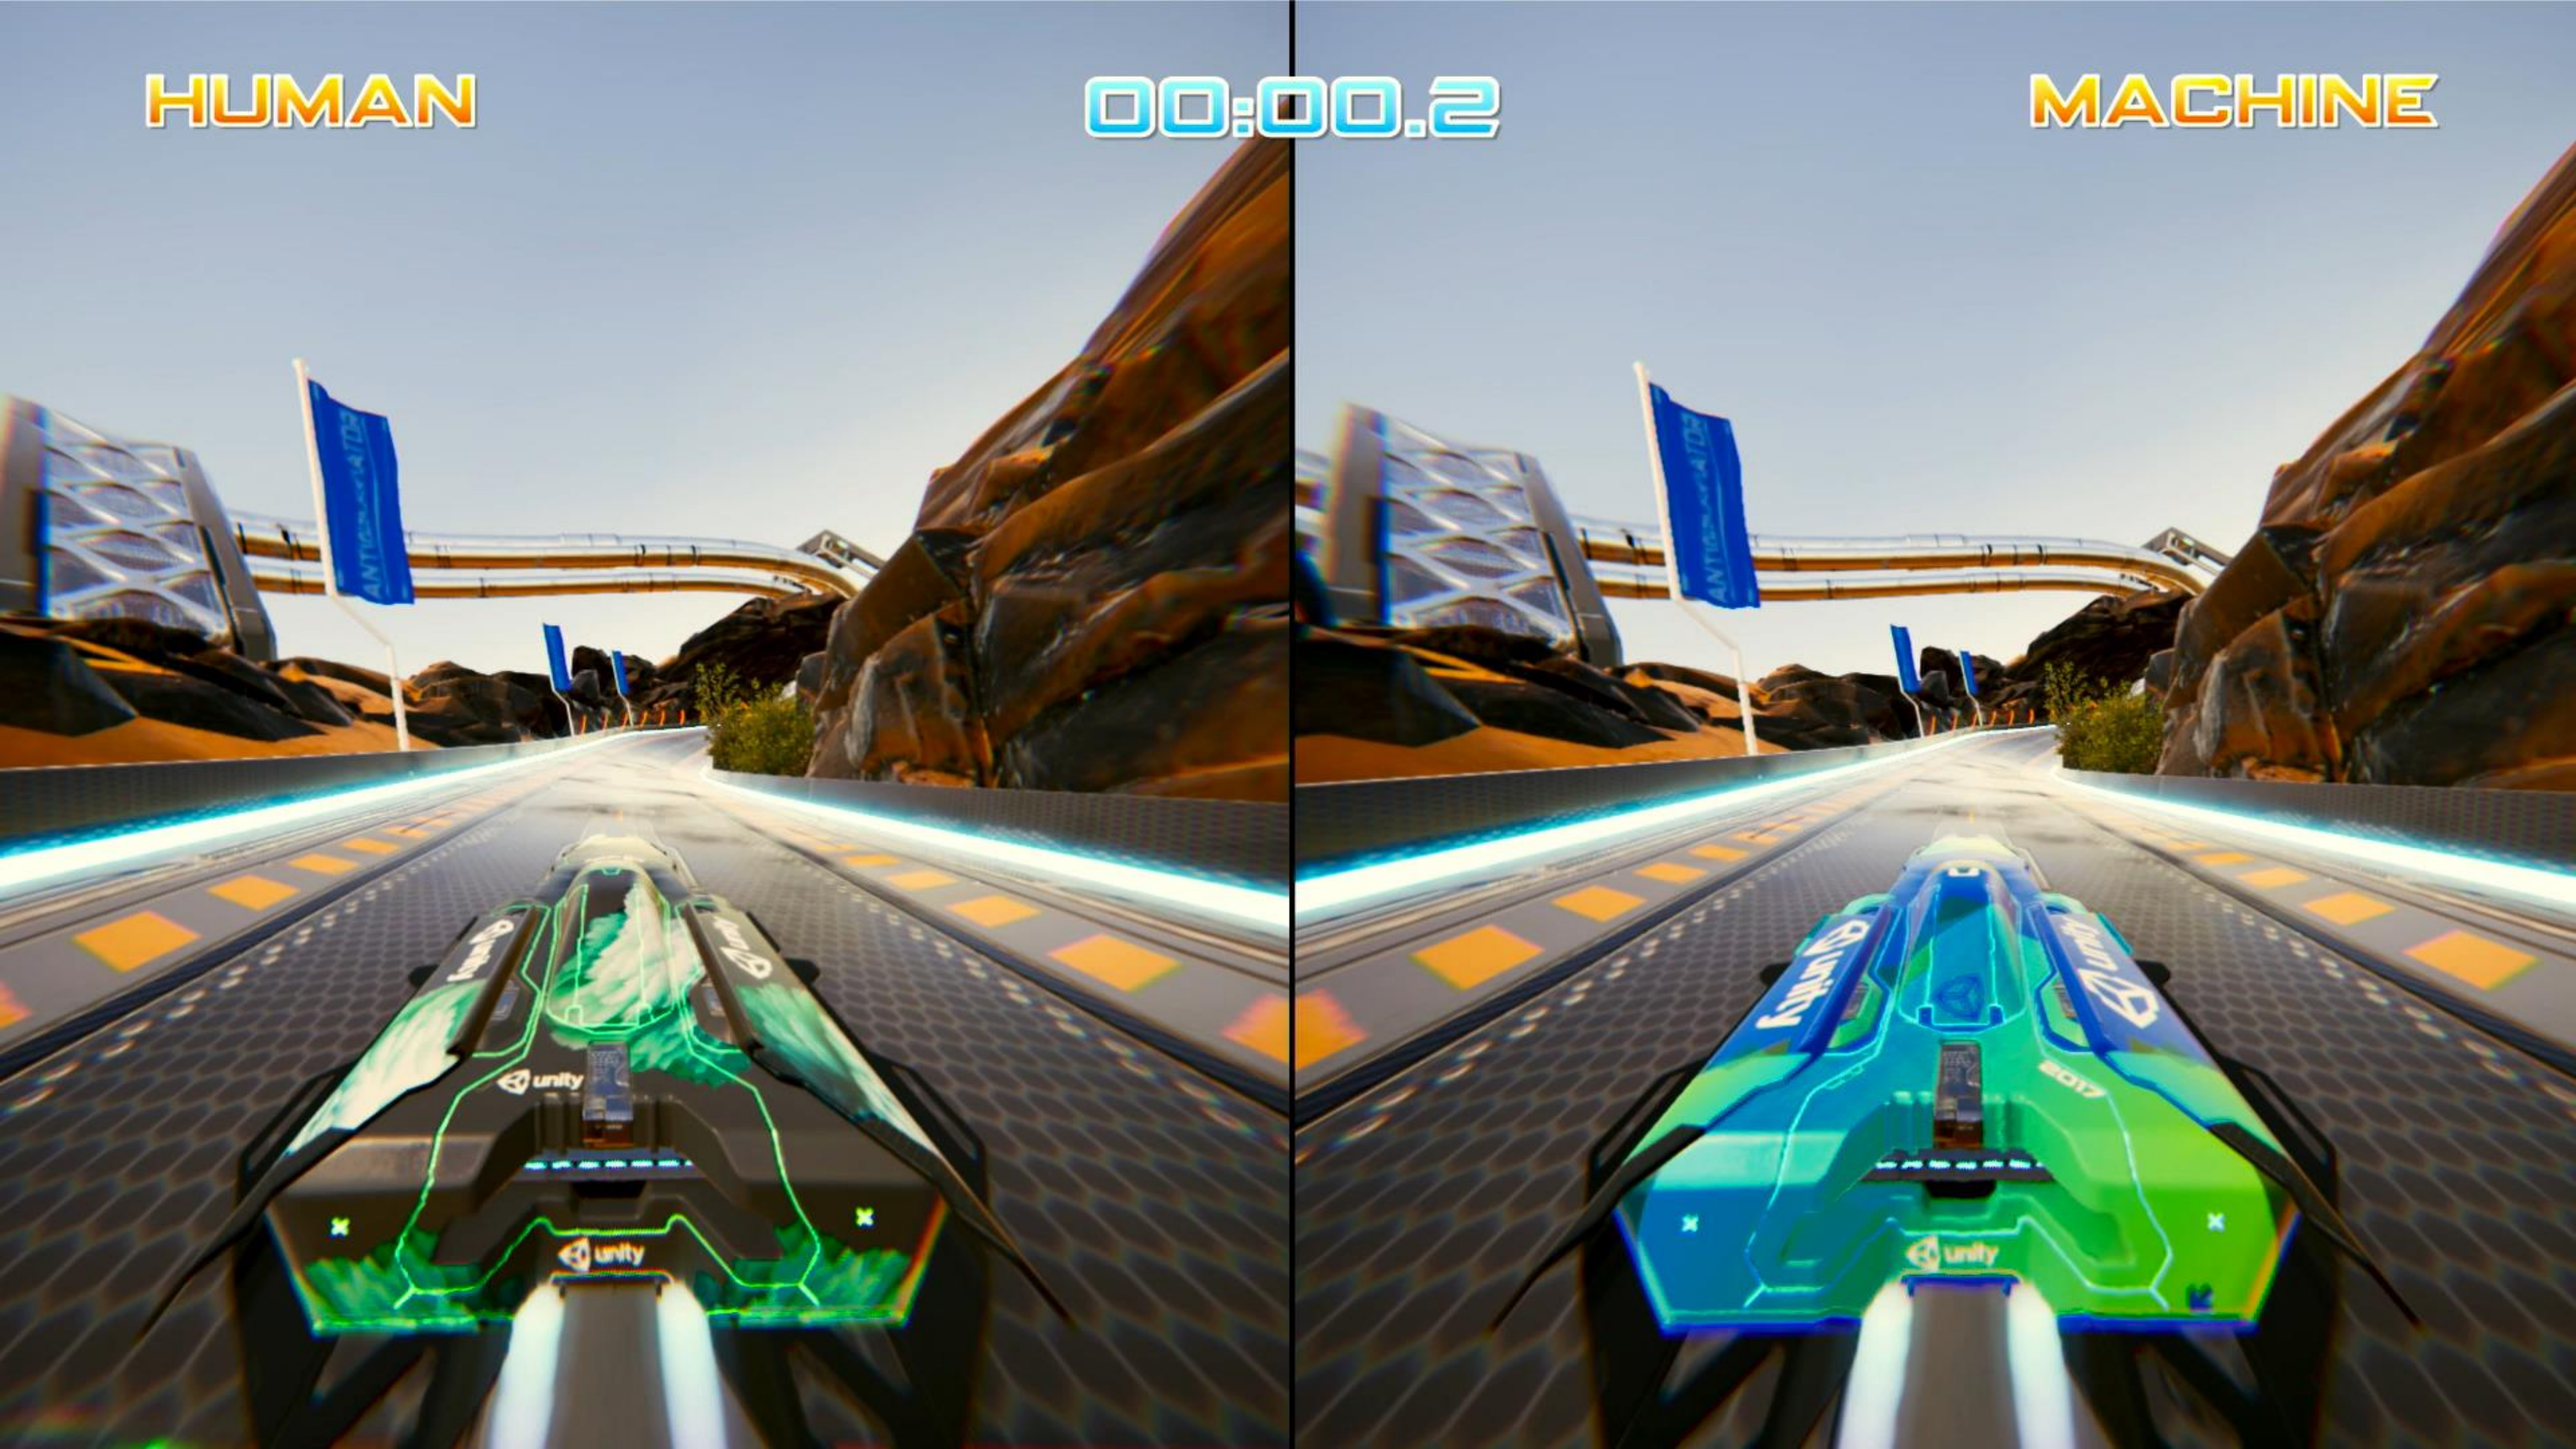
HUMAN
00:00.2
MACHINE

Unity ML Agents Workflow

Set Up Game for Training

Train Agents

Embed Agents

Embedding agents into a game (experimental)

- Import *.bytes* file into the Unity project (this the model file that is created from training agents)
- Set corresponding brain component to “Internal”
- The agent will run in the game or scene based on model created
- Inferencing is a very challenging, industry wide problem

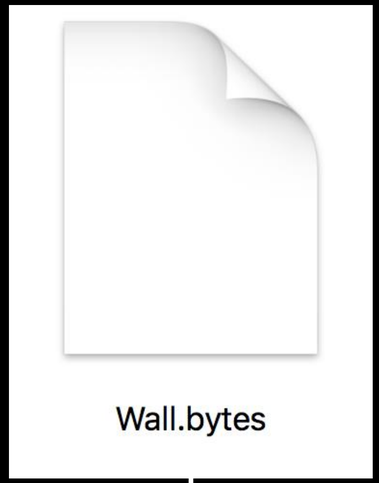

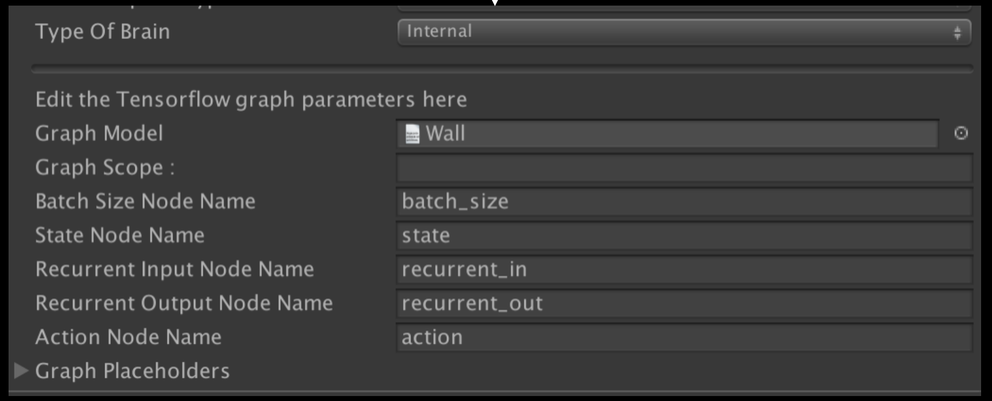

We are collaborating with different industry leaders

Microsoft WinML – Windows devices

Google Tensforflow Lite – Android devices

Apple CoreML – Apple iOS and OSX devices

Other platforms planned for the future – Nintendo, Sony

Task Possibilities

3D Balance Ball

Our first environment to use ML-Agents

Goal	Balance ball as long as possible
Observations	Platform rotation, ball position and rotation
Actions	Platform rotation (in x and z)
Rewards	Bonus for keeping ball up

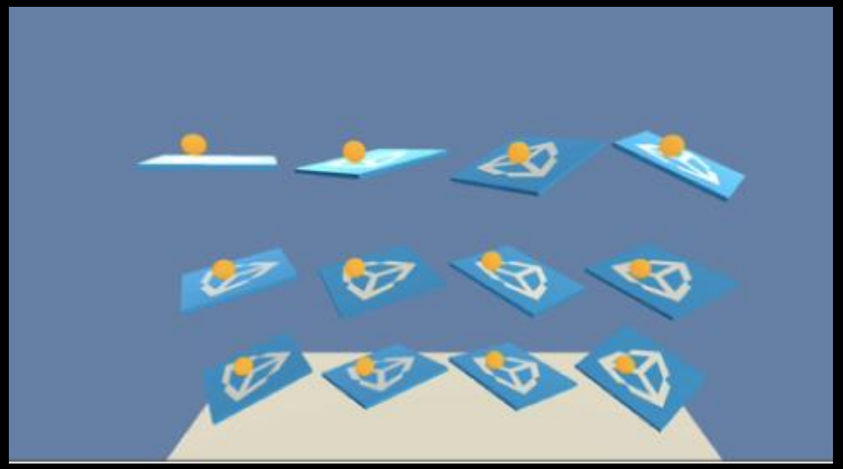

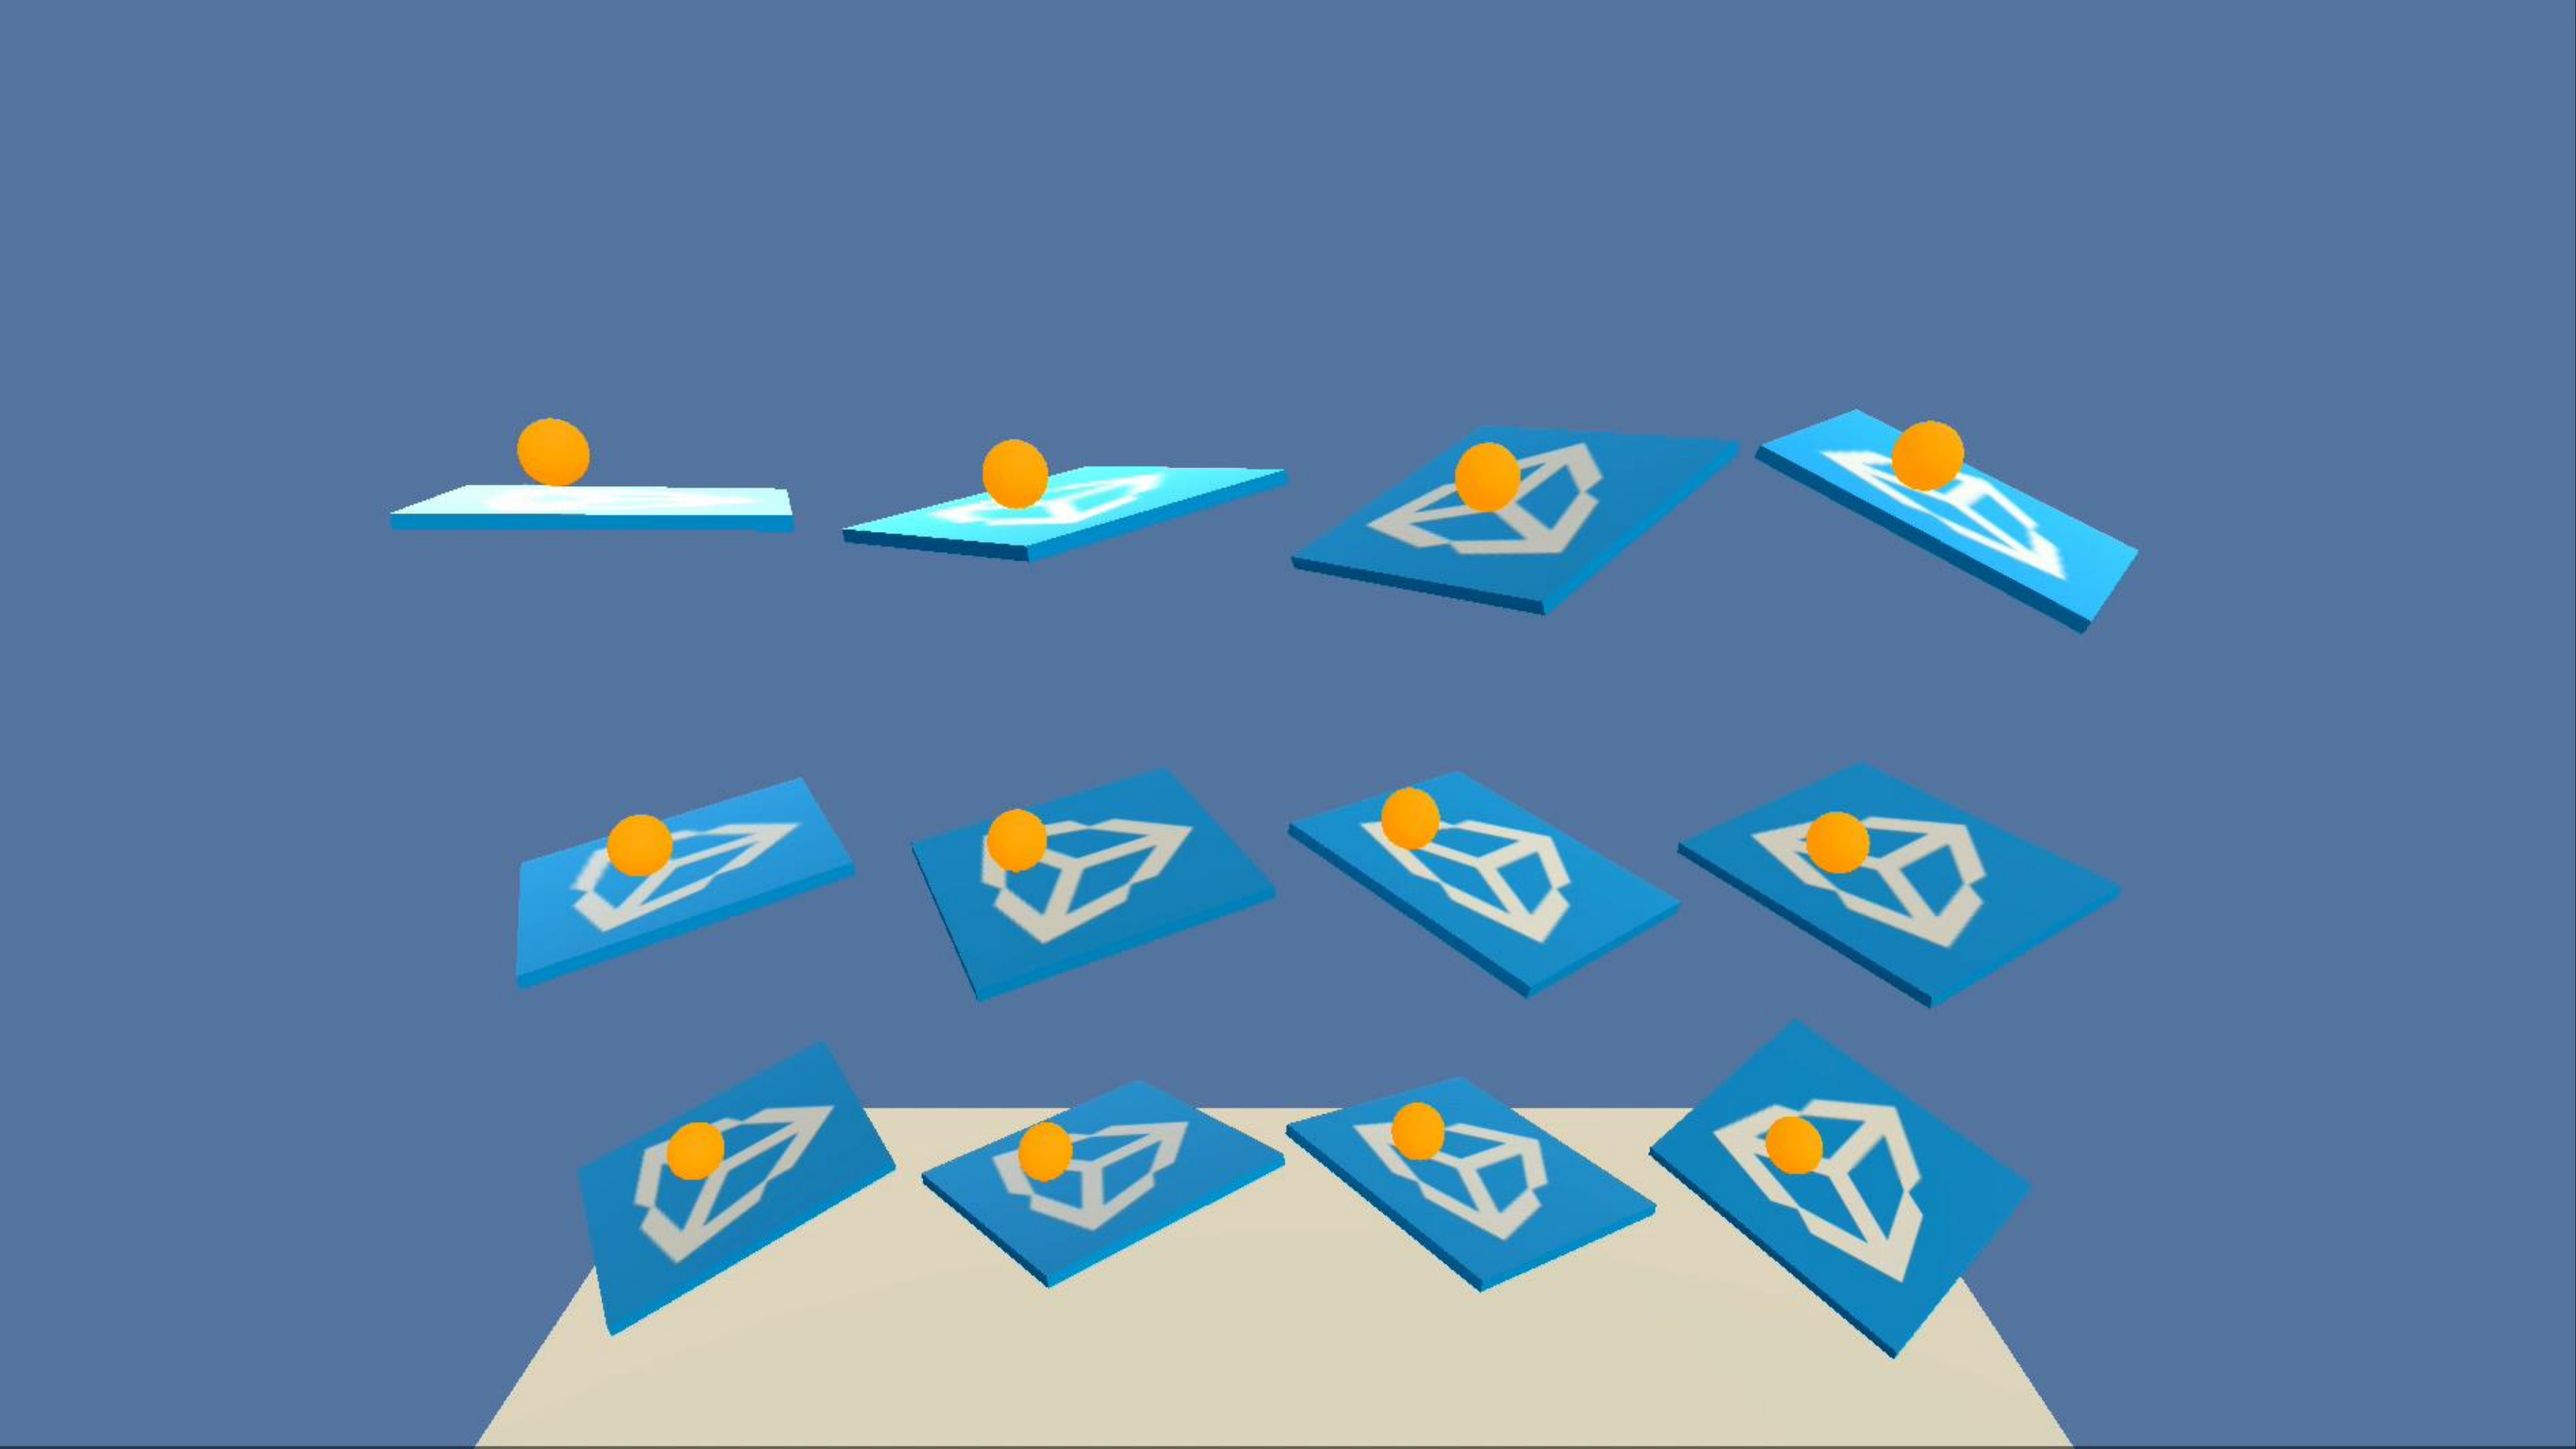

Curriculum Learning

- Agents learn from simpler exercises and combines the learning to tackle much more difficult tasks

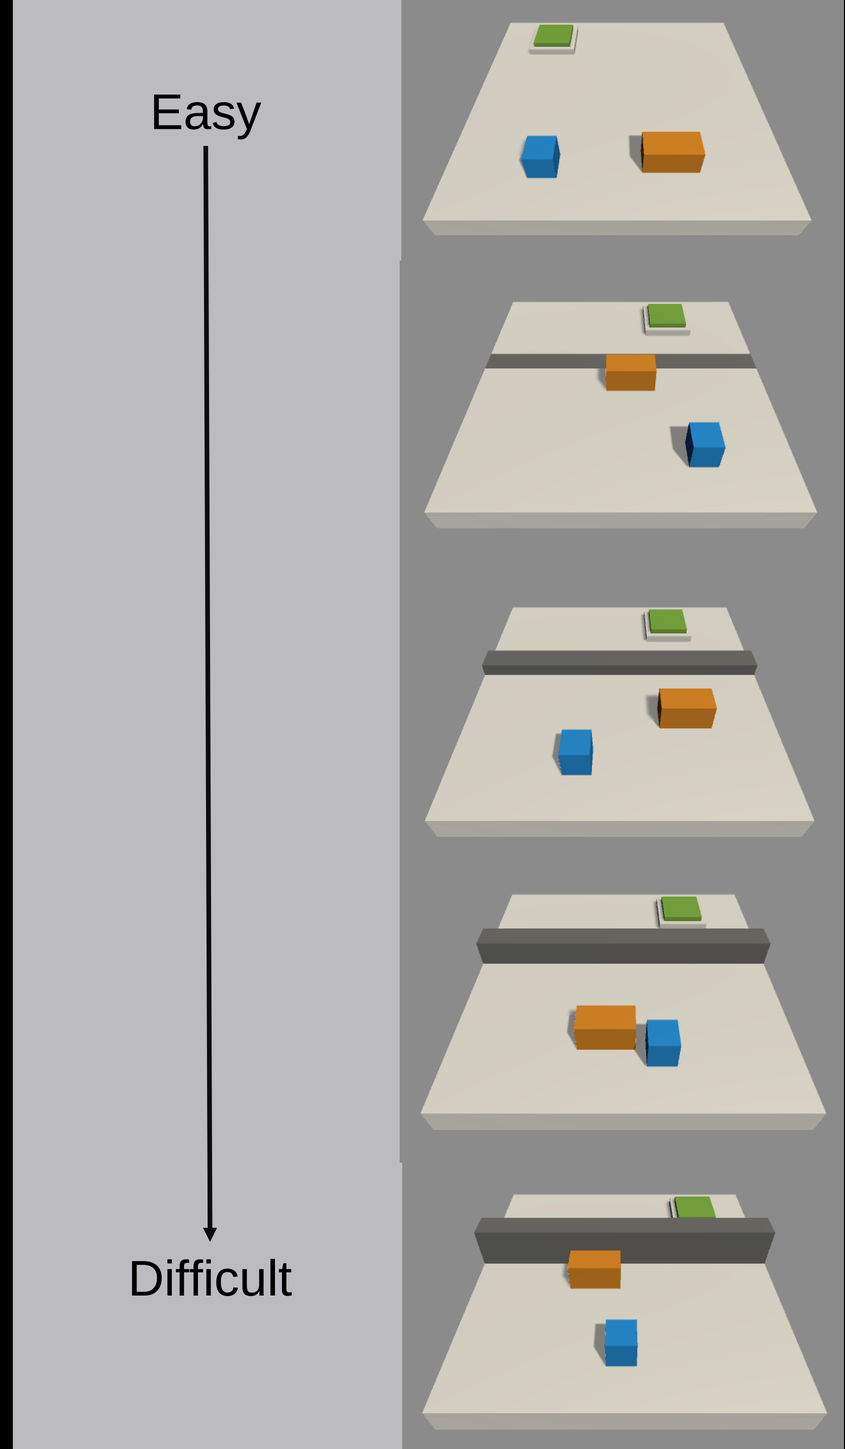

Final Outcome

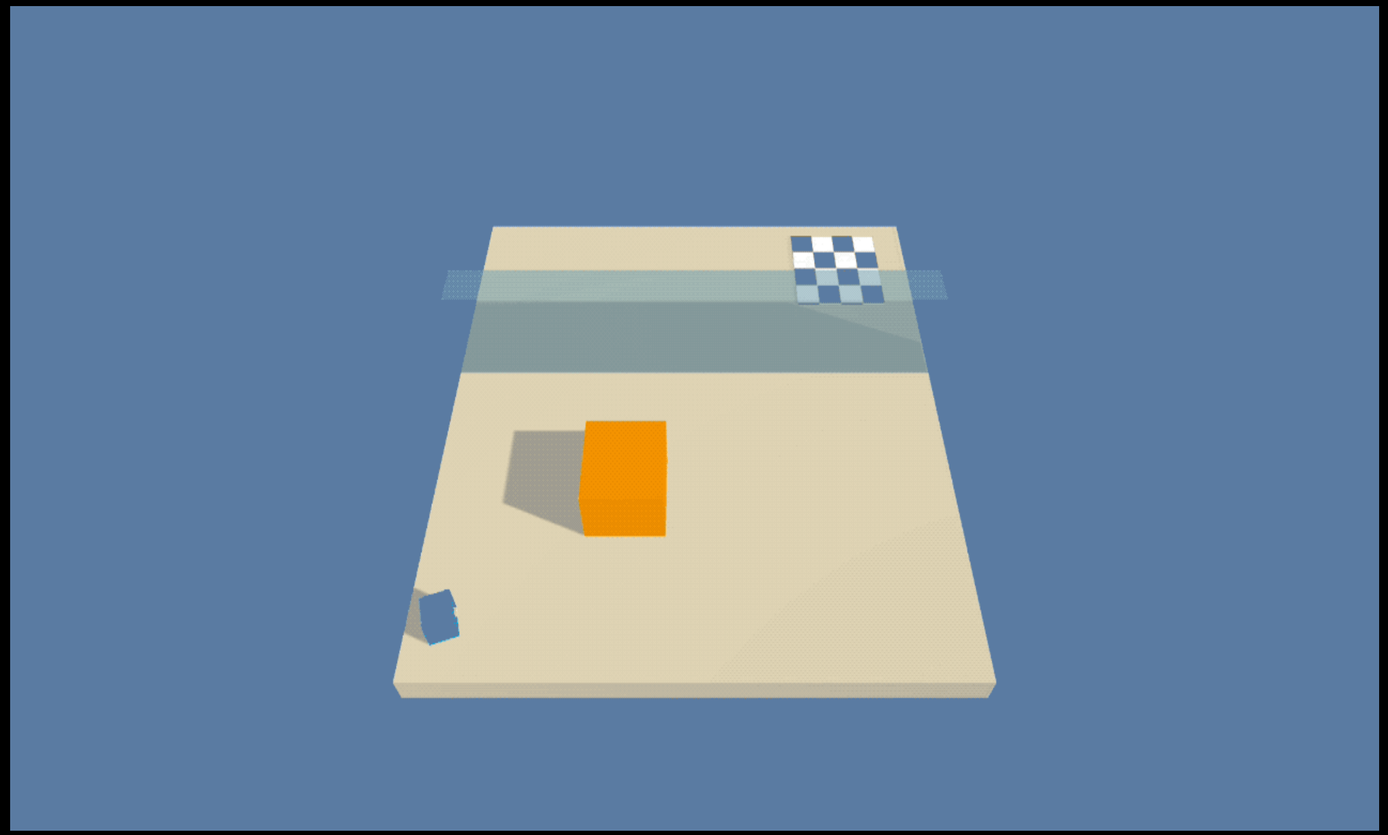

Easy

Difficult

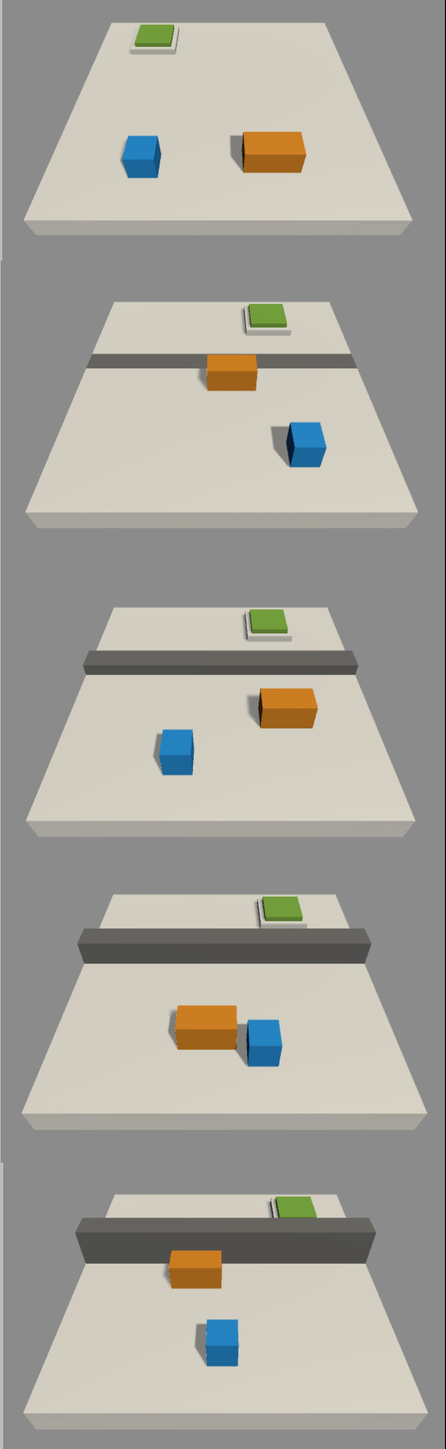

Enabling Long Short-Term Memory

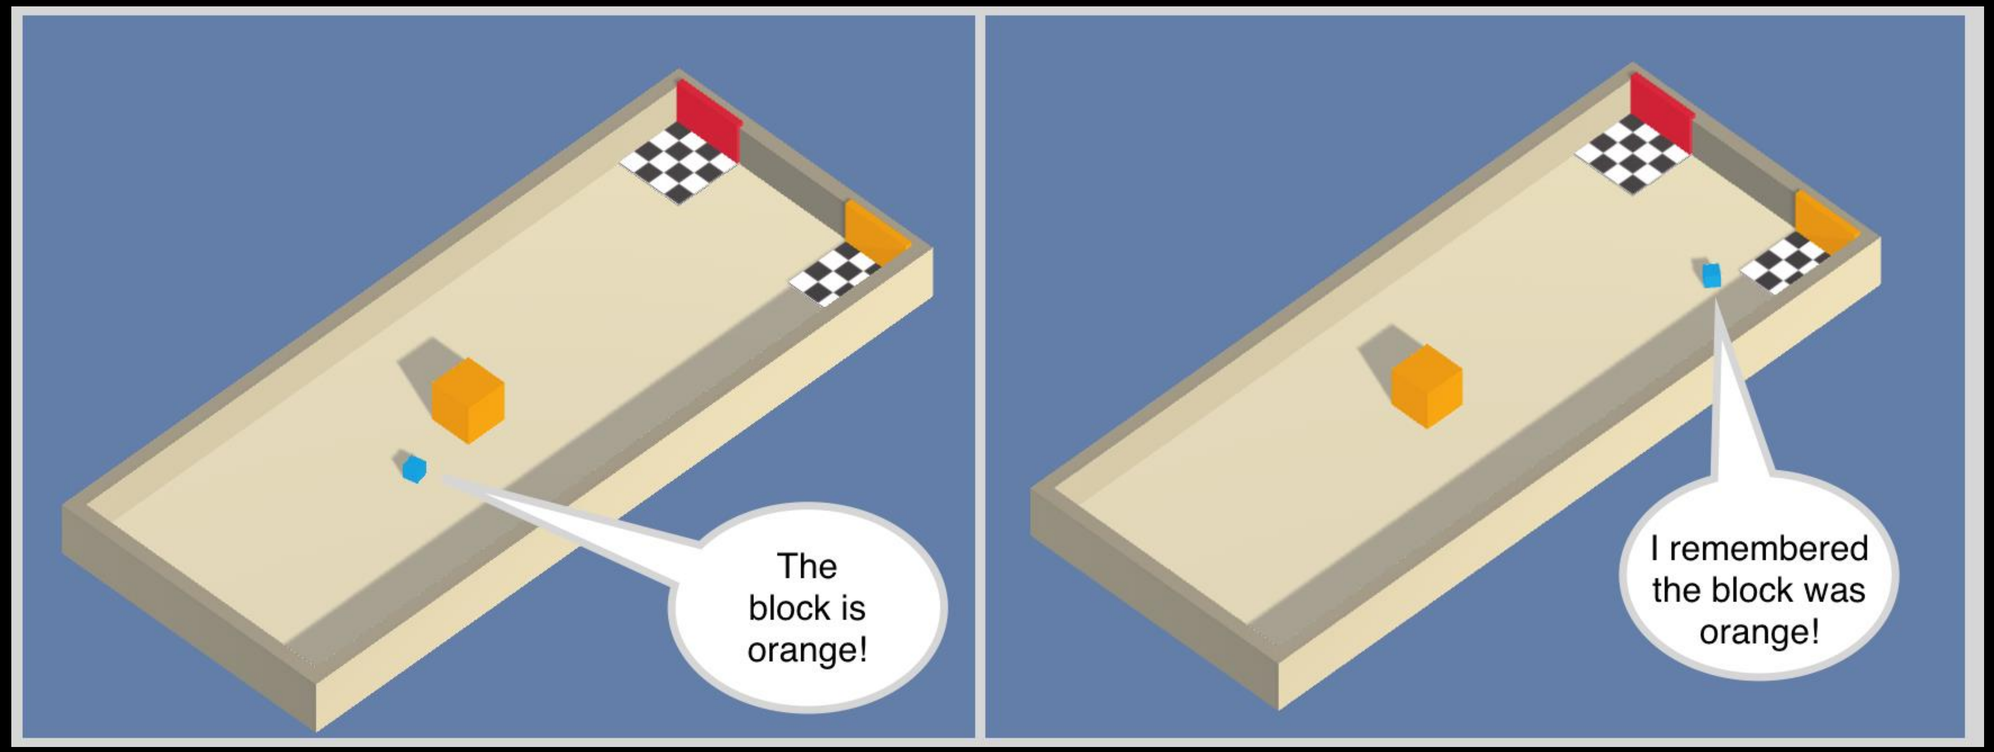

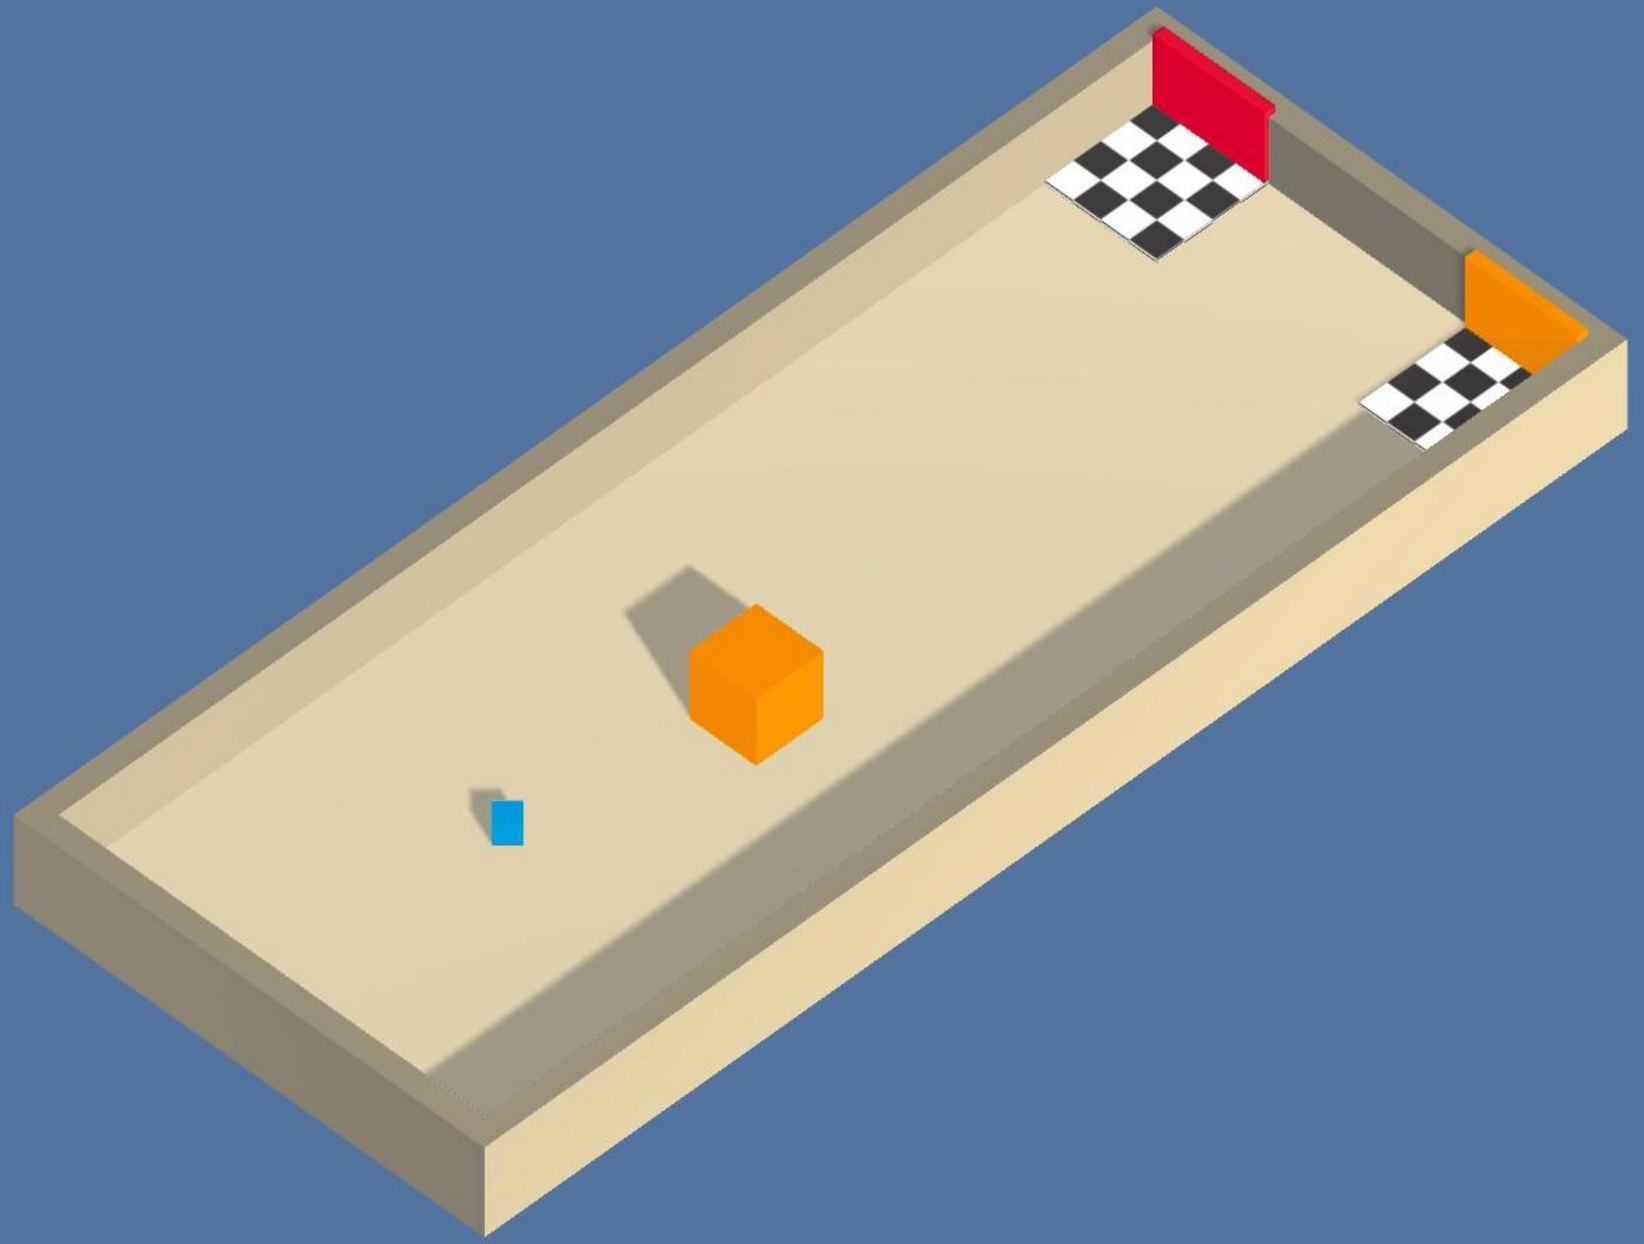

Multi-Agent Soccer Training

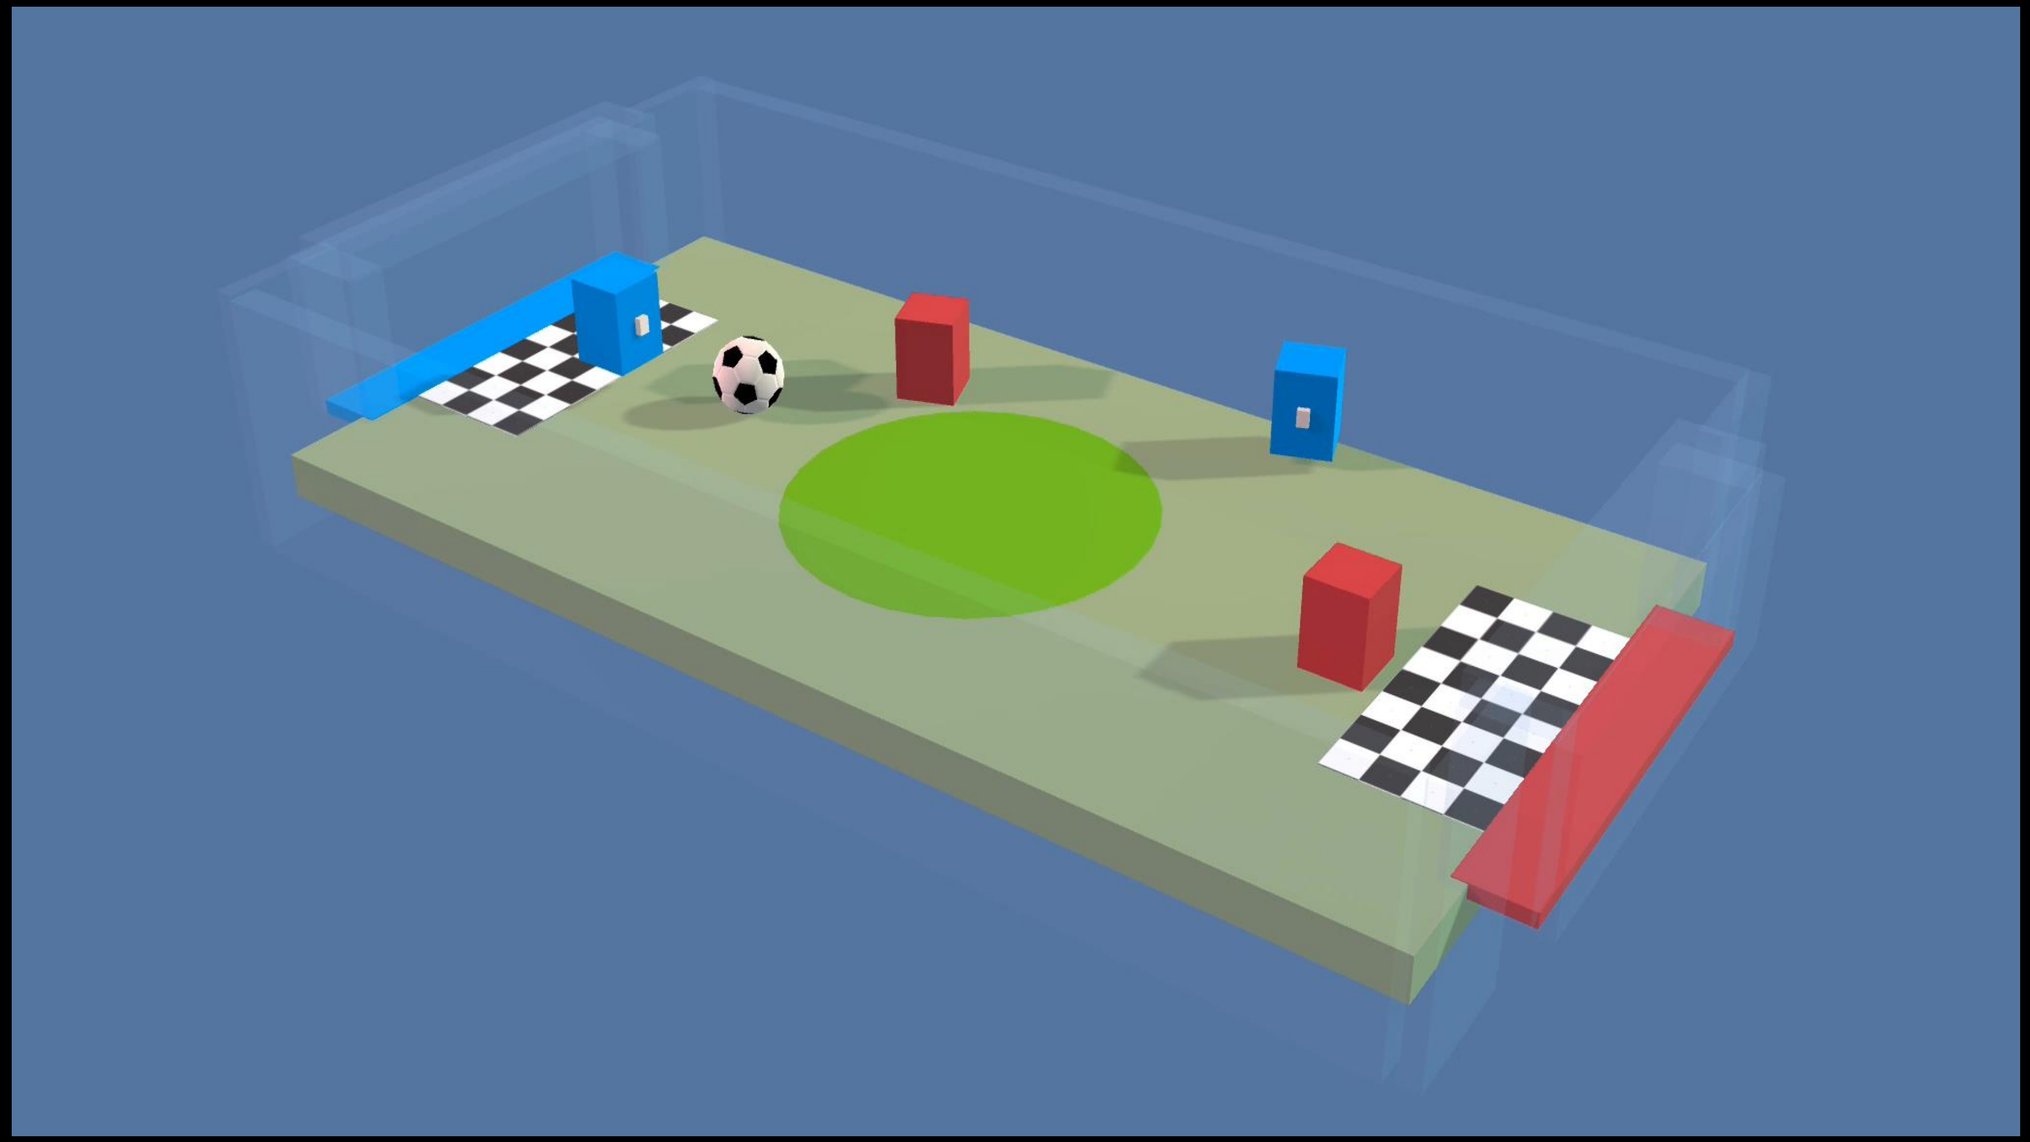

Multi-Agent Banana Collectors

Tested

Trained	Plentiful	Scarce
Plentiful	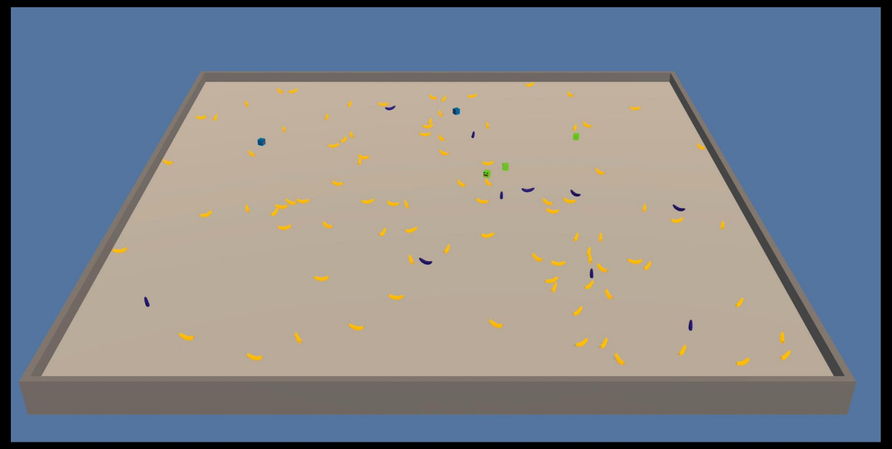	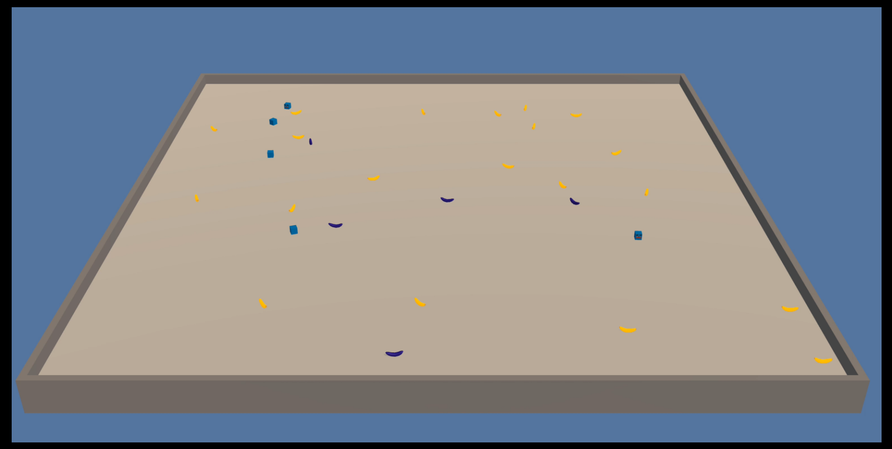
Scarce	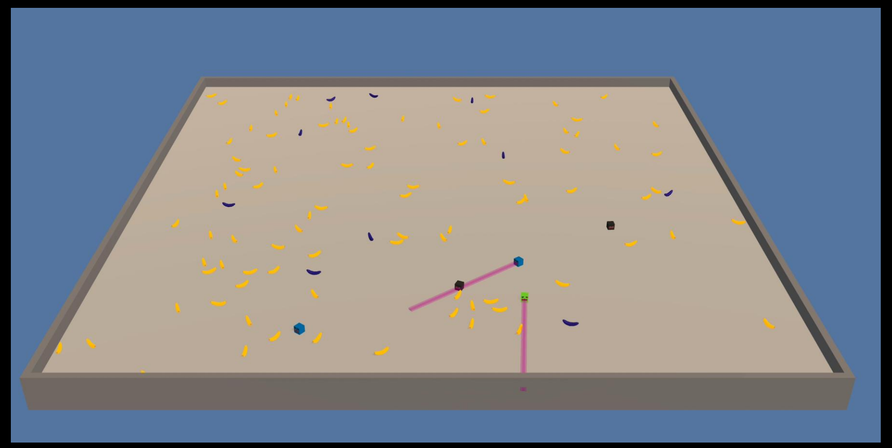	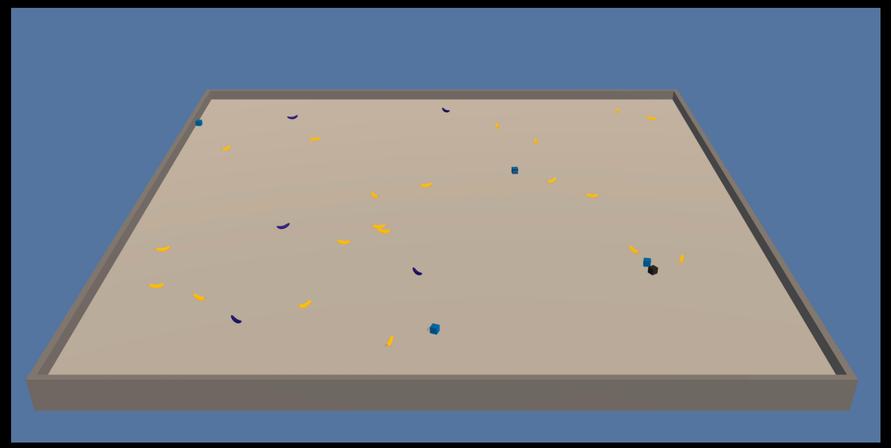

Mujoco Continuous Control Tasks in Unity

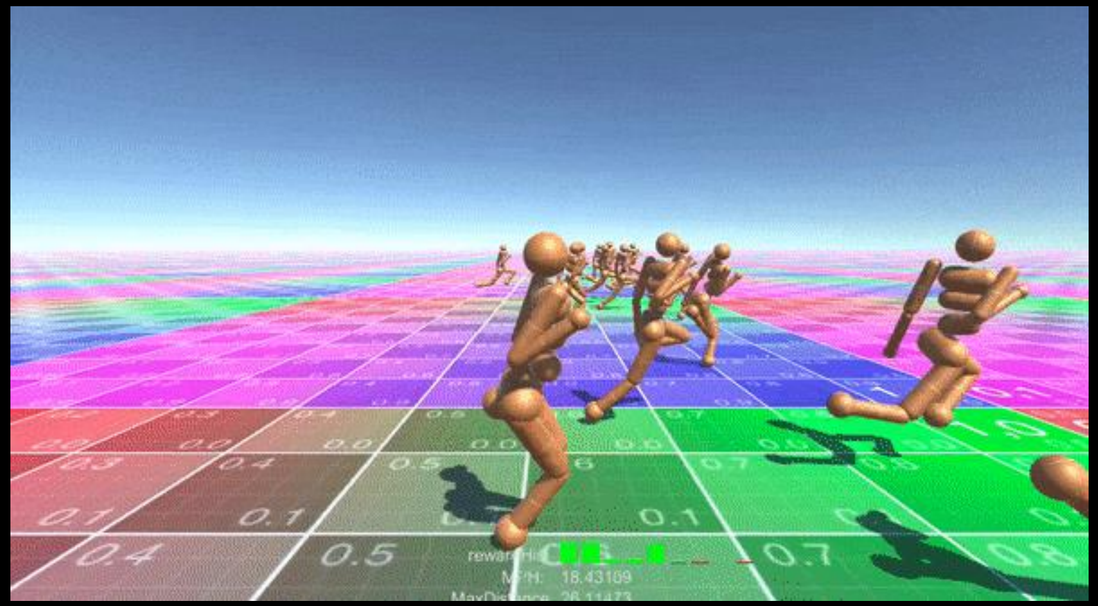

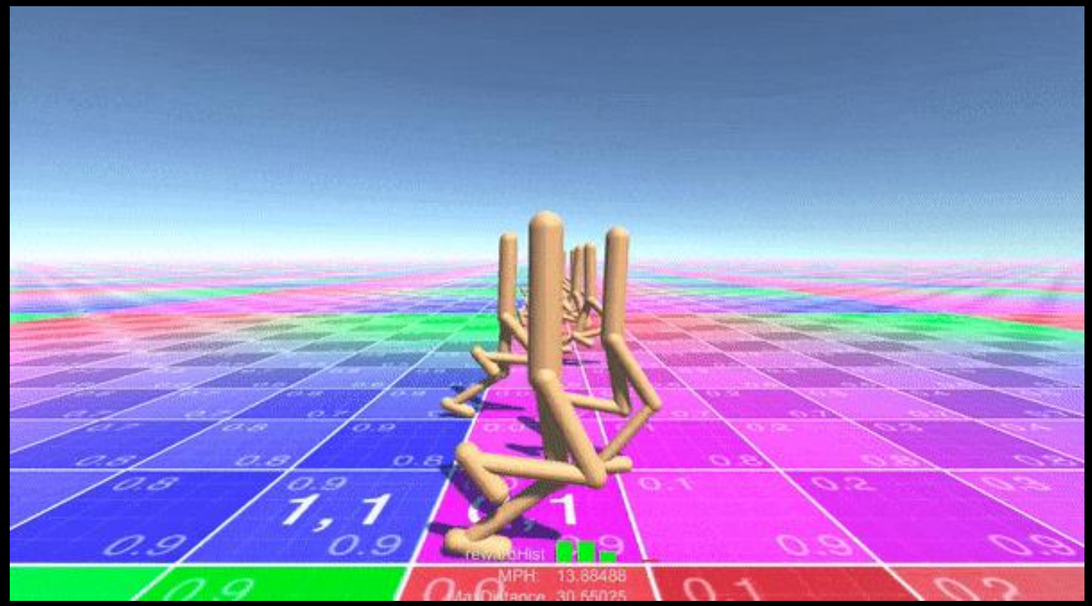

Unity Continuous Control Tasks in Unity

Curiosity-Driven Exploration

- In some environments the rewards are sparsely distributed
- “+1 for accomplishing goal”
- Intrinsic reward can encourage exploration
- Reward agent based on experienced surprise in outcome of actions

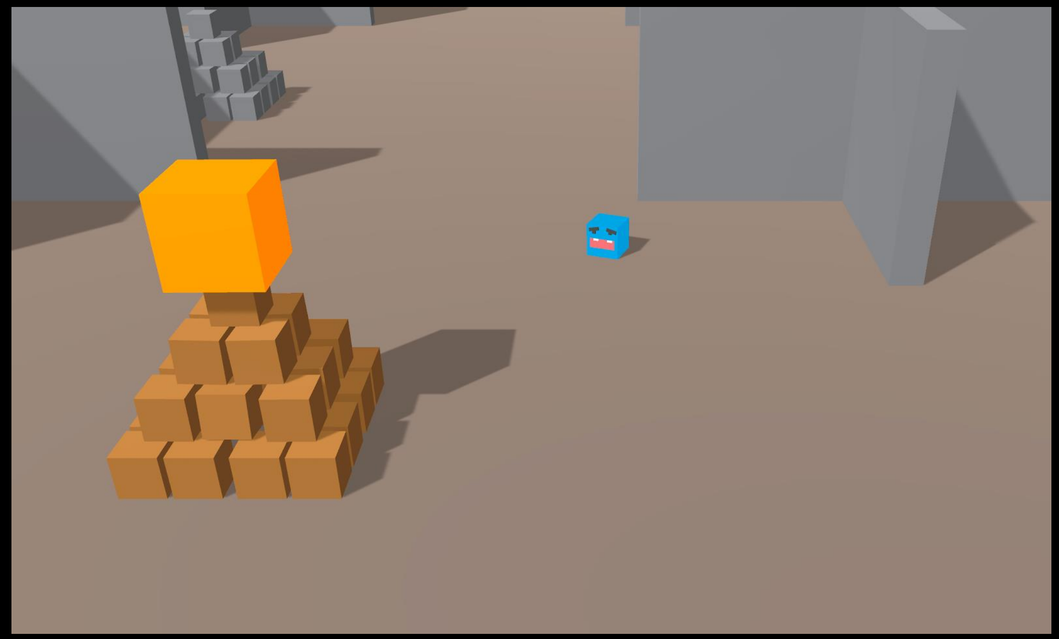

Implementation of: “**Curiosity-driven Exploration by Self-supervised Prediction**“
Pathak et al., 2017

Pyramids Environment

- Nine rooms
- One switch
- Six stone pyramids
- Once switch pressed, brick pyramid spawned
- Gold brick on top of brick pyramid provides +2 reward.

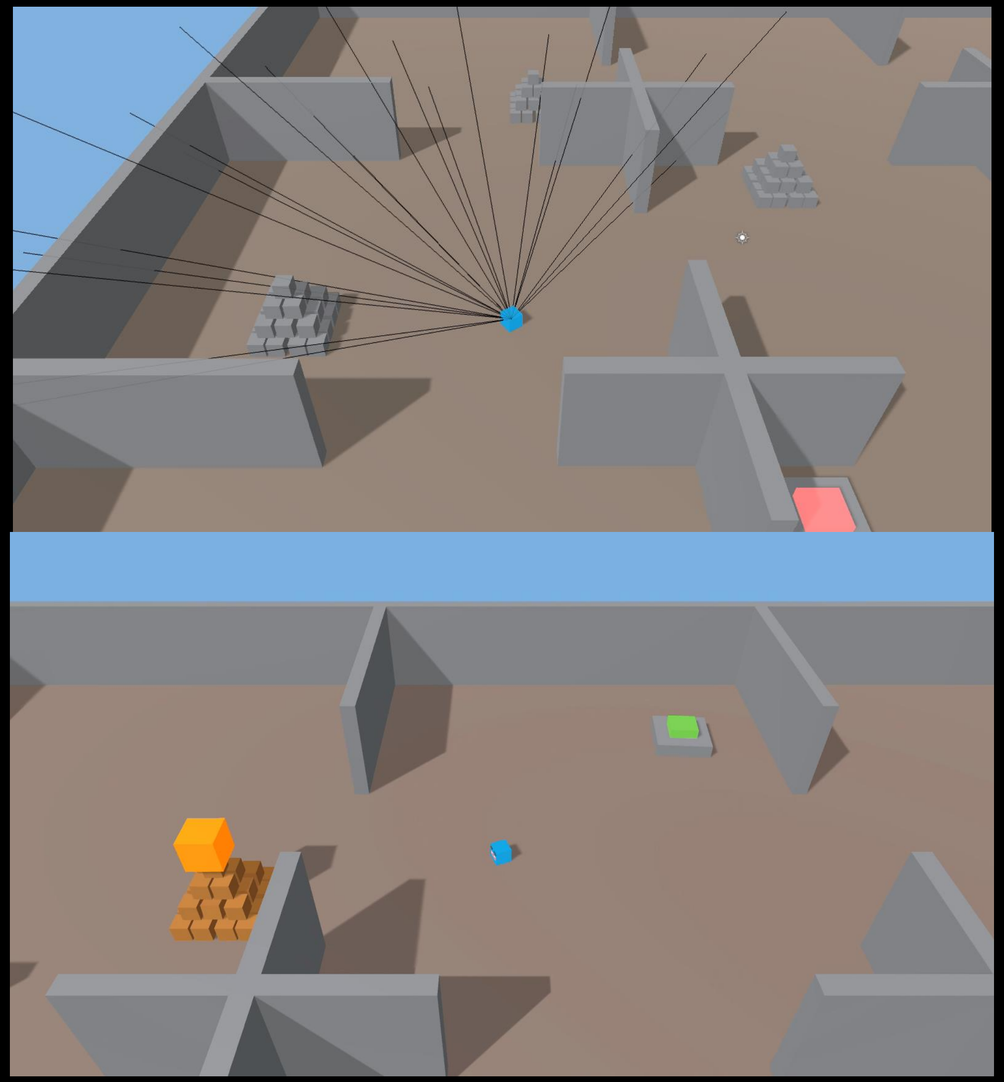

Intrinsic Curiosity Module

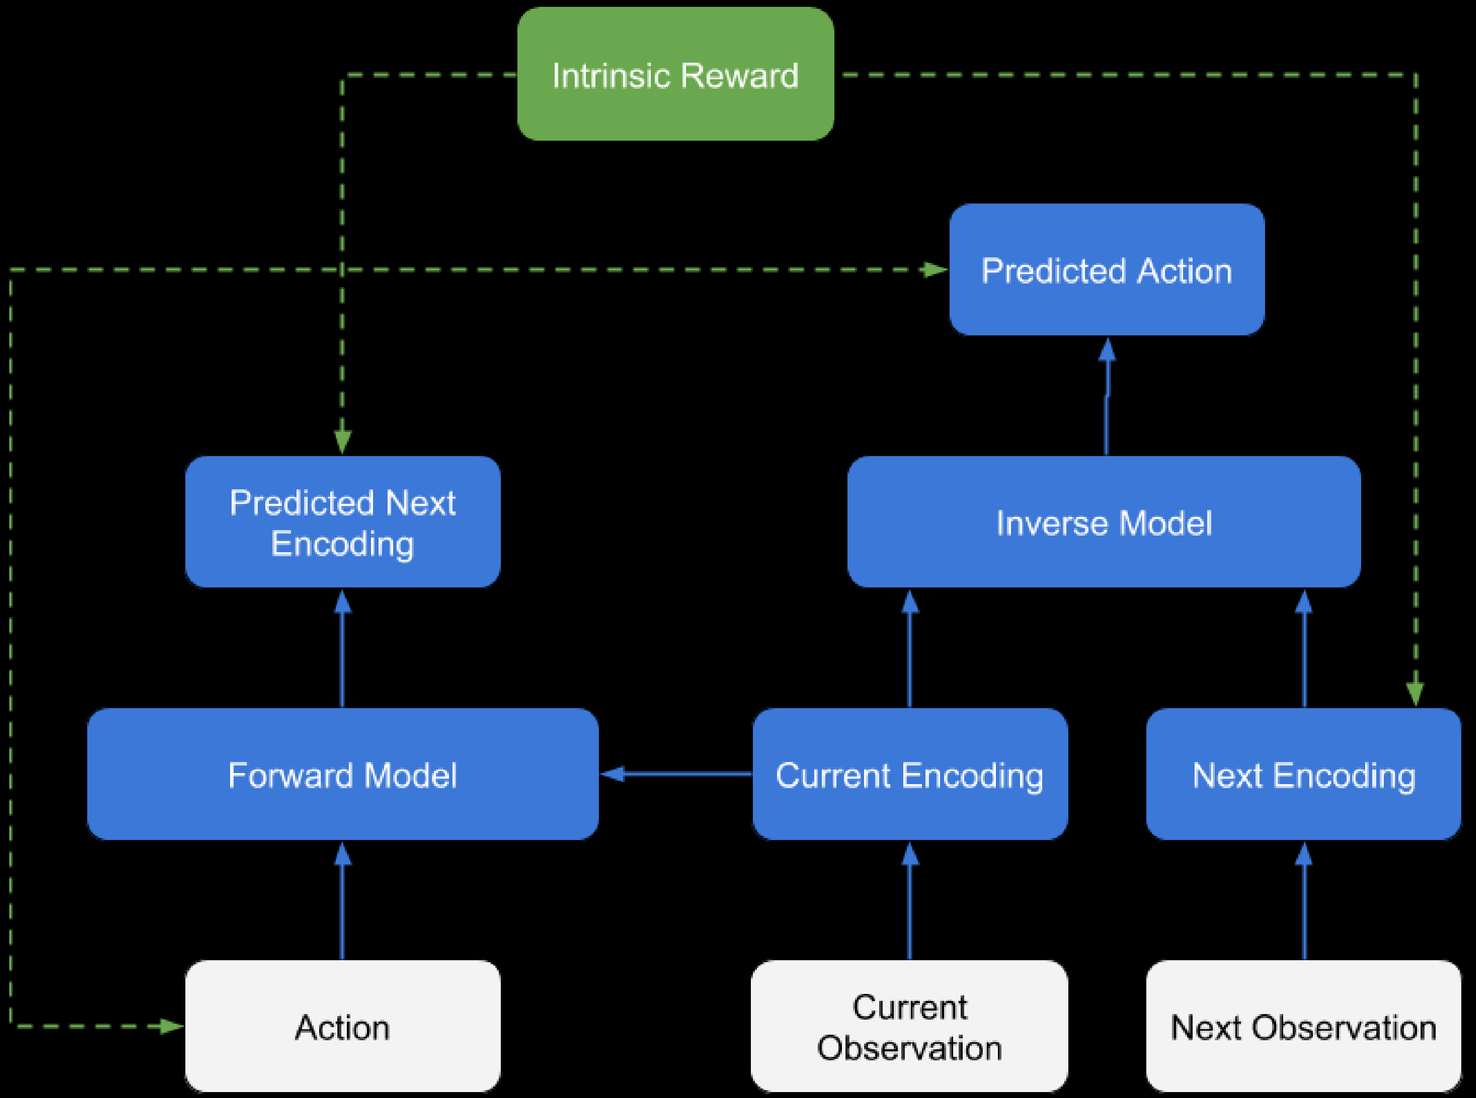

Extrinsic Reward Only

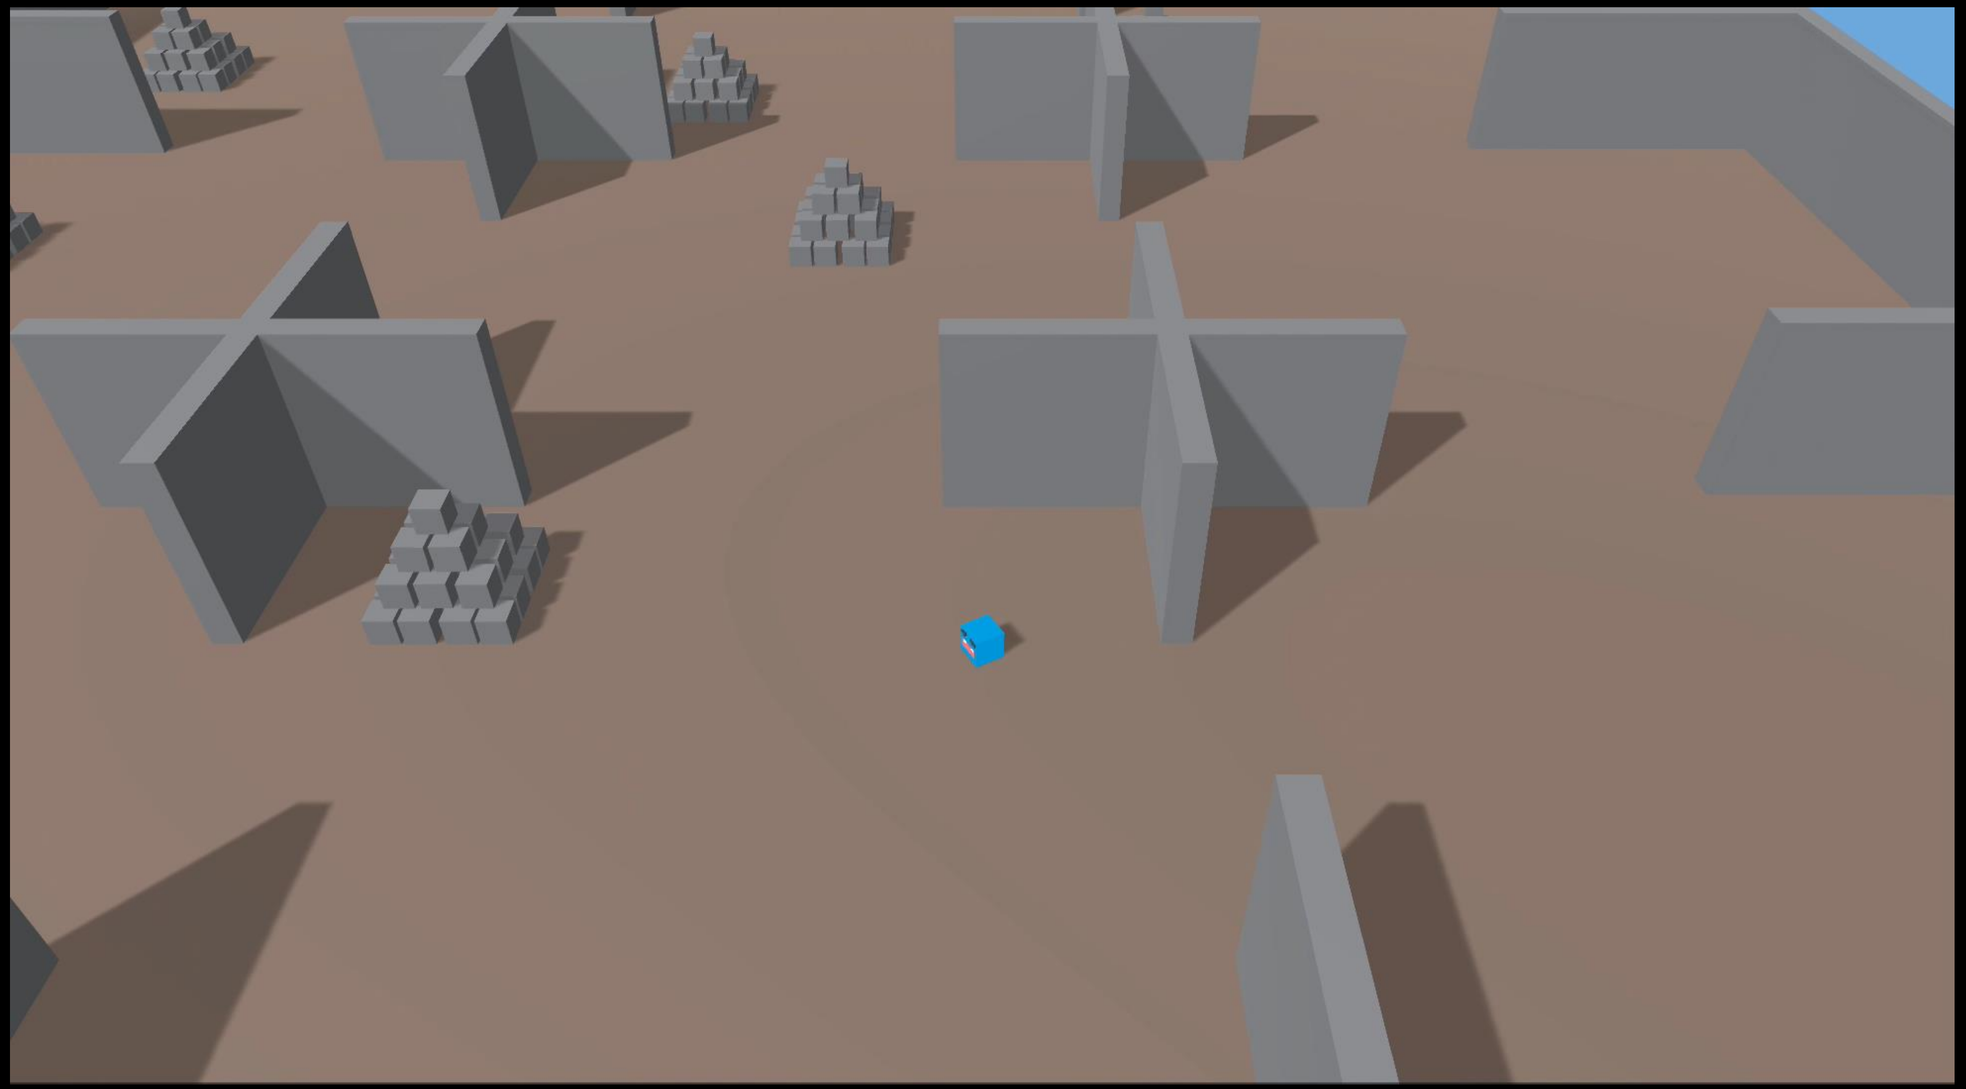

Intrinsic Reward Only

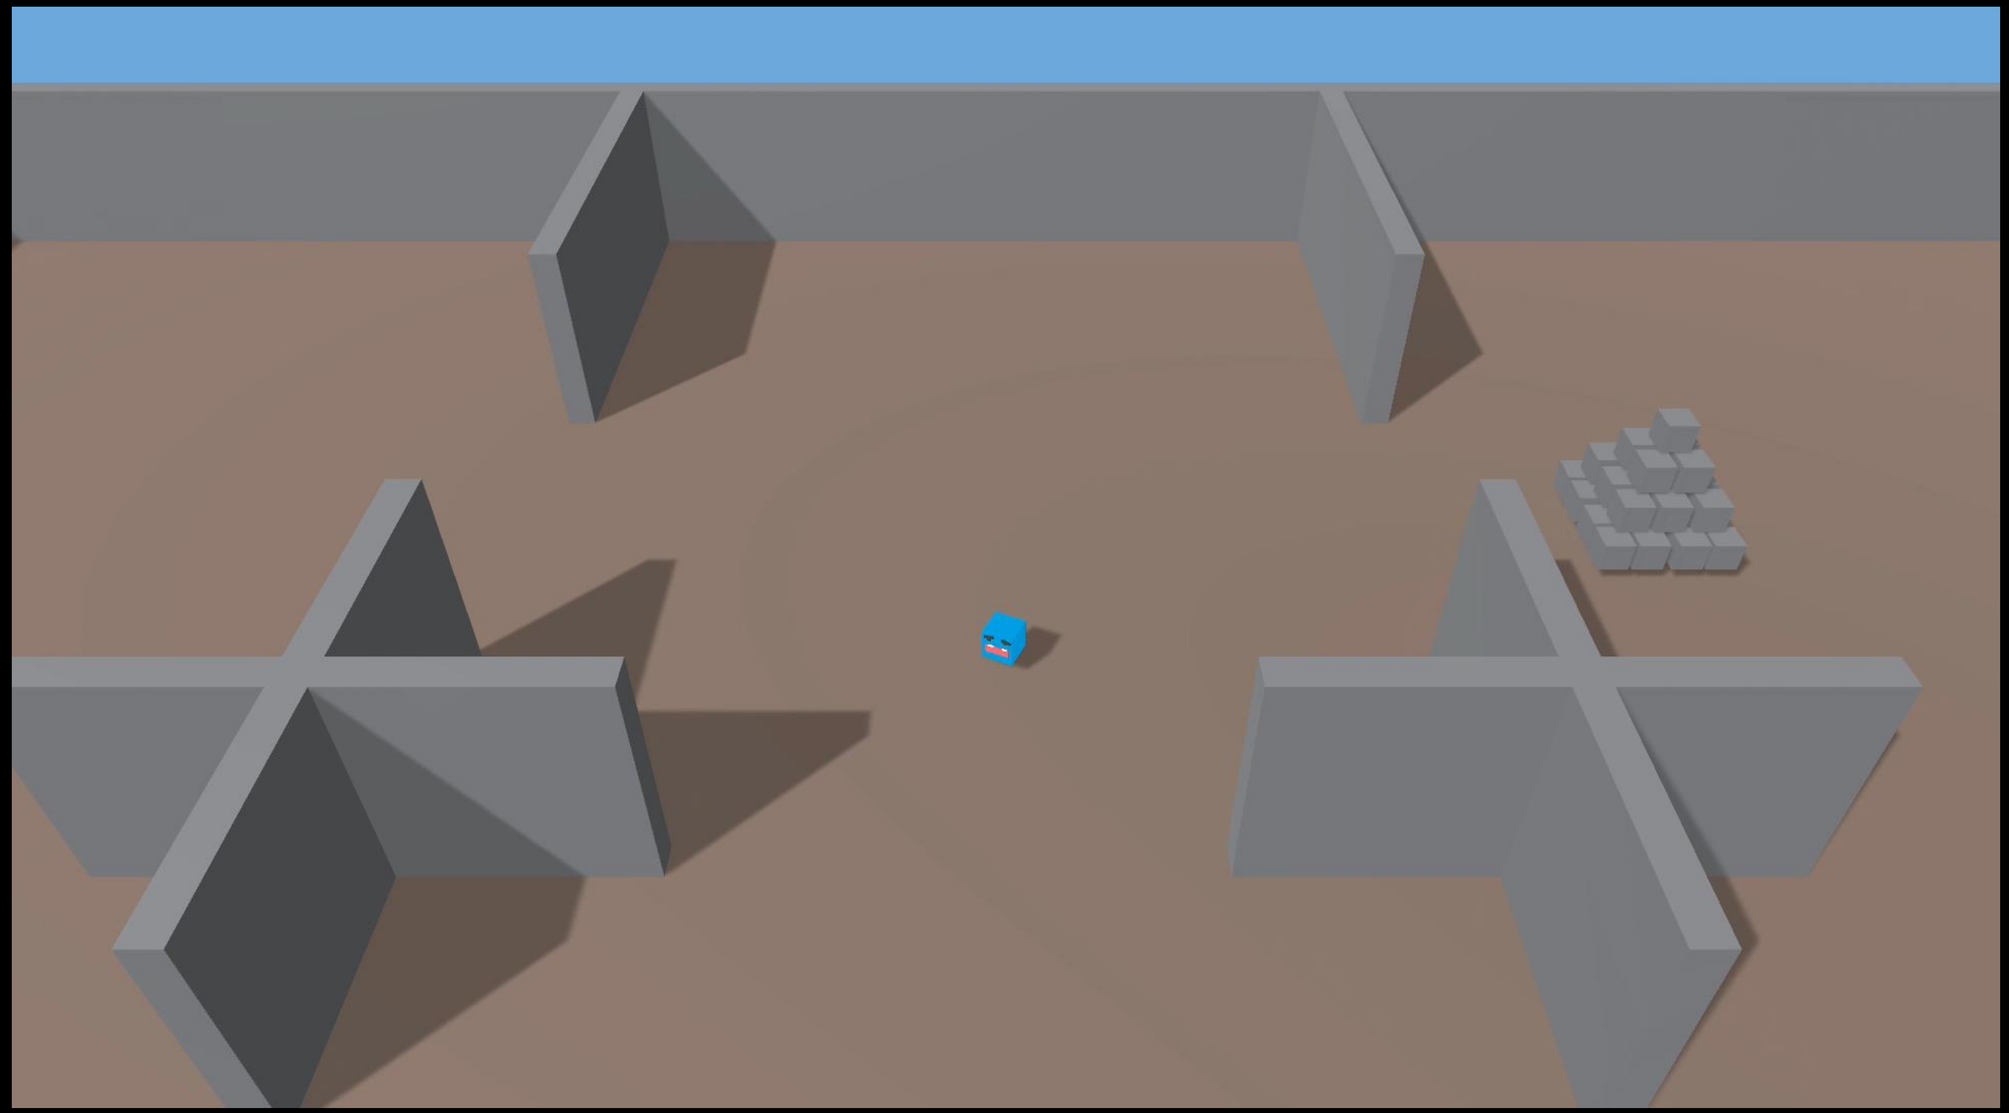

Both Rewards

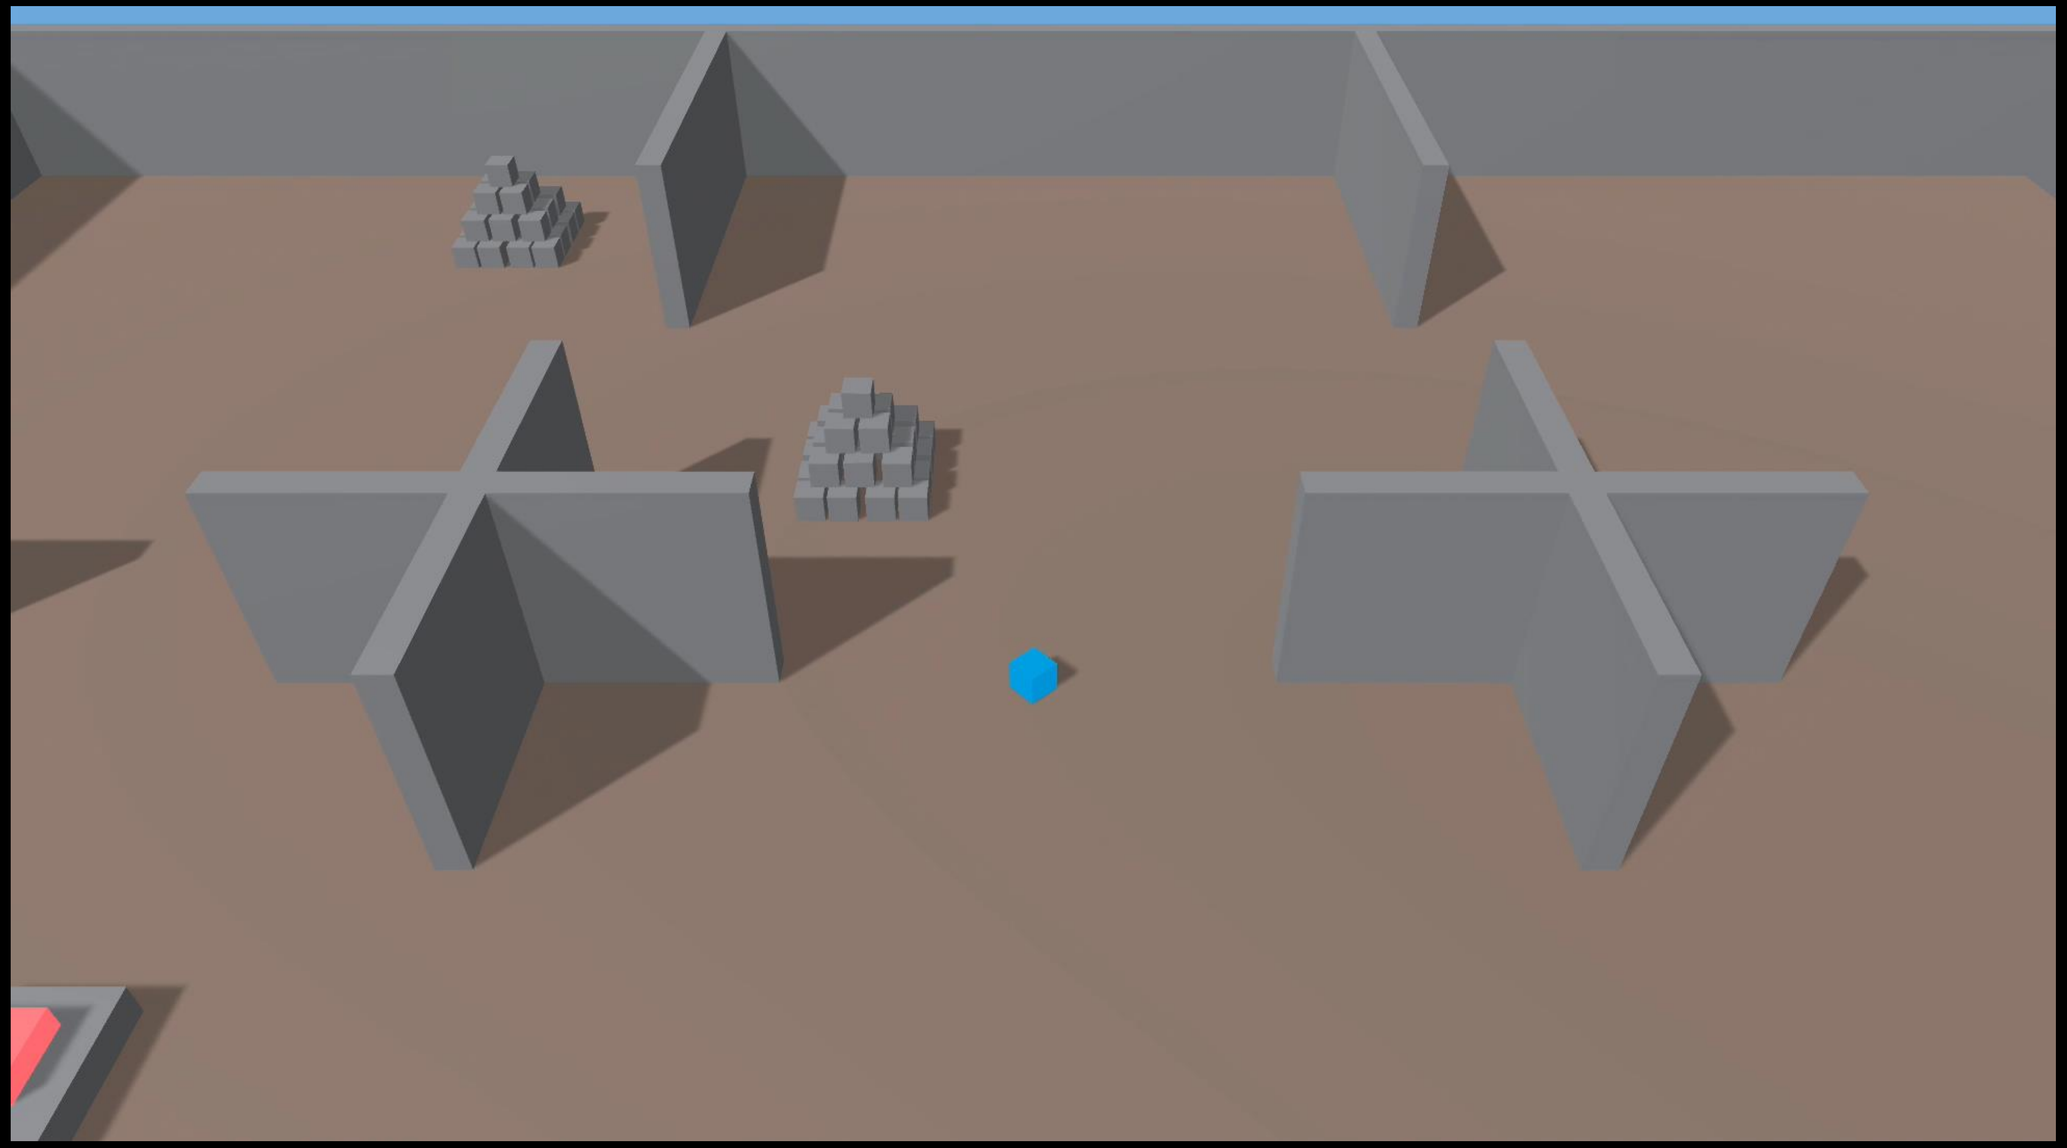

Results


PPO+Curiosity
PPO
Reward
1.5
1.0
0.5
0.0
−0.5
−1.0
25000
50000
75000
100000
125000
150000
175000
200000
Steps

In Progress - Hierarchical Control

Cognitive Agent

- Observes “Visual” information
- Acts on target direction
- Reward: goal contact

Motor Agent

- Observes proprioceptive information; target direction
- Acts on Joint torques
- Reward: target direction alignment

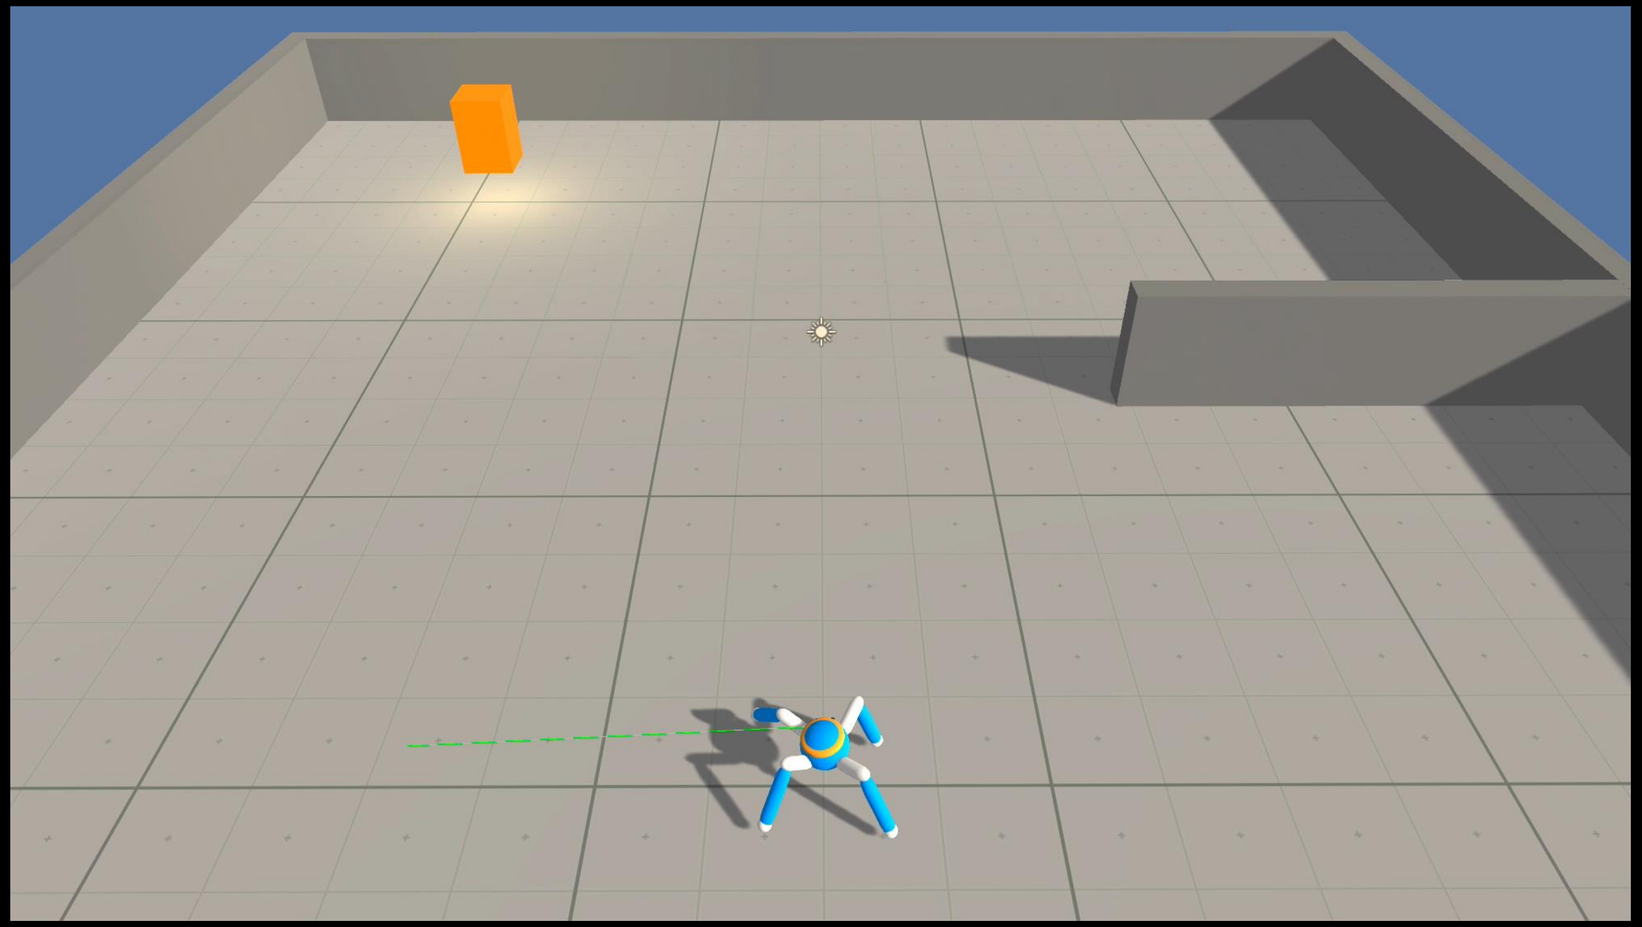

Get ML-Agents at GitHub Now
github.com/Unity-Technologies/ml-agents

Please share your feedback!

Contact Us
ml-agents@Unity3d.com

Arthur Juliani
Machine Learning Engineer
arthurj@unity3d.com
@awjuliani