goto; chicago

On the Road to Artificial General Intelligence

Danny Lange



follow us @gotochgo

goto; chicago





Click ‘Rate Session’
to rate session
and ask questions.

follow us @gotochgo

On the Road to Artificial General Intelligence

Danny Lange
VP of AI and Machine Learning, Unity



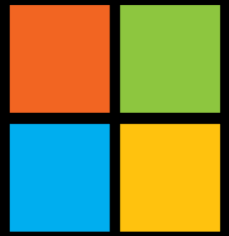
Generative Art – Made with Unity







UBER



About Unity Technologies

- Real-time 3D Development Platform
  - Game, AR/VR, film, automotive
- 60% of top 1,000 games
- On over 4 billion unique devices
- 1.8 billion monthly active players
- 2,400 employees with HQ in San Francisco



What is Artificial Intelligence?

Siri and Alexa

Amazon & Netflix Recommendations

Fraud Detection Services

Equity Trading

Facebook Feed

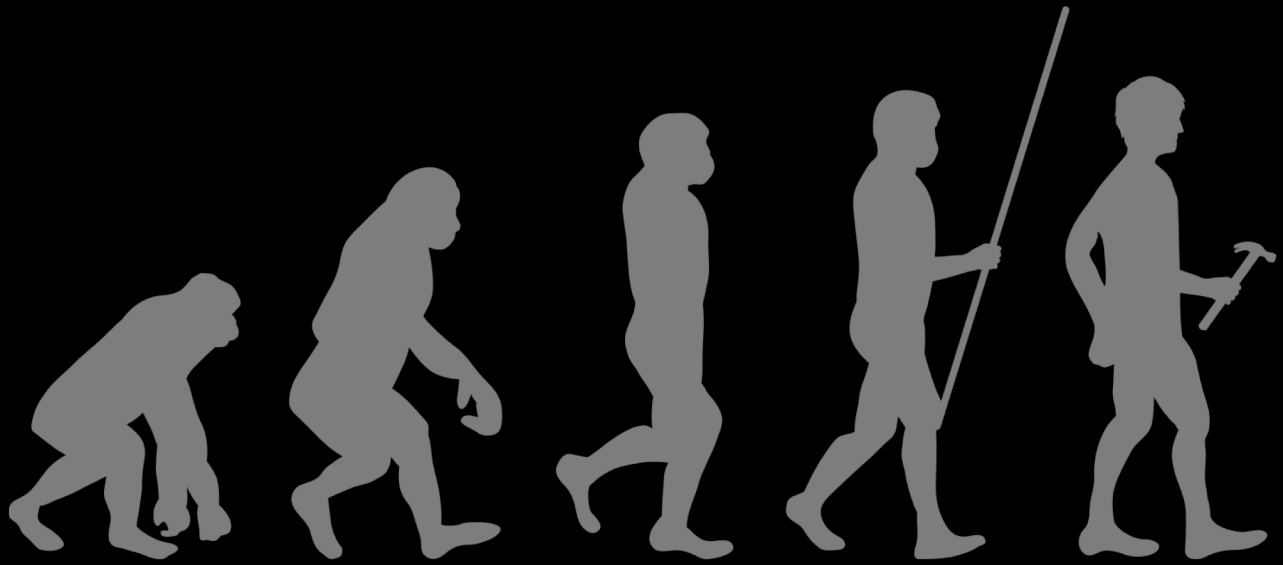
Every other Job Title on LinkedIn

What is *Real* Intelligence?

What is *Real* Intelligence?

In·tel·li·gence *[in ˈteləjəns]* – the ability to acquire and apply knowledge and skills.

What is *Real* Intelligence?



What is the only Real Intelligence you Truly Know?

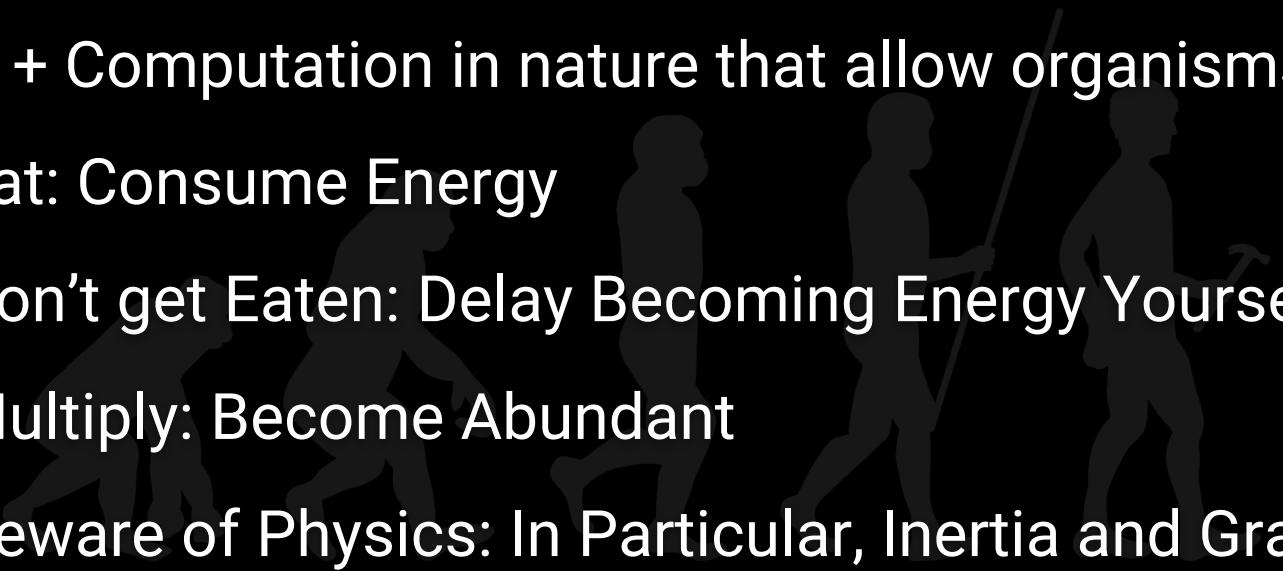


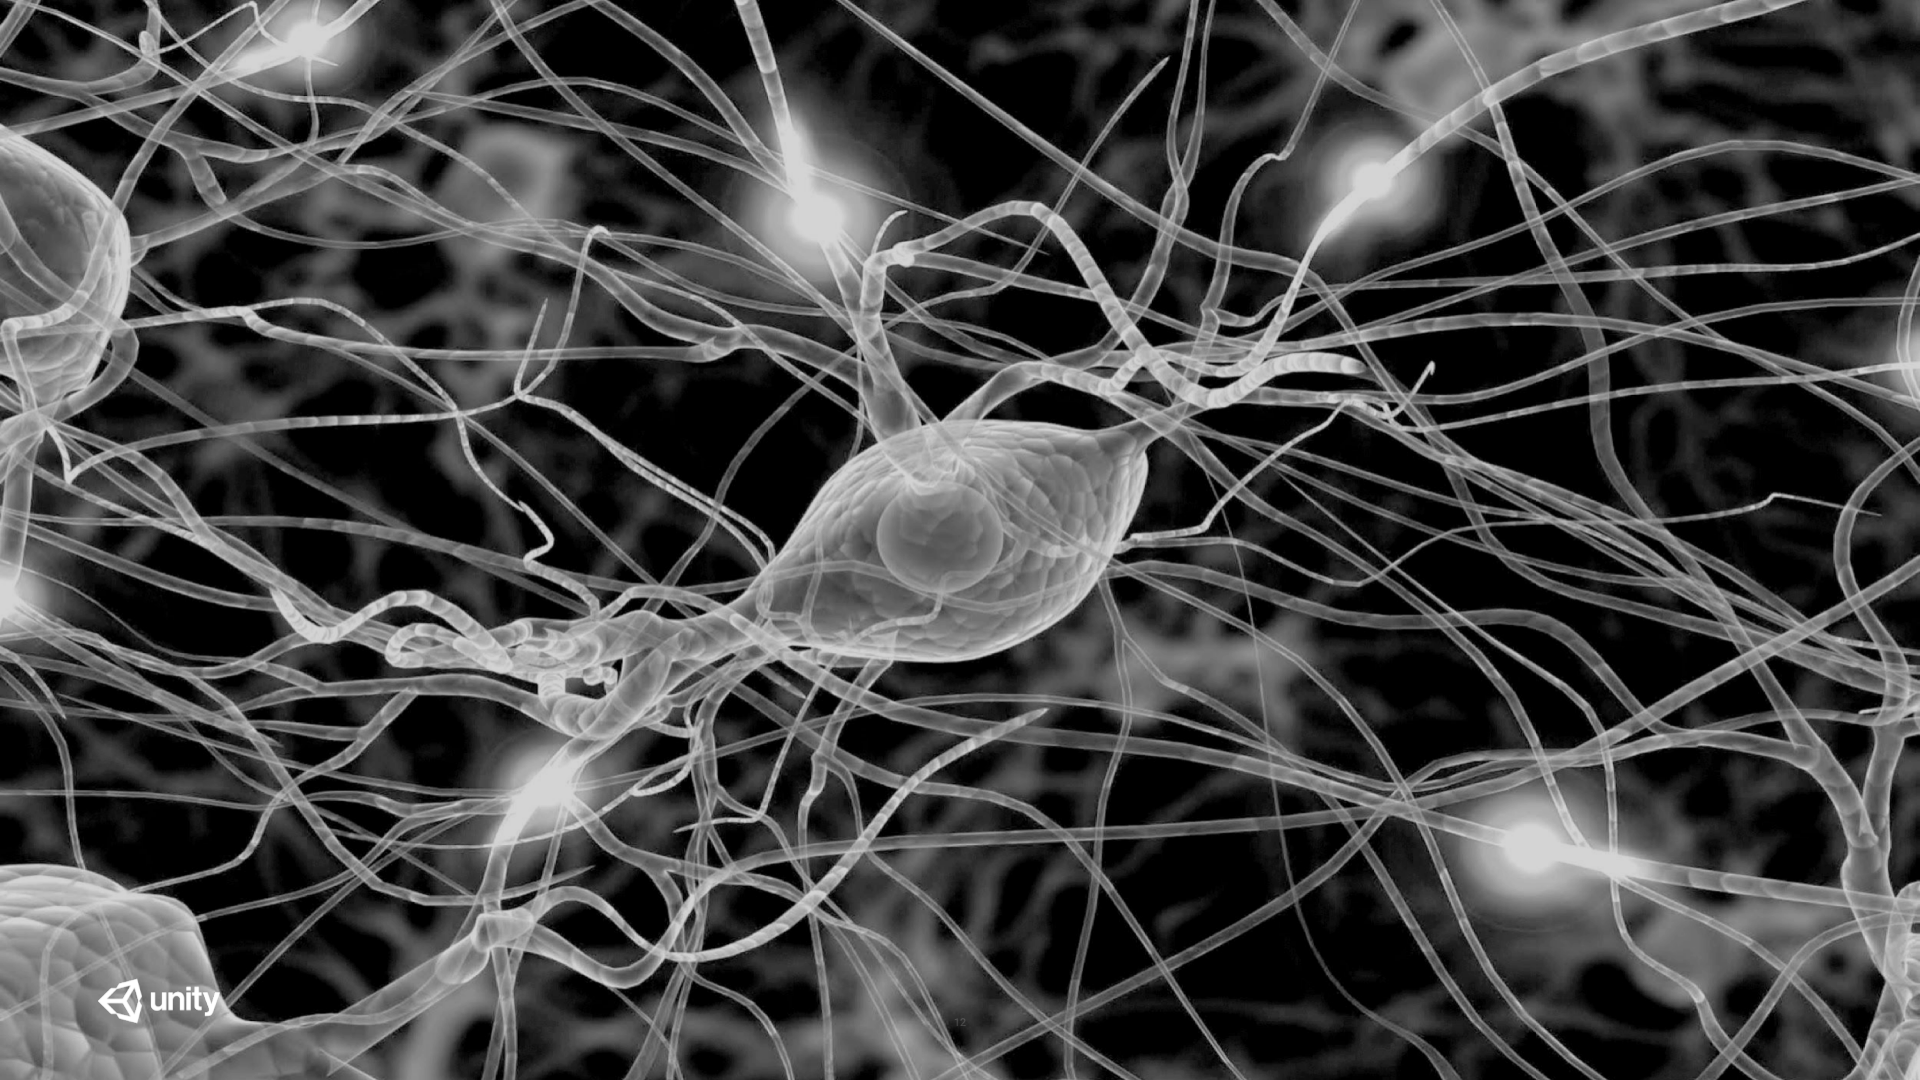
unity

Intelligence in Biological Systems

Senses + Computation in nature that allow organisms to:

- Eat: Consume Energy
- Don’t get Eaten: Delay Becoming Energy Yourself
- Multiply: Become Abundant
- Beware of Physics: In Particular, Inertia and Gravity
- Agency: The Ability to Act upon the Environment

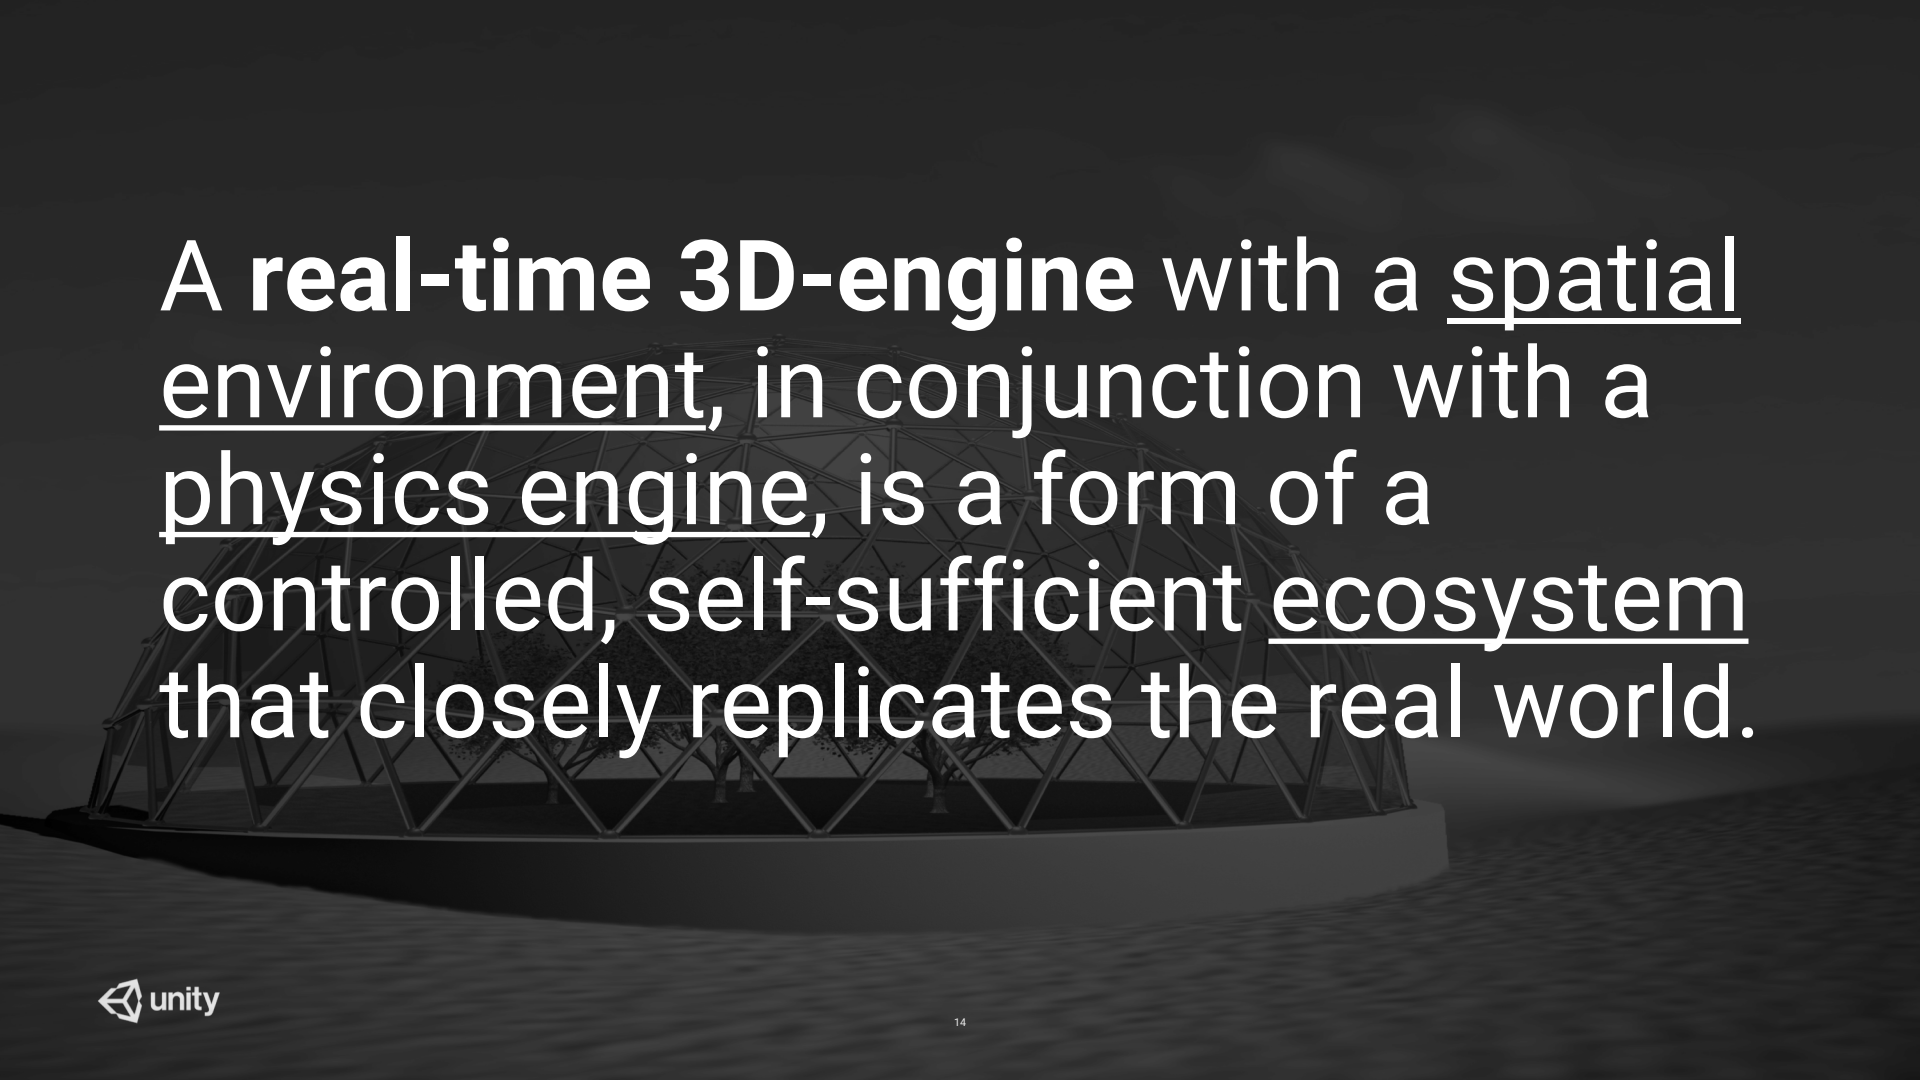


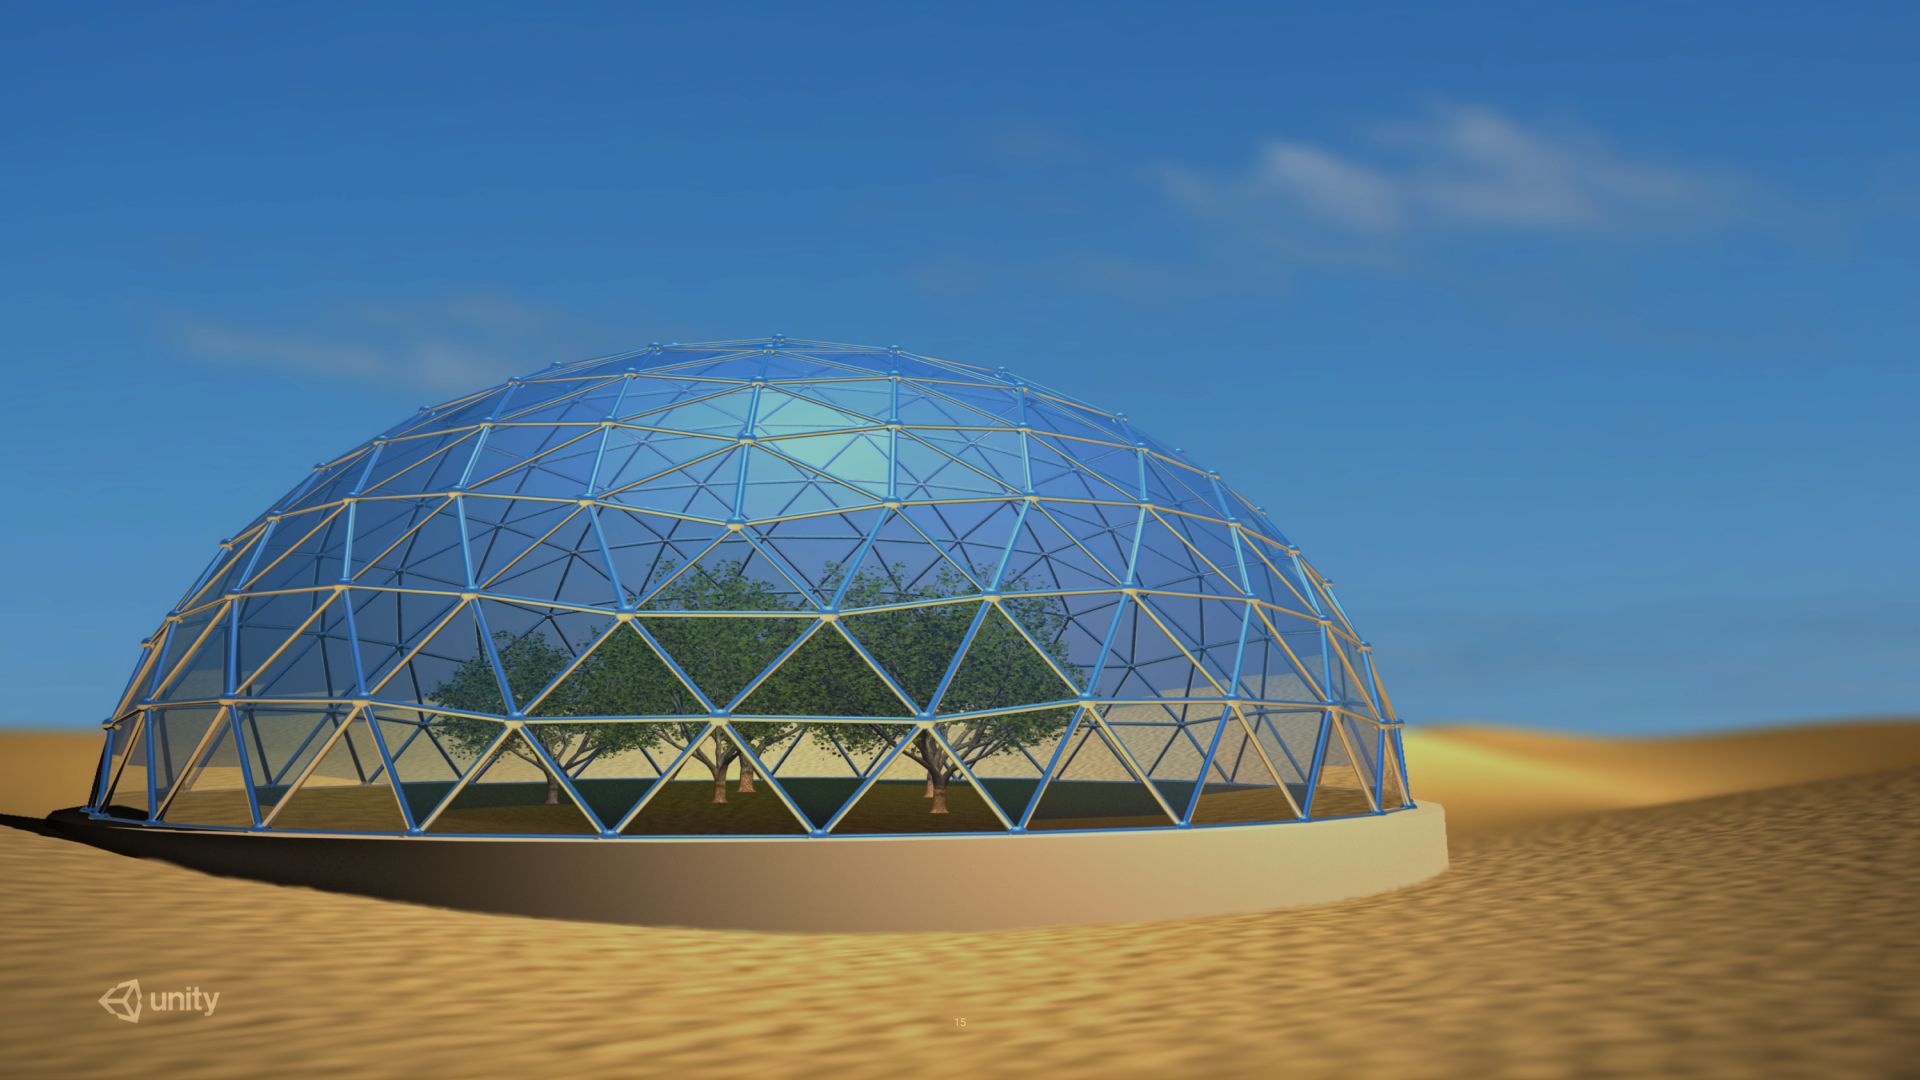

unity

Intelligence Is Achieved from Infrastructure

Nature implemented intelligence

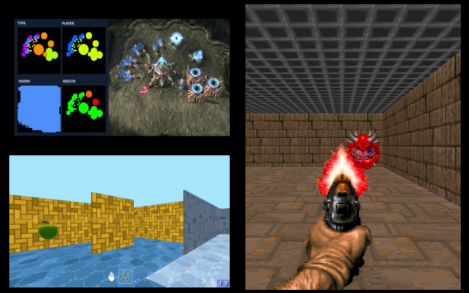
- Chemical Mechanisms
- Cellular Structures
- Multicellular Organisms with Messaging Systems
- Movement: Neuromuscular Junction
- Sensors: Touch, Sight, Hearing, Taste, and Smell


A **real-time 3D-engine** with a spatial environment, in conjunction with a physics engine, is a form of a controlled, self-sufficient ecosystem that closely replicates the real world.

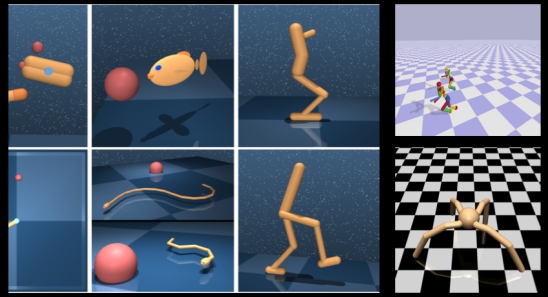



Introducing
Unity ML-Agents Toolkit
Your Private AI Biodome
unity

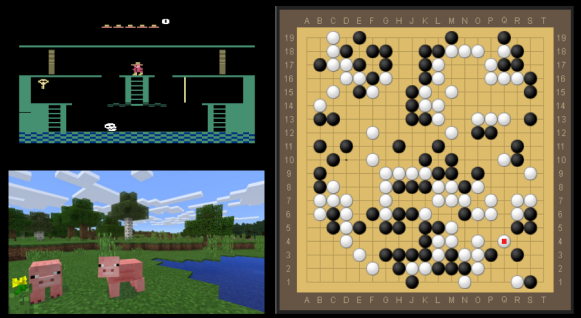
ML Training Environments



Visual



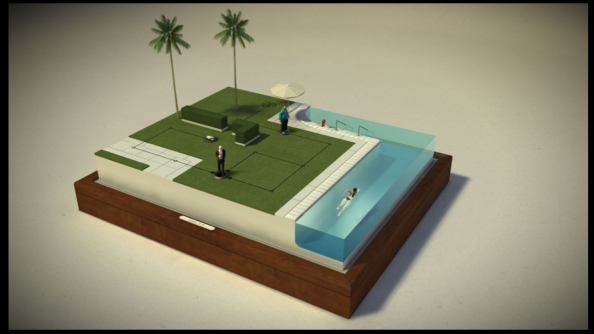
Physical

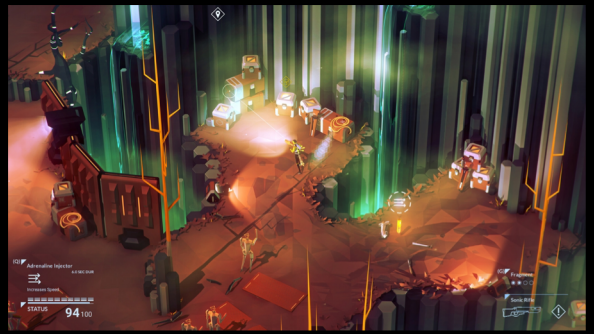


Cognitive

The Unity Ecosystem















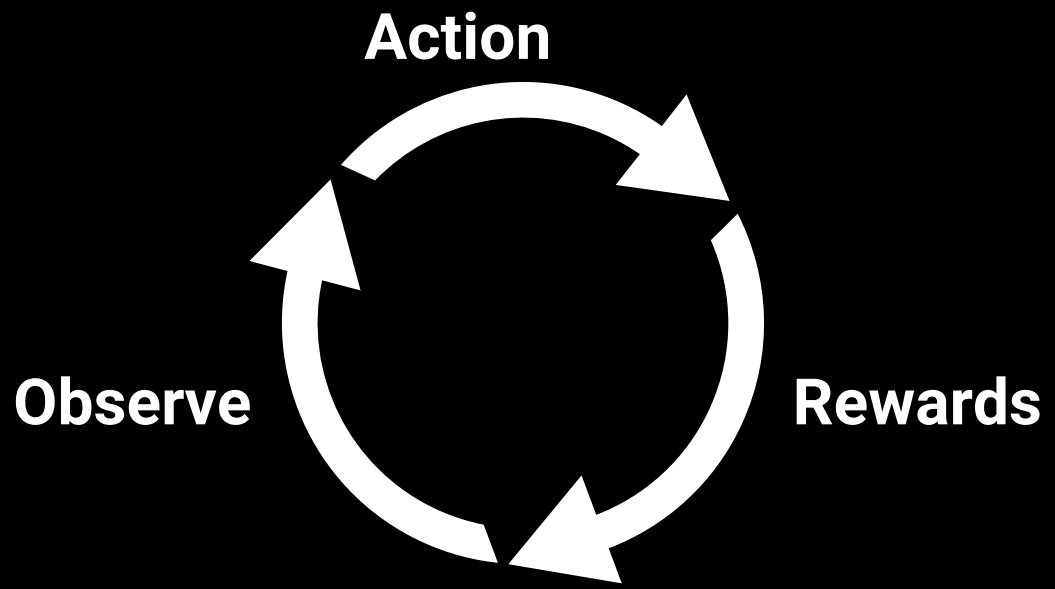


Demis Hassabis - co-founder and CEO of DeepMind

*As a former video game designer myself, I couldn't be more excited to be collaborating with **Unity**, creating virtual environments for developing and testing the kind of smart, flexible algorithms we need to tackle real-world problems.*

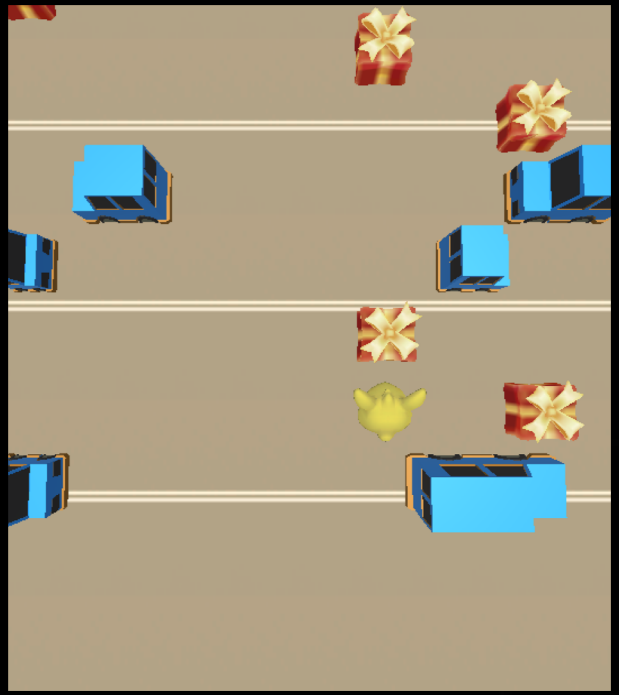


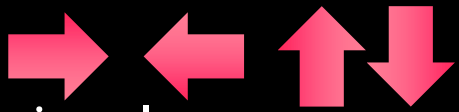
Nature’s Learning Method: Reinforcement



Explore → Exploit

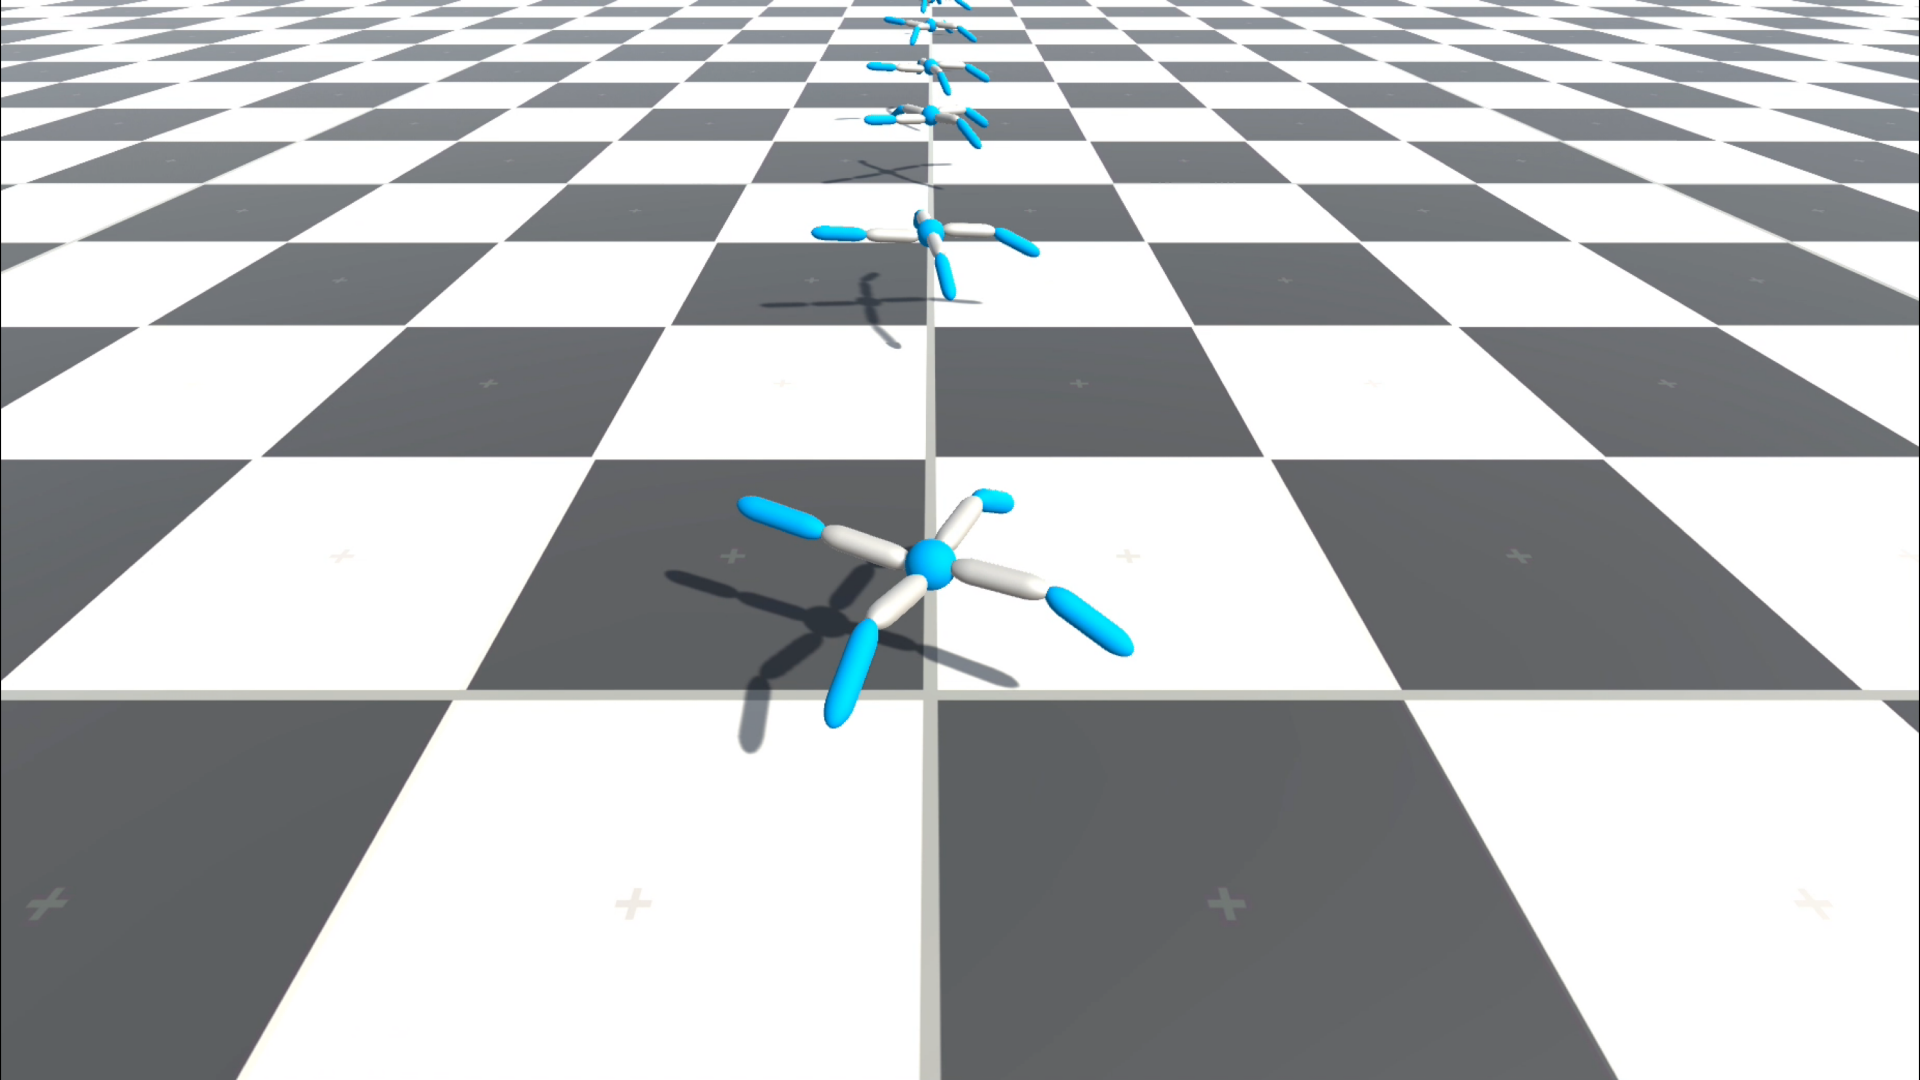
Chicken Crossing the Road: Tabula Rasa

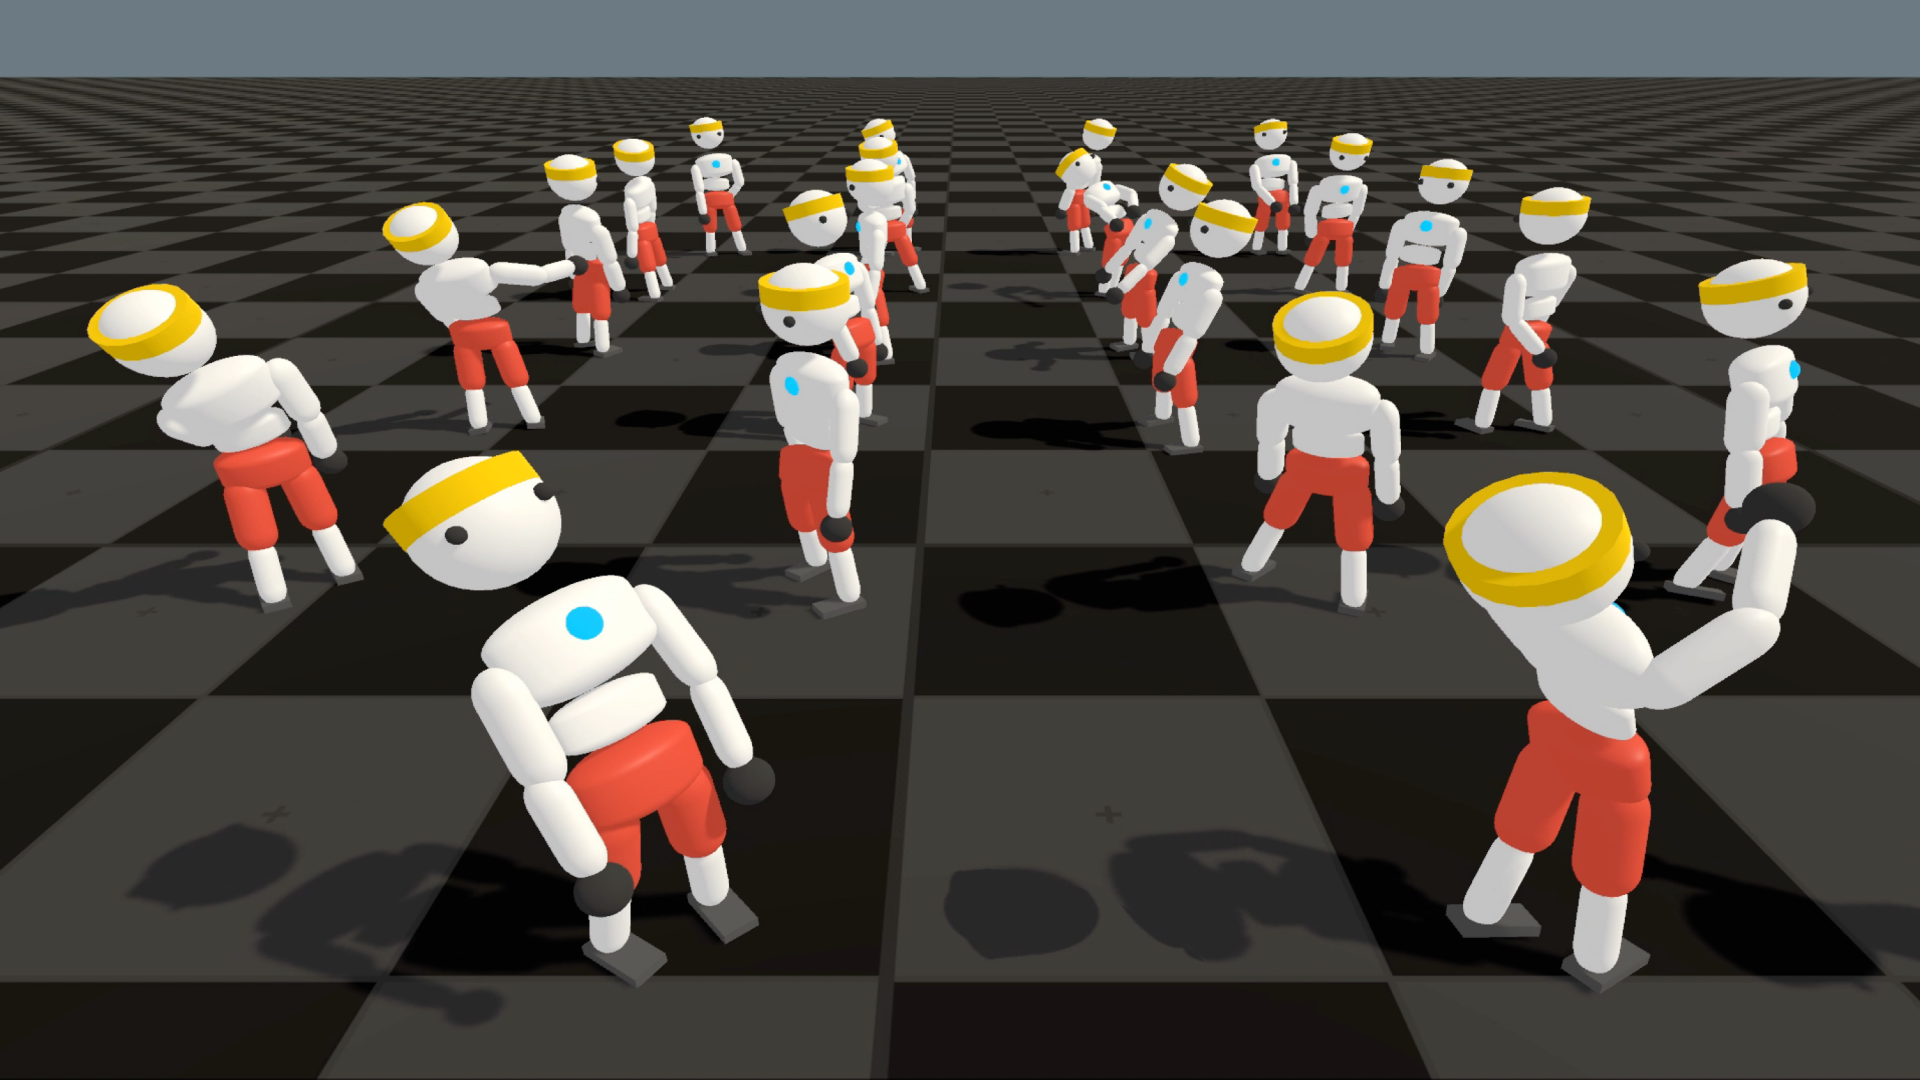


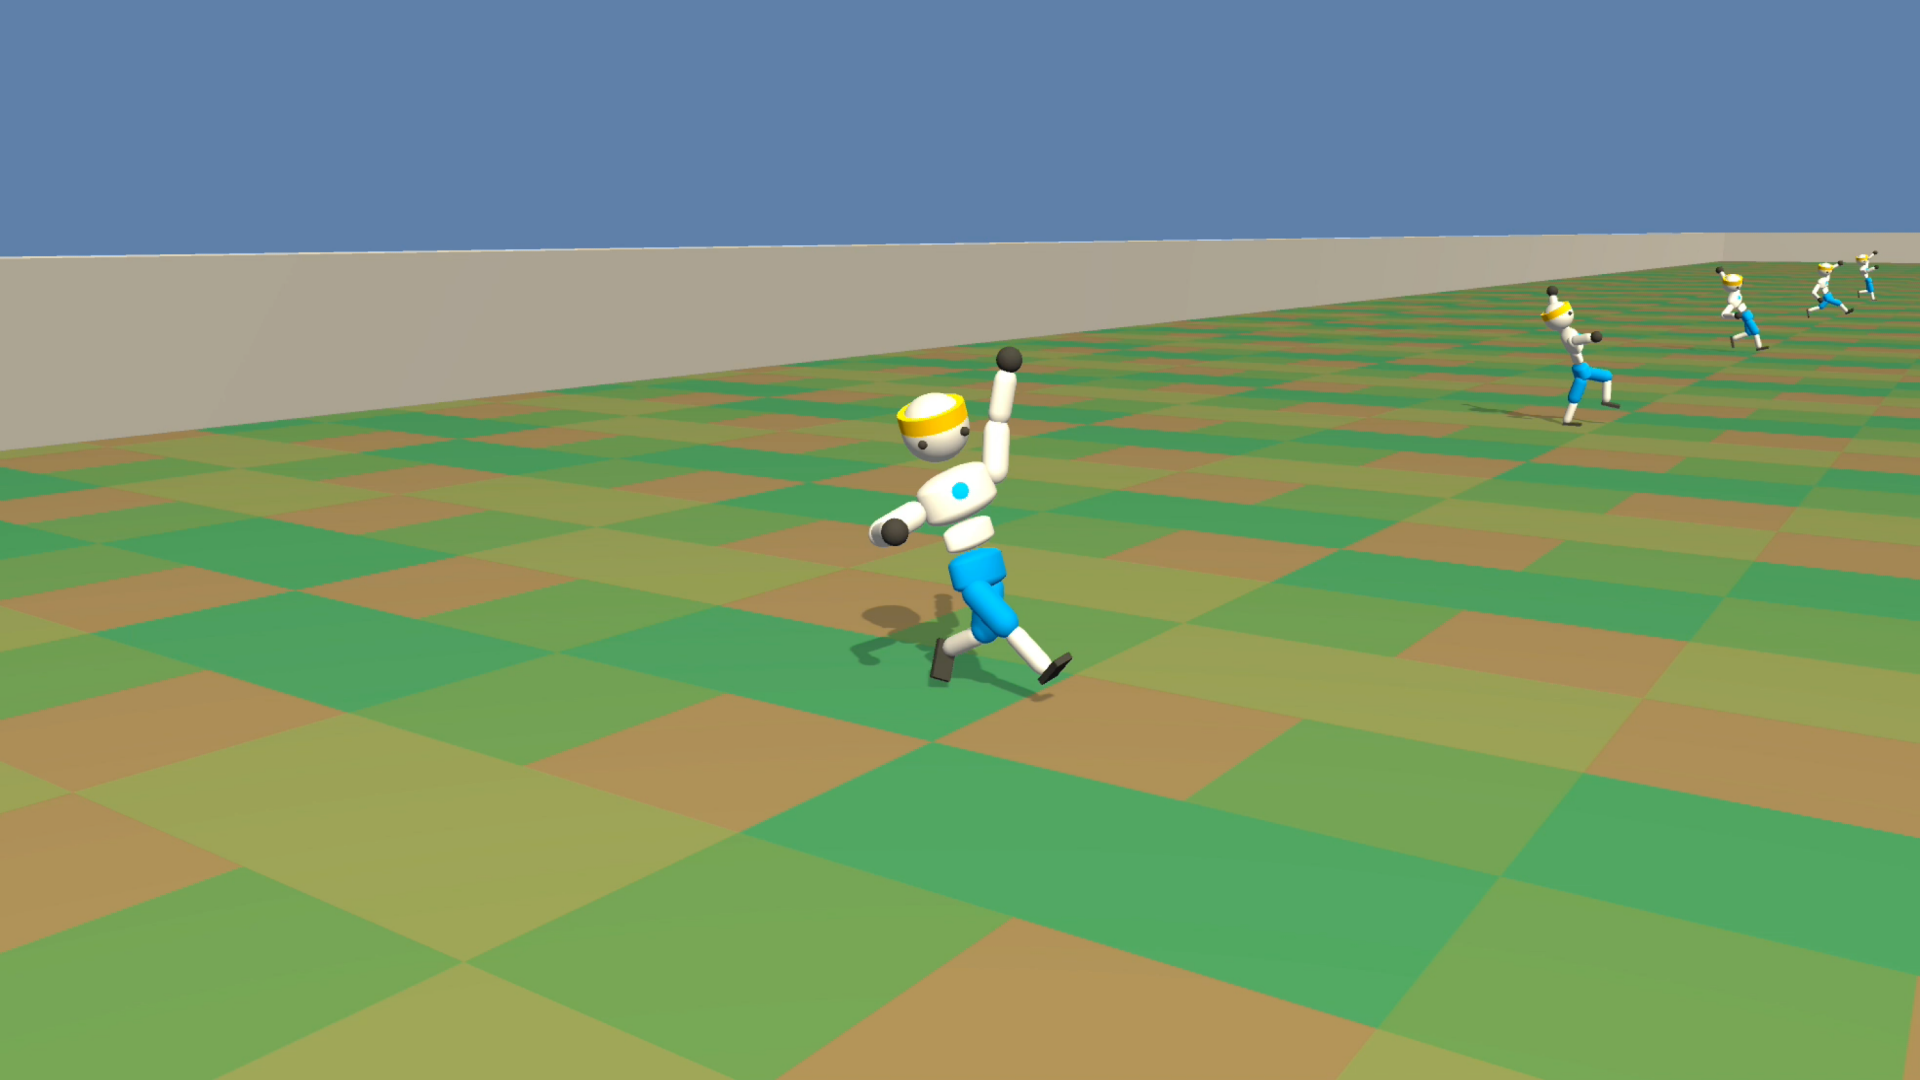
- Observe: Pixels in frame
- Actions: 
- Rewards signal
  - Negative for being hit
  - Positive for gift pickup


Learning Scenarios

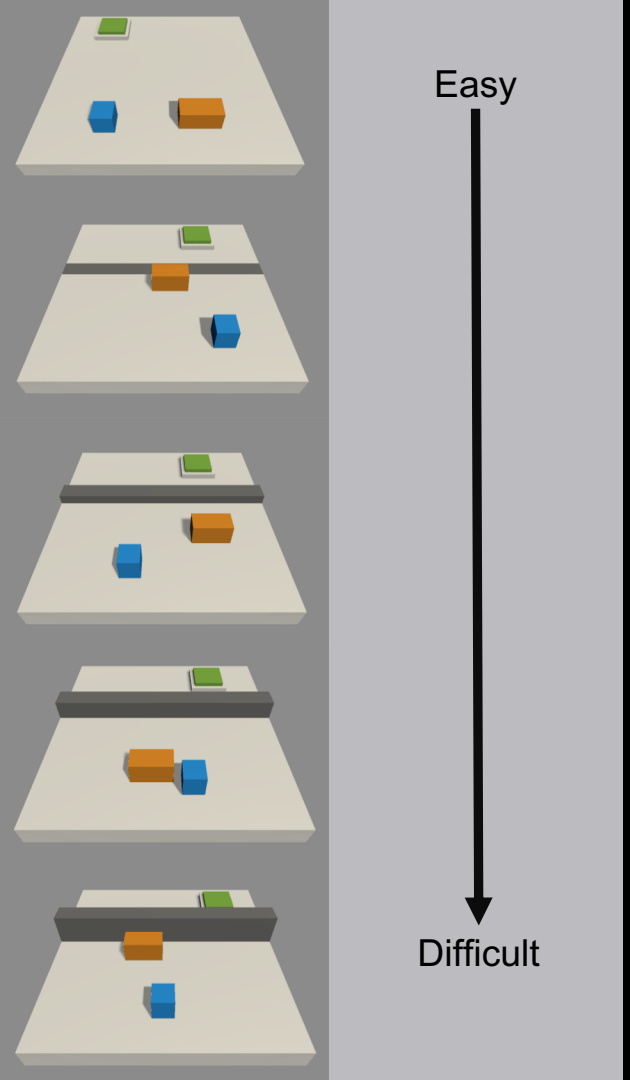
Control Learning







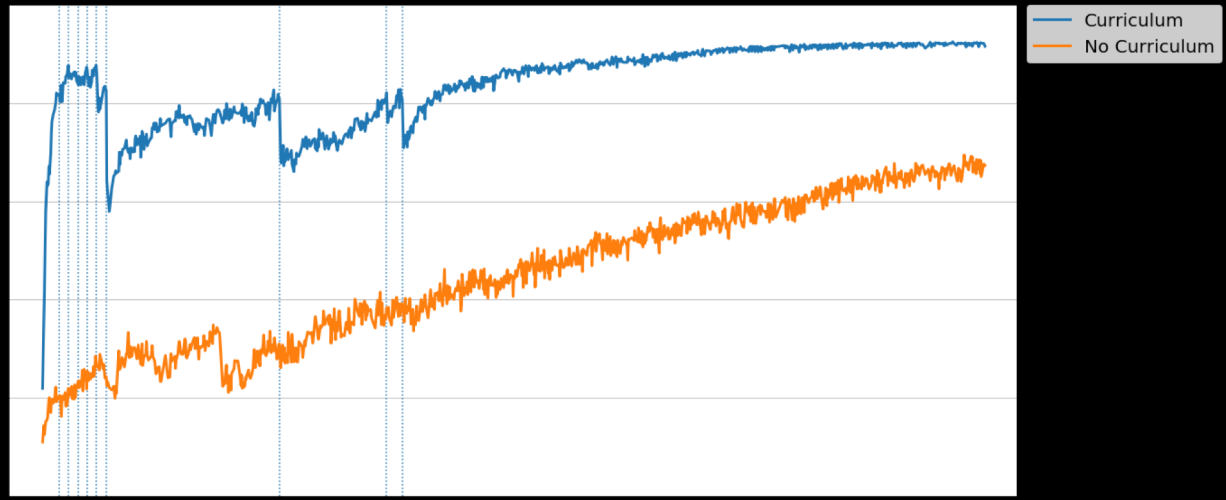
Curriculum Learning

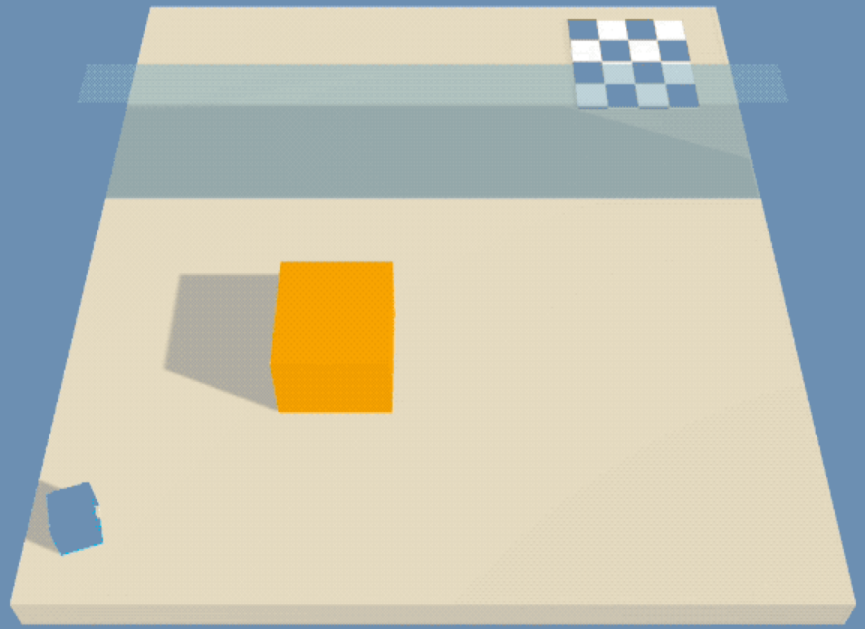


Easy

Difficult

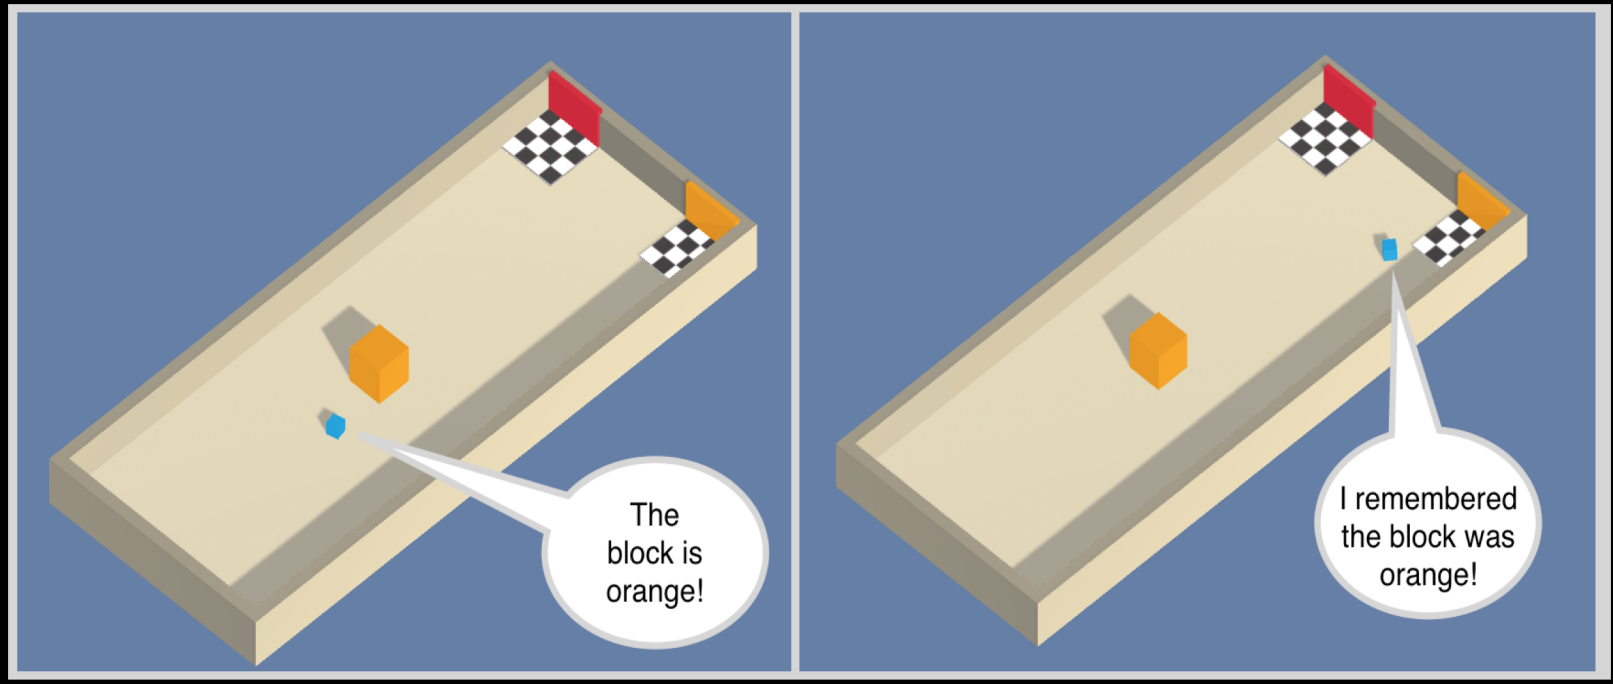
- Start easy
- Incrementally harder tasks
- Progress depends on graduation

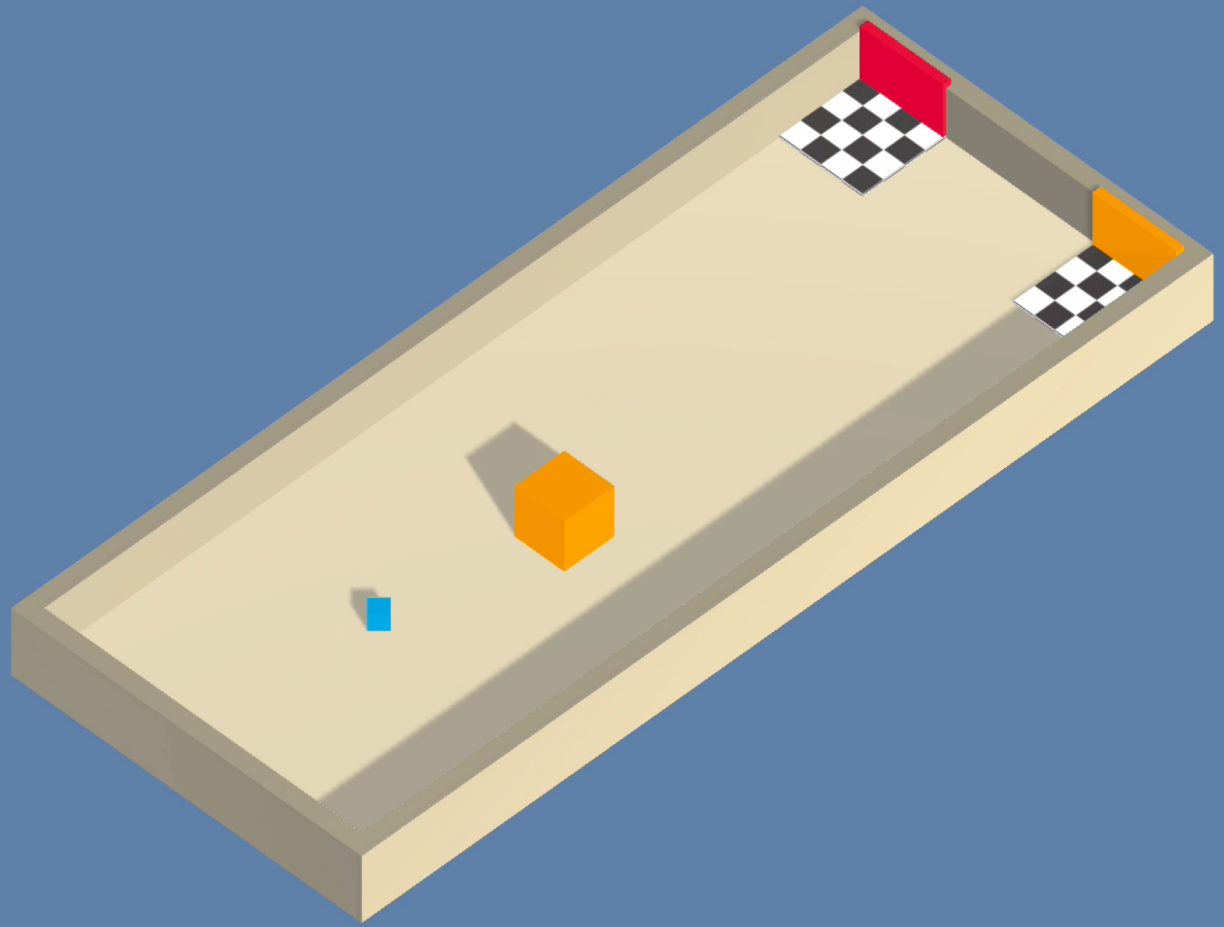




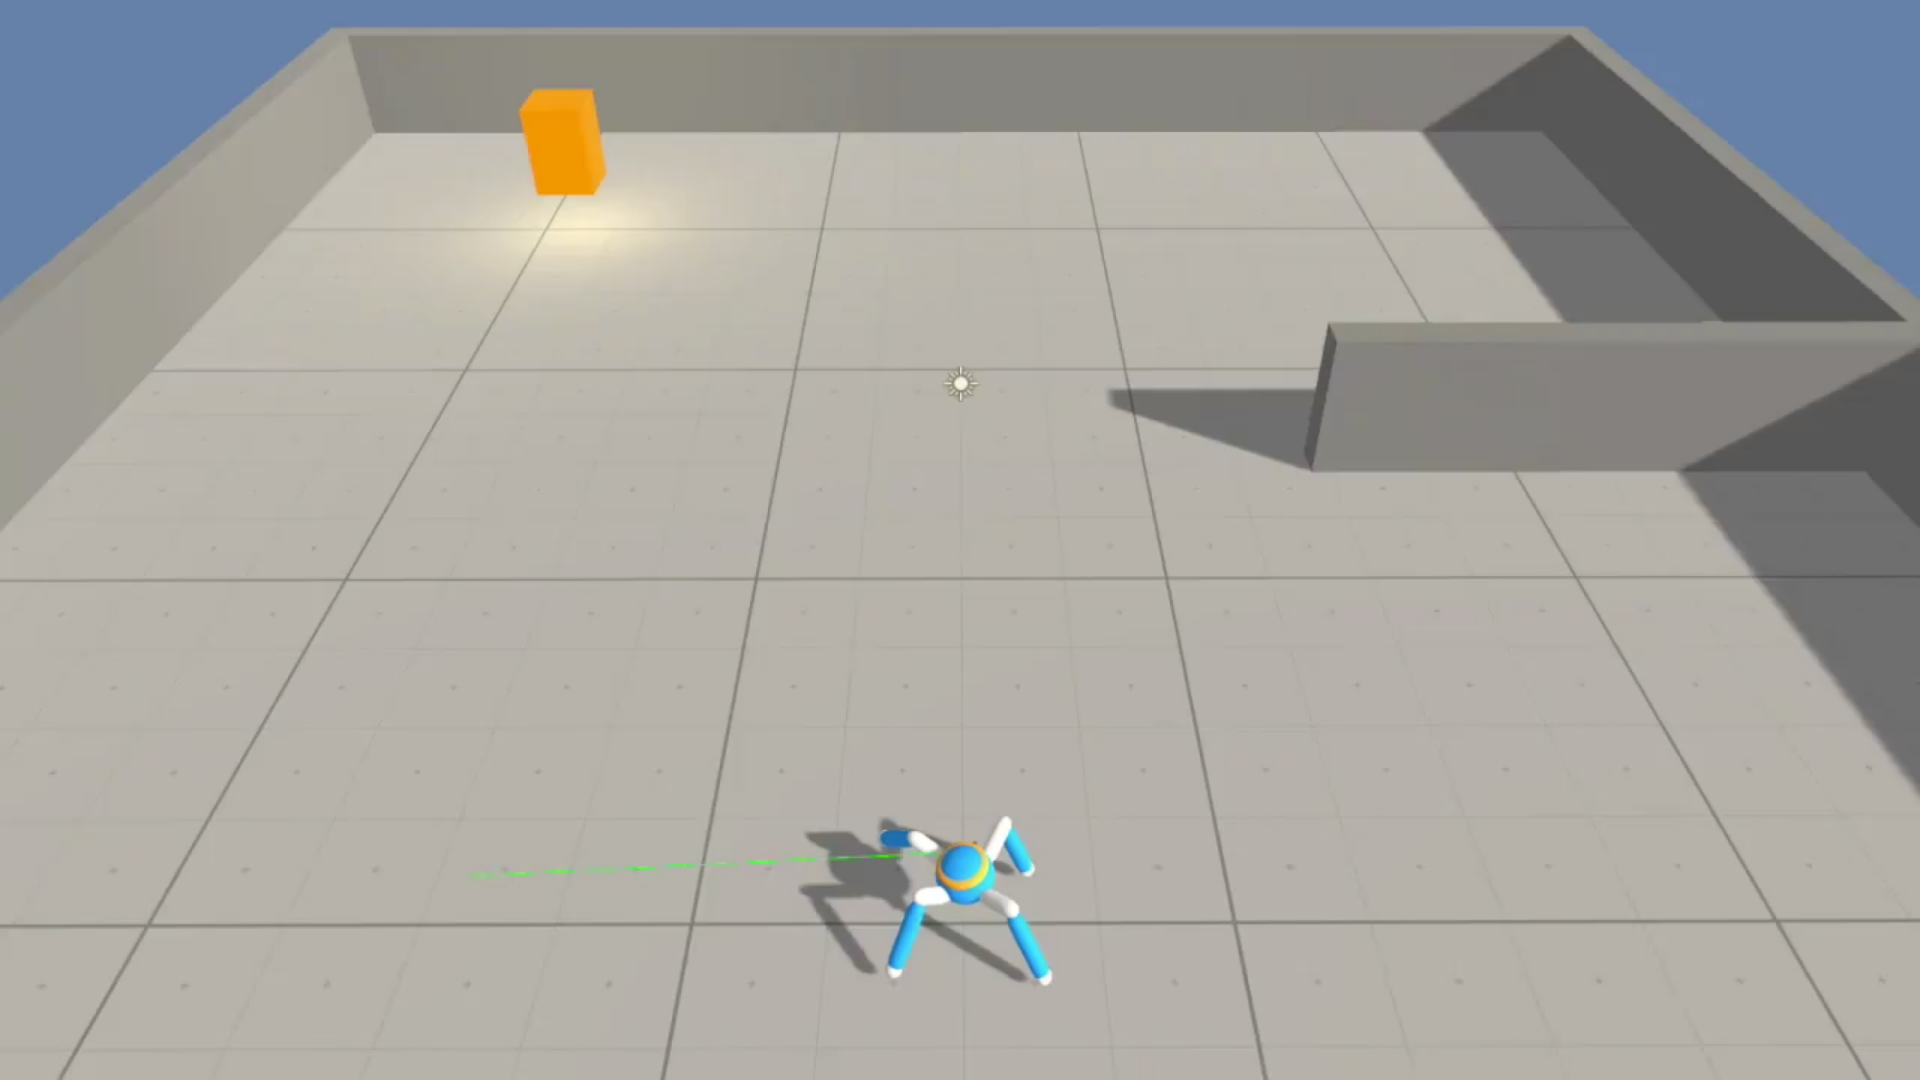
Memory in Learning

Enabling Long Short-Term Memory (LSTM)





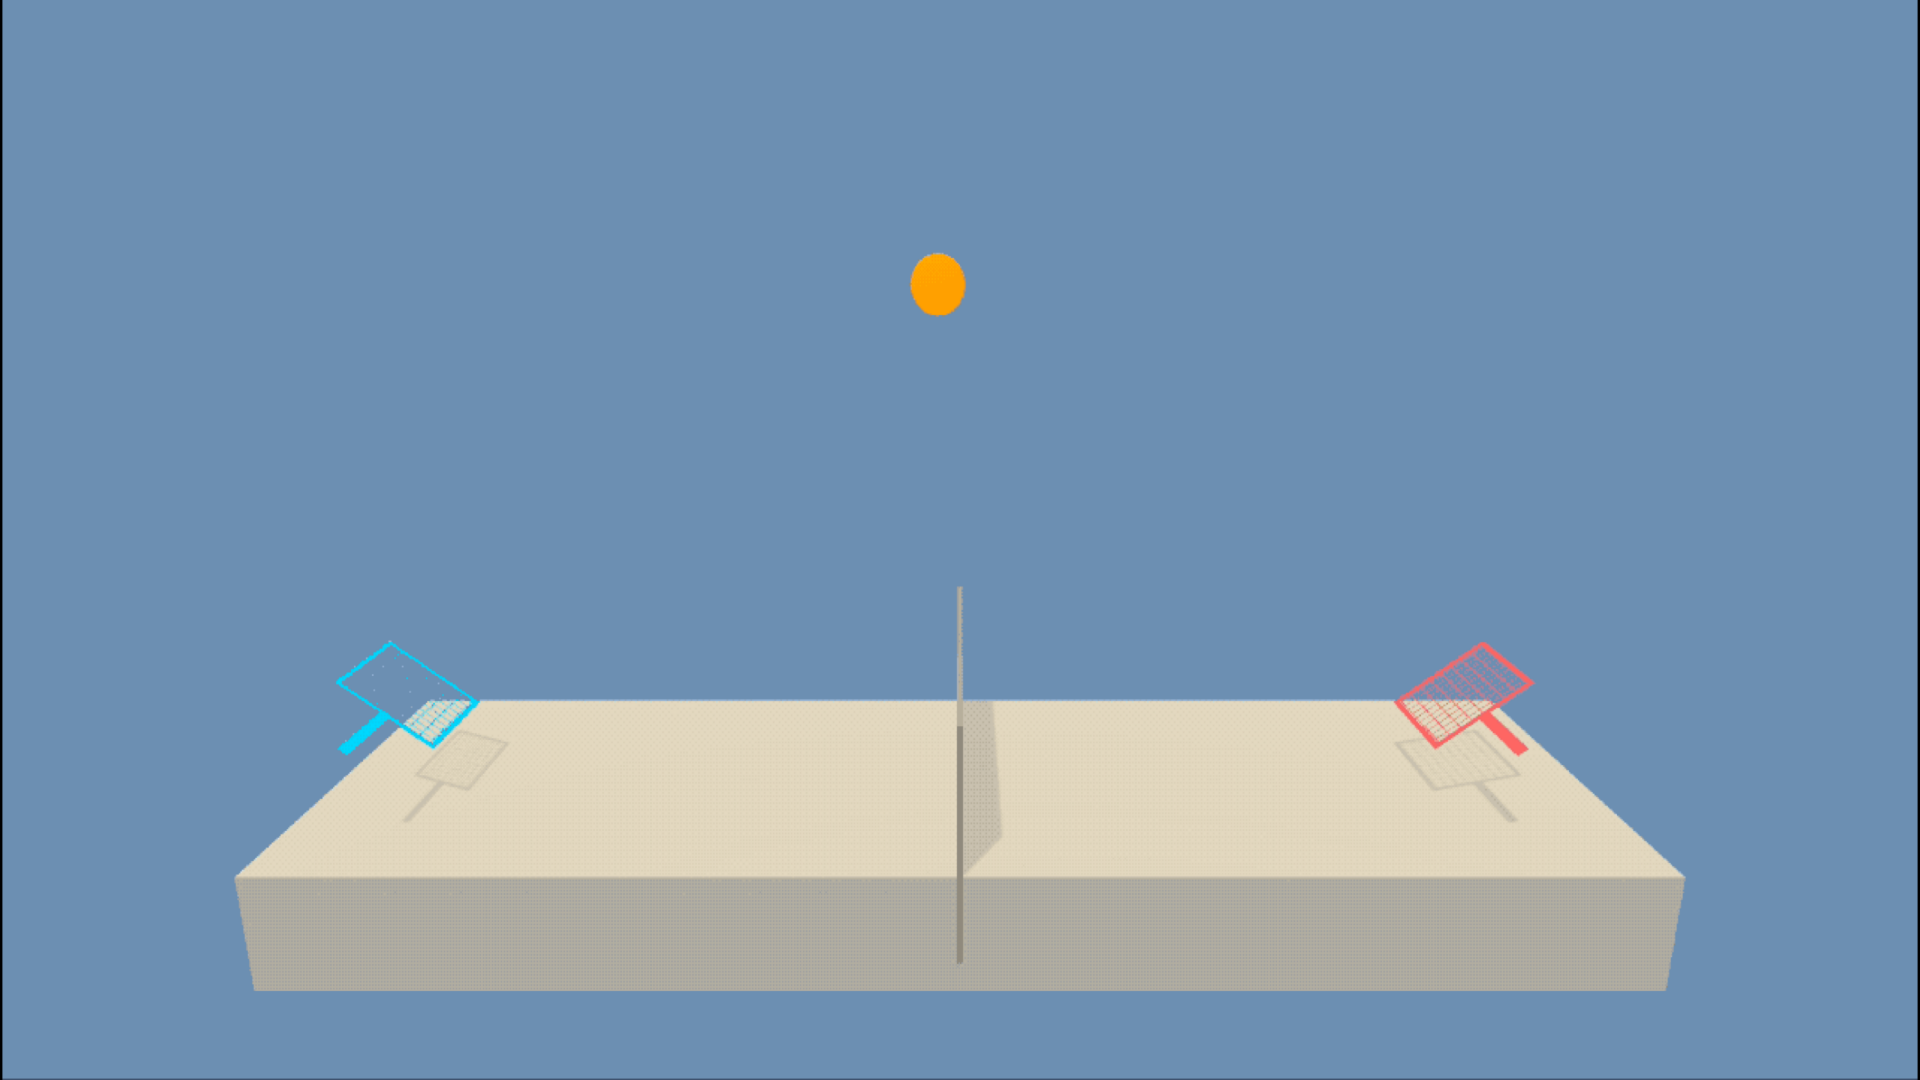
Hierarchical Learning Navigation + Control



Multi-agent Learning

Cooperative Rewards

- Single Brain with two agents
- Goal: Keep the ball off the ground
- Observations: Positions and velocities of racket and ball



Competitive Rewards

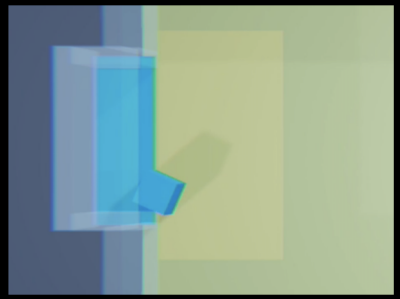
- Two Brains with each two agents
- Striker Objective: Score in your opponent’s goal
- Defense Objective: Keep your goal clean
- Observations: Local ray-cast perception on objects

Multi-Stage Soccer Camp

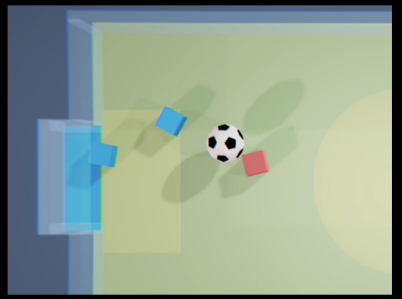
Offense
Train offensive brain with positive reward for ball entering opponent's goal

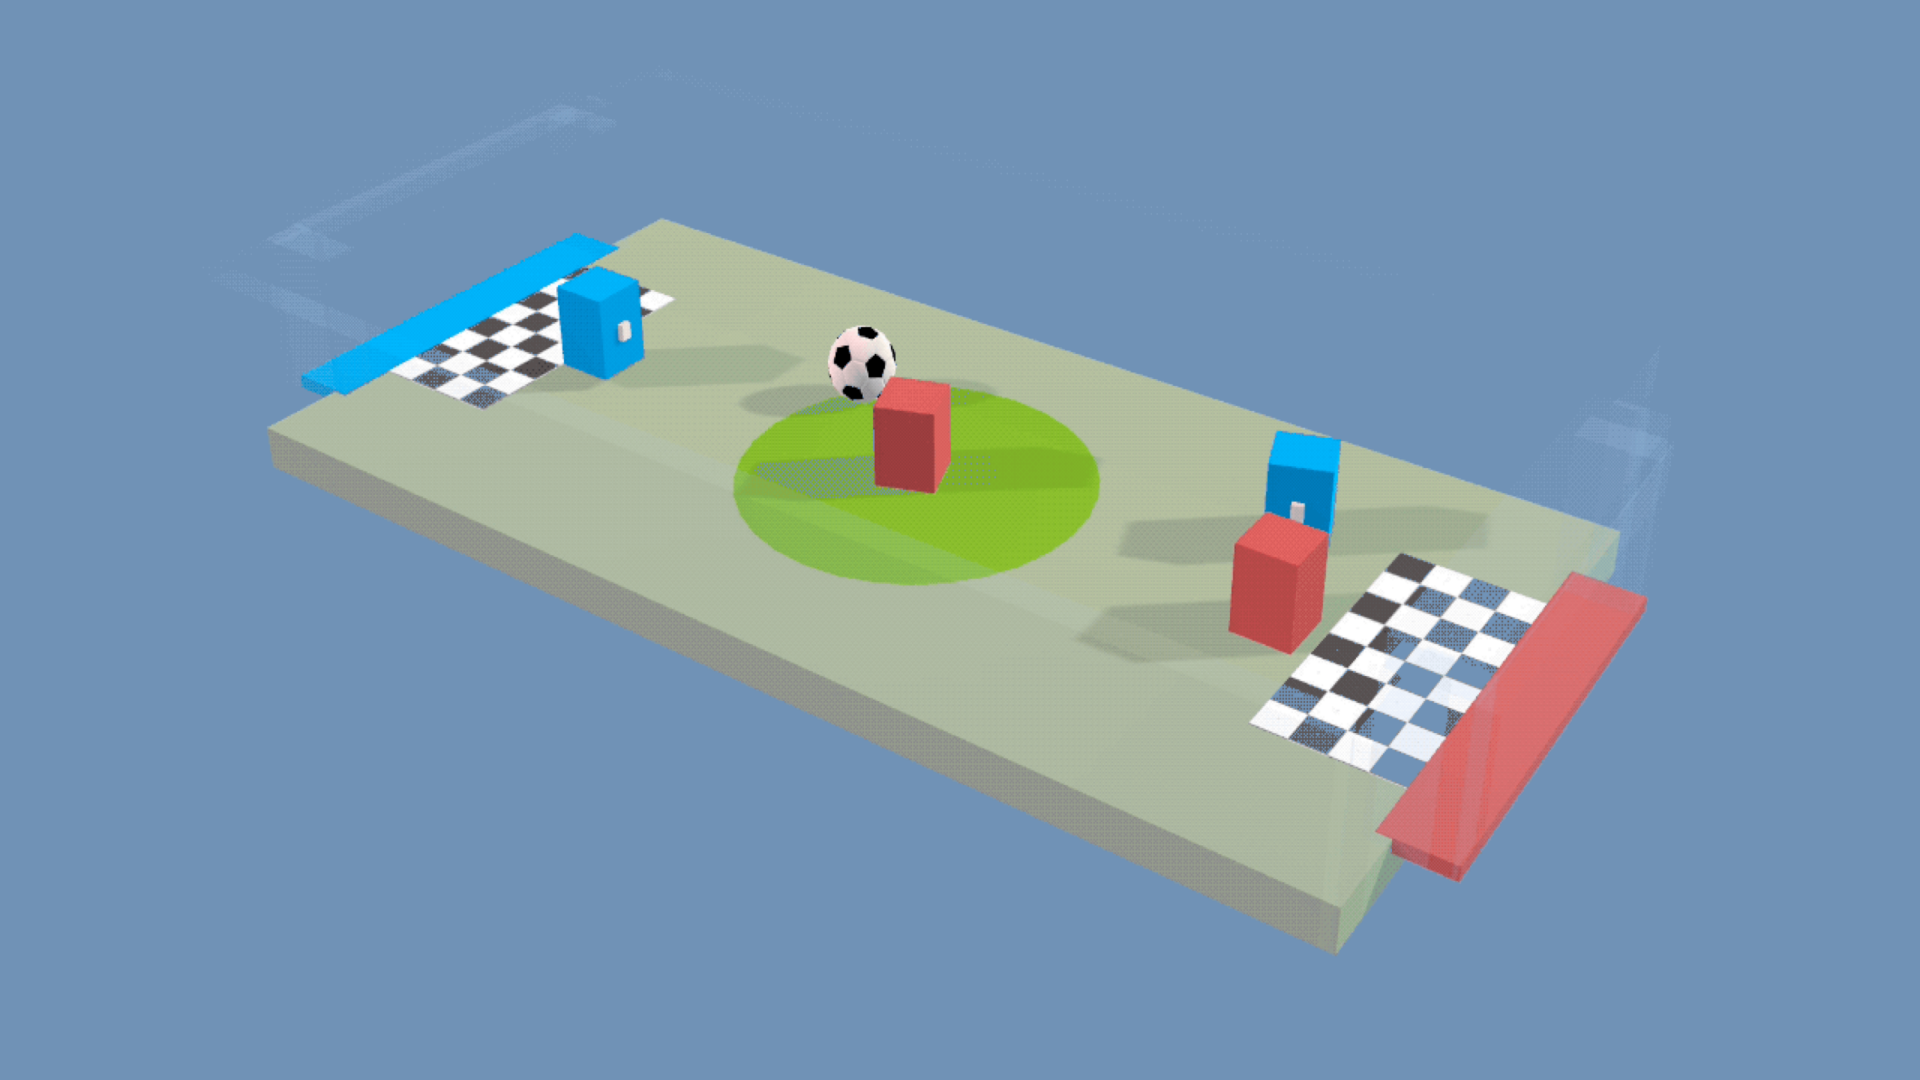


Defense
Train defensive brain with negative reward for ball entering own goal



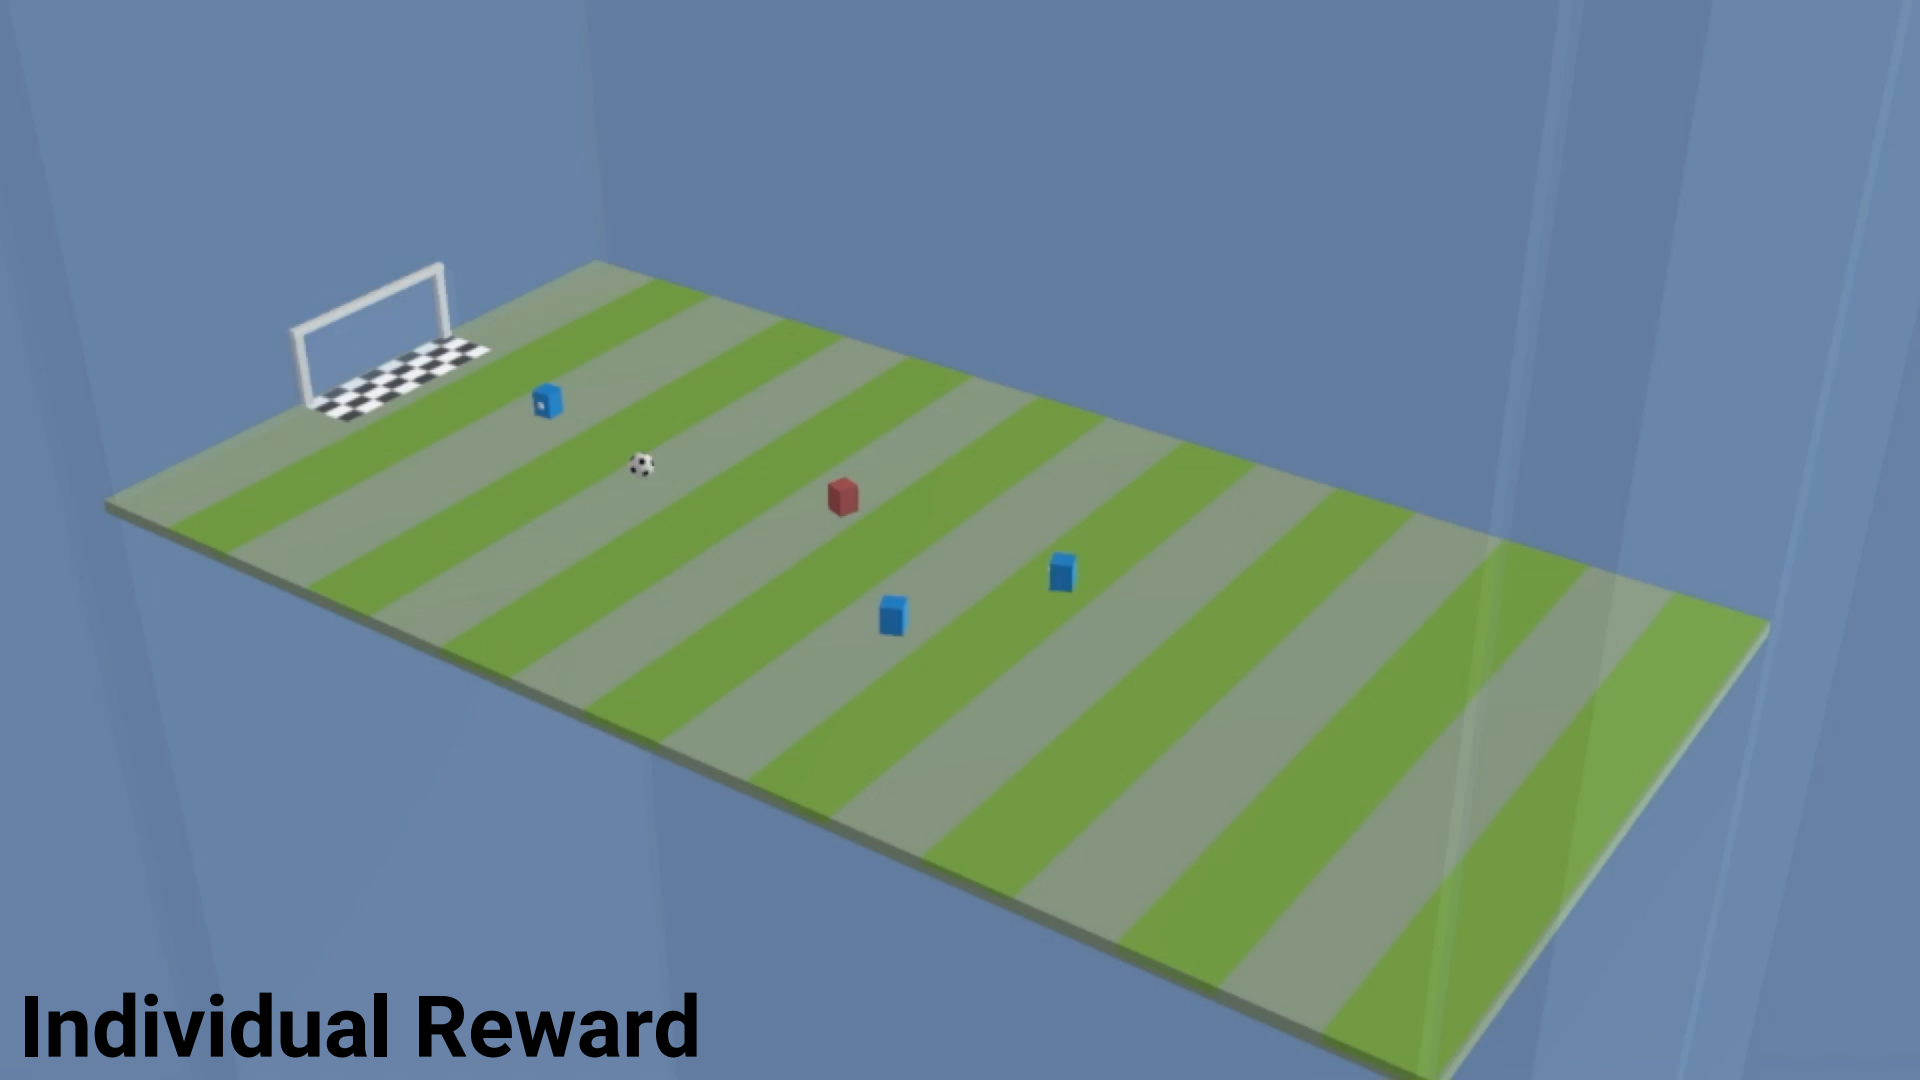
Combined
Train both brains together to play against opponent team



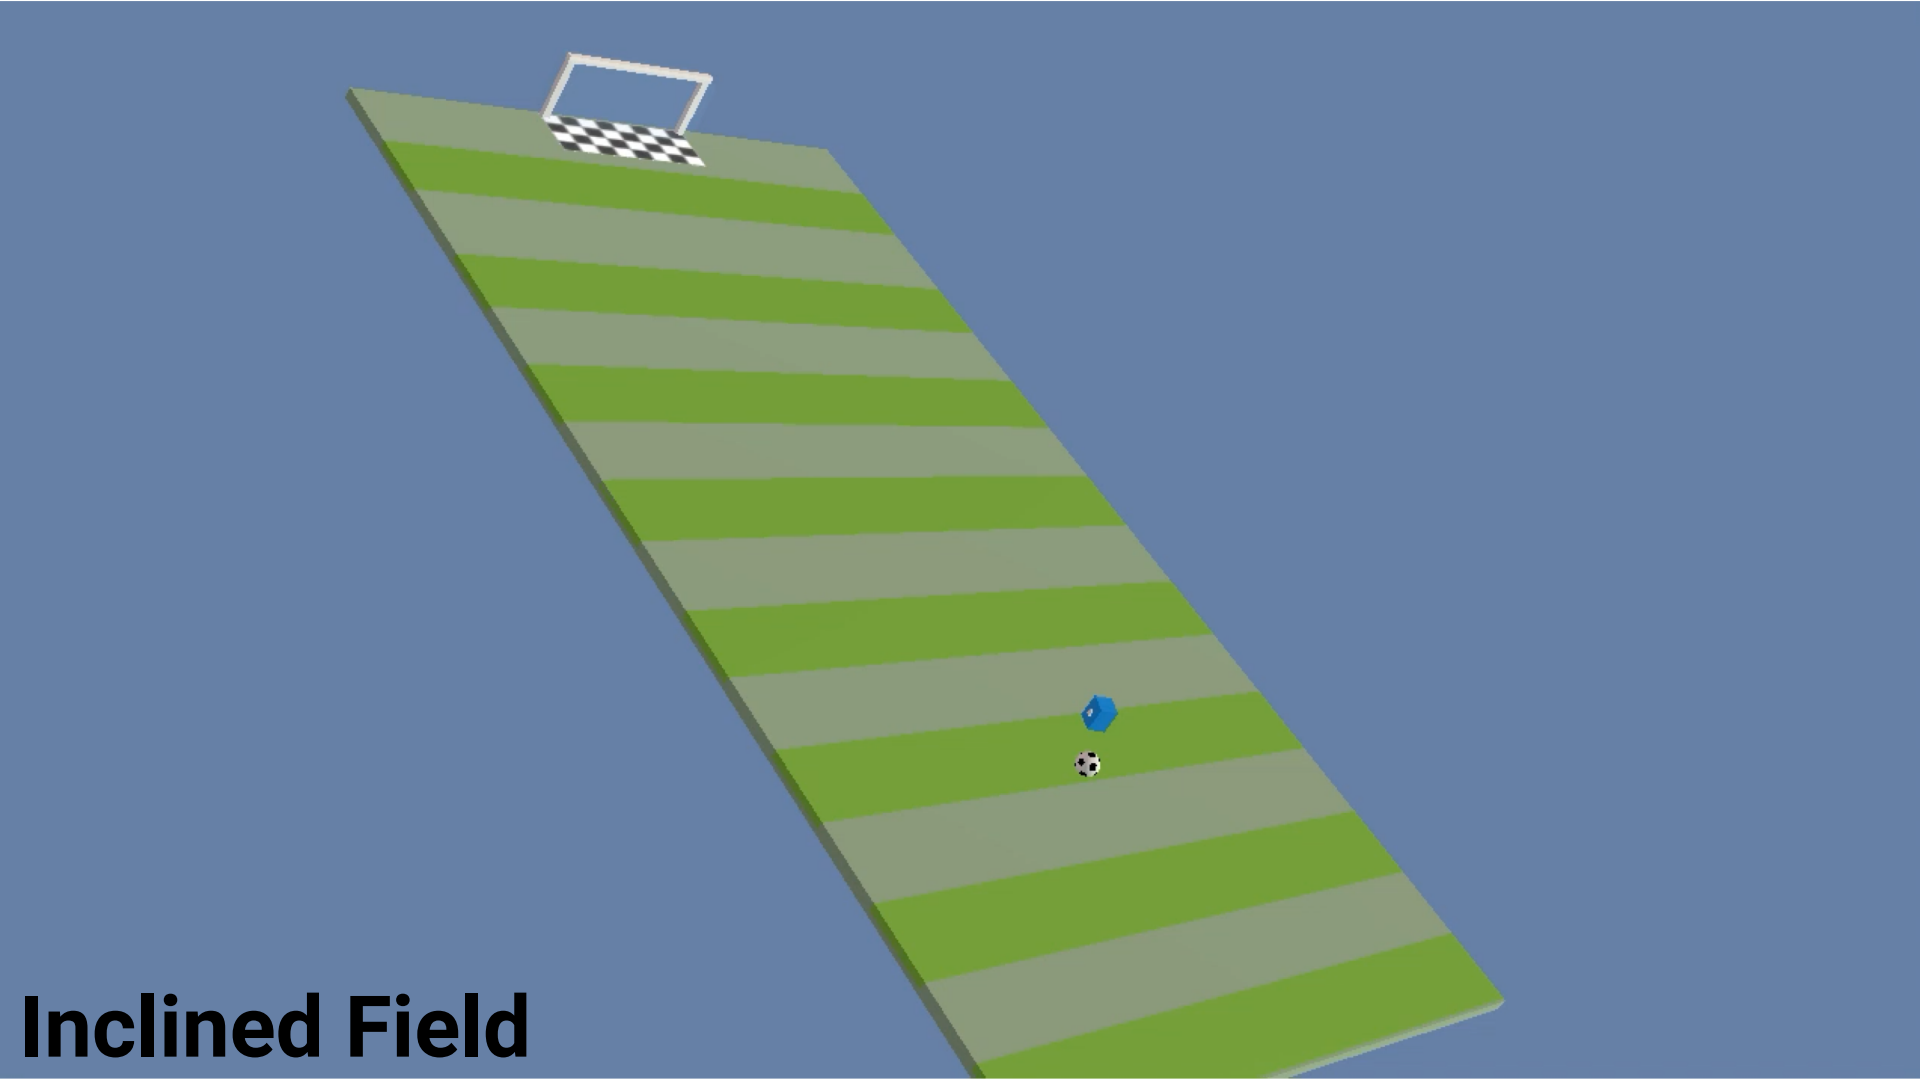


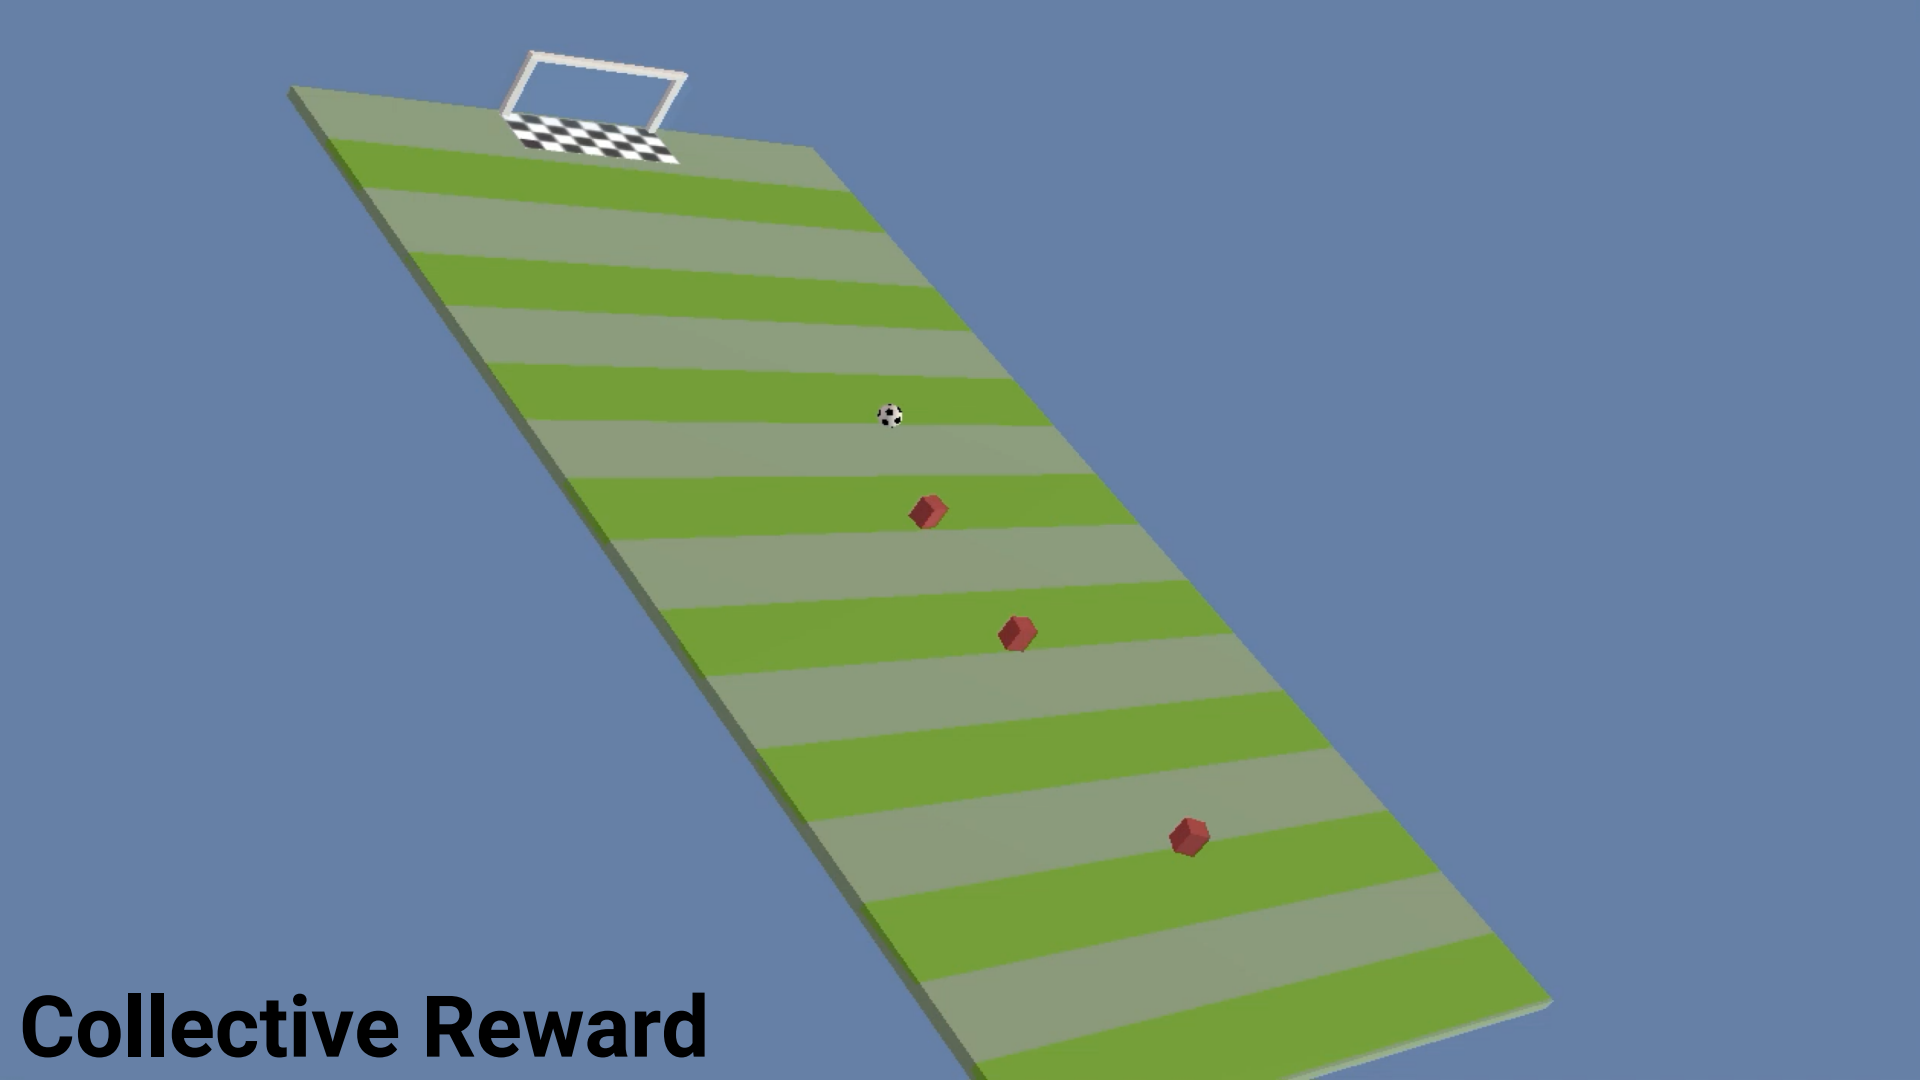
More Soccer-playing Agents

- Single brain with one or more agents
- Blue Agent: Free to move around
- Red Agent: Two seconds penalty after kick (cost)
- Individual Reward versus Collective Reward
- Observations: Local ray-cast perception on objects



Individual Reward


Inclined Field


Collective Reward

Traits in Learning

Extrinsic and Intrinsic Rewards

Extrinsic Rewards

- Examples: Capture, achieve, collect,…
- Specific to the environment
- *“Getting Rich”*

Intrinsic Rewards

- Examples: Curiosity, (im)patience, happiness, love, empathy,…
- Specific to the agent
- *“Getting Happy”*

Limits to Standard Reinforcement Learning

Sparse Rewards Spaces

- Agatha Christie’s event chains
- Ridiculous improbable scenarios
- Agent → Rooms → Button → Pyramid → Tumble → Collect

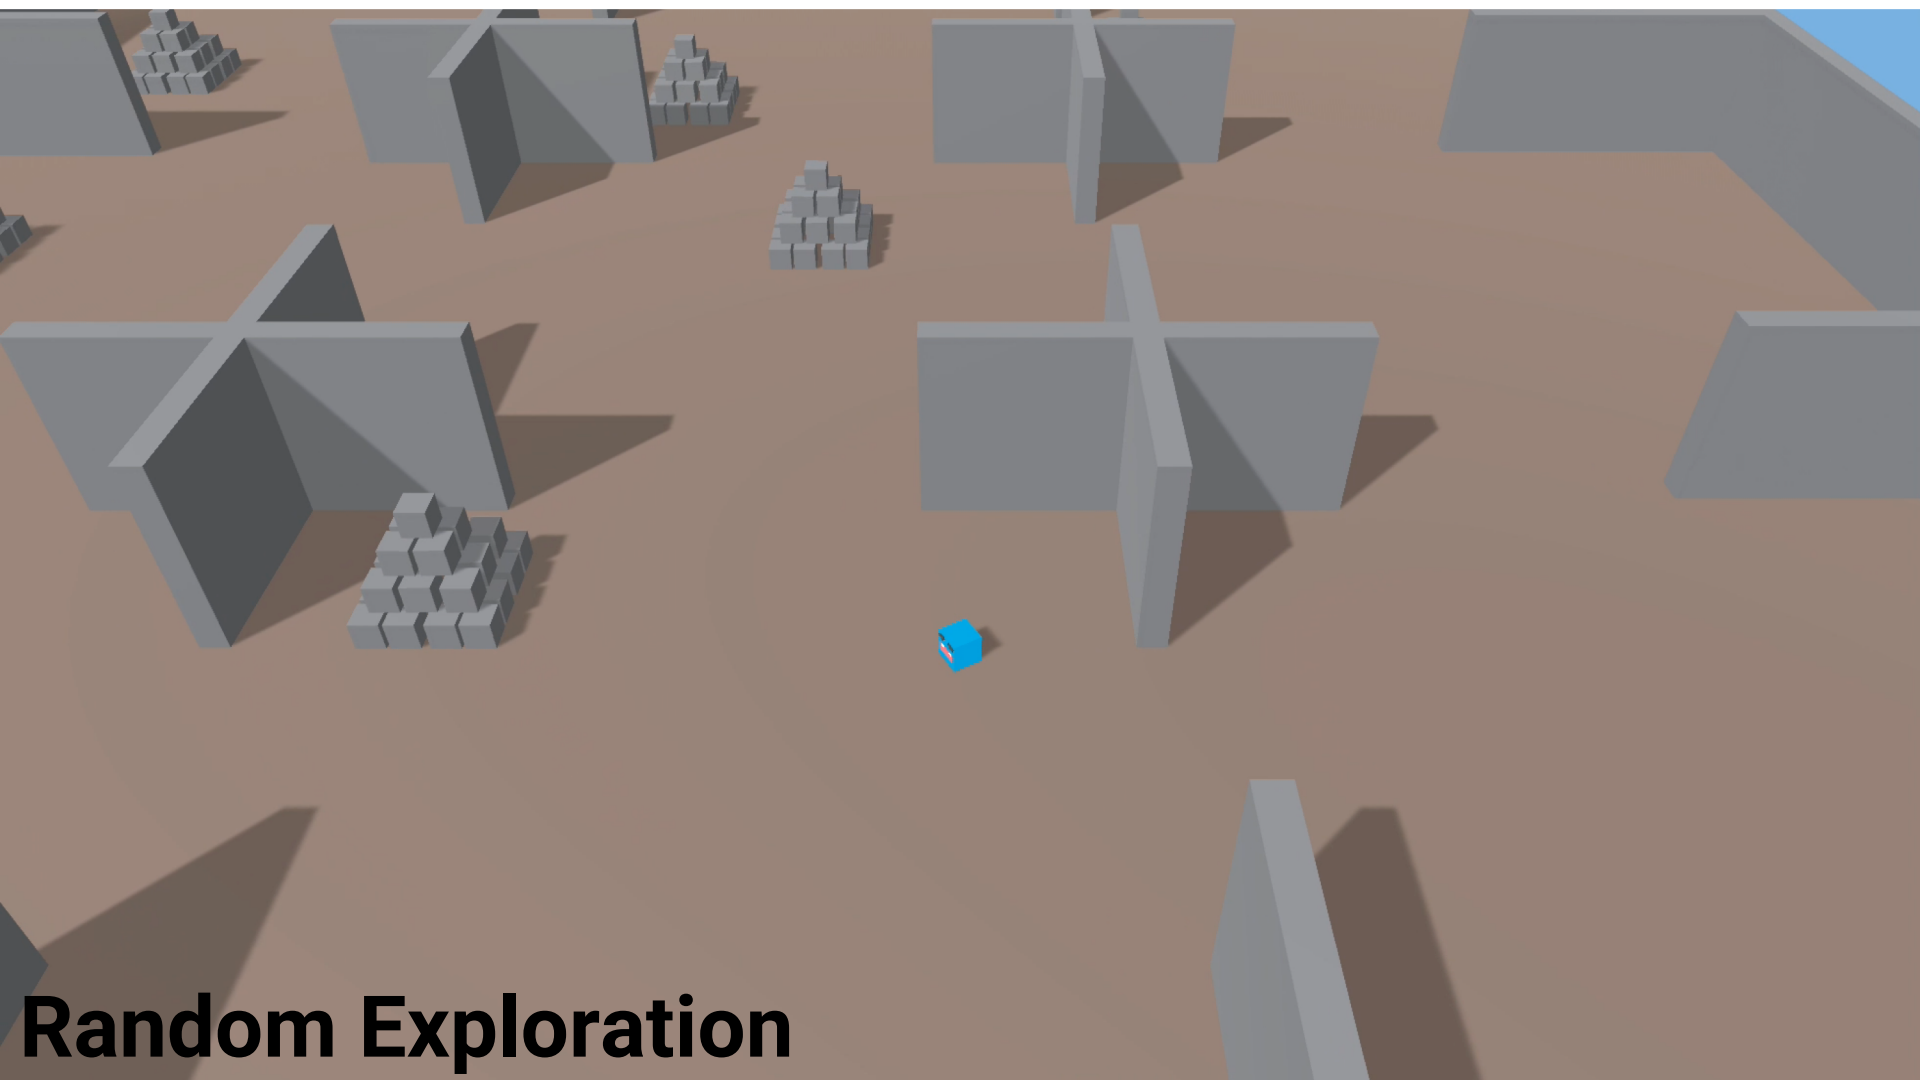
Solution is to add Strategy to Exploration

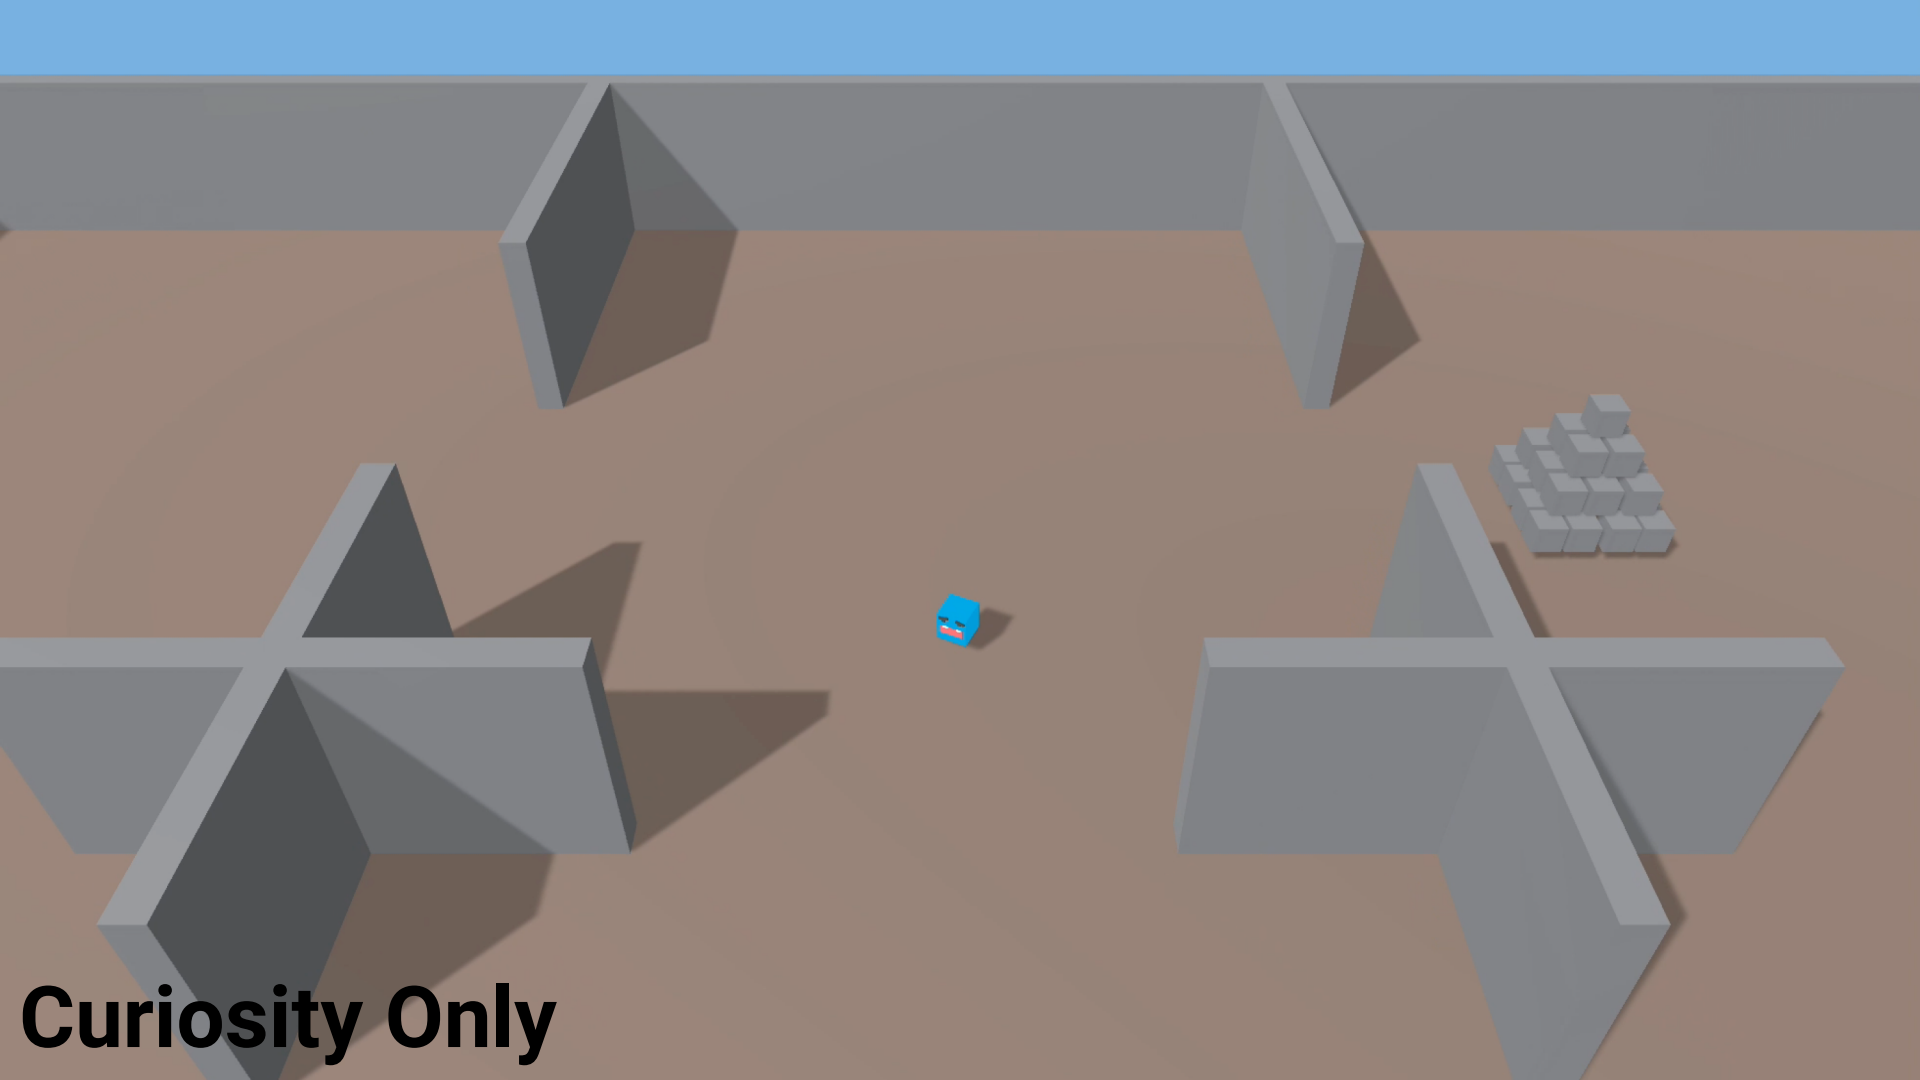
- Favor Agency over Randomness
- Key Innovation: Intrinsic Rewards

The Quest for Surprisal: Curiosity (in Math)

- Observations x_t and x_{t+1}
- *Action* a_t such that x_t transitions to x_{t+1}
- Embedding $\phi(x)$
- Prediction $p(\phi(x_{t+1}) \mid x_t, a_t)$
- **Reward $r_t = -\log p(\phi(x_{t+1}) \mid x_t, a_t$**
- Train to maximize r_t
- Agent now favors transitions with high prediction error




Random Exploration


Curiosity Only


Curiosity + Extrinsic Reward

Make an NPC Learn (Non–playable Character)


Welcome Puppo

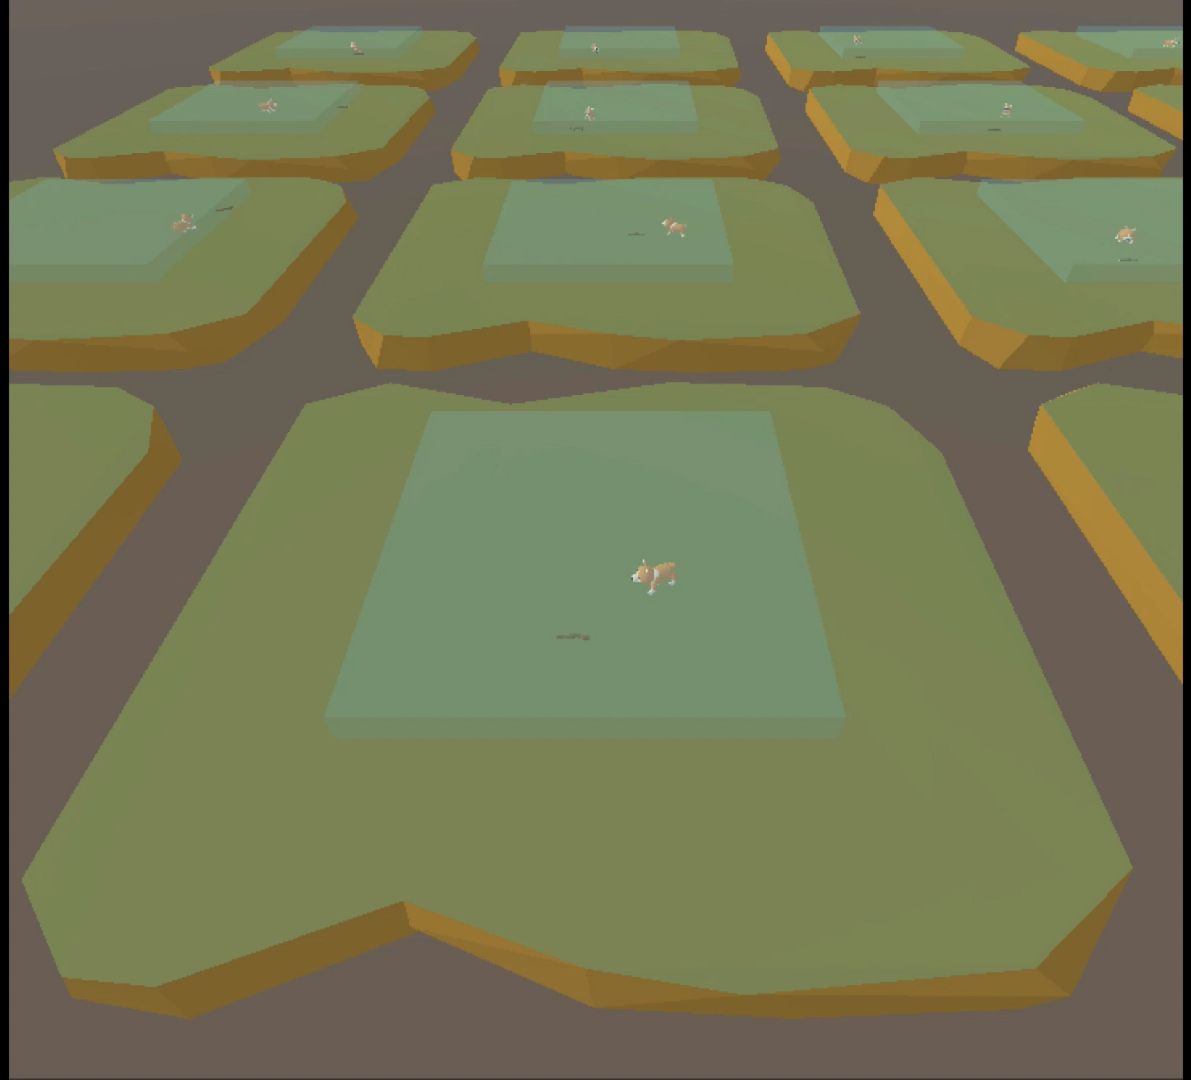
The Good Puppy, Bad Puppy Method

Stop coding and start training your NPC (game character)

- Learn to Walk, Run, Turn, Jump, and Fetch
- Spatial environment
- Physics engine
- Reinforcement Learning
- The Reward: Return the stick!



unity




PUPPO
THE CORGI!
START


PUPPO
THE CORGI!
START


1 DOG
10 DOGS
50 DOGS
BACK

Make a Virtual Player Learn (Play a Playable Character)

Virtual Game-playing

Virtual Game Player Benefits

- What you see is what you play
- Highly scalable – not a human
- Reinforcement Learning
- Continuous testing
- Game balancing in A | B Testing



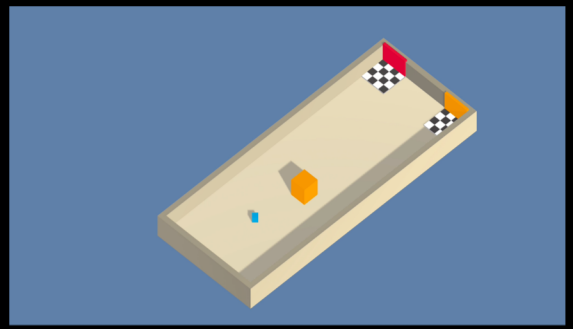

0/12
0
1
2
3
4
5
6

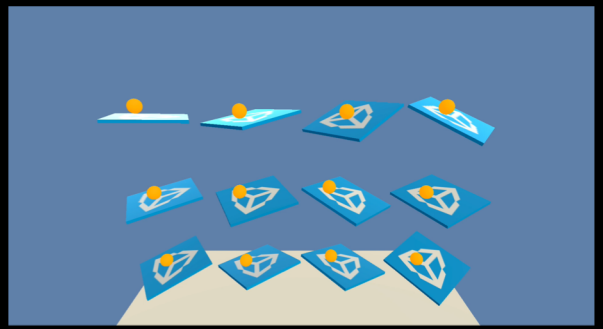

Unity ML-Agents

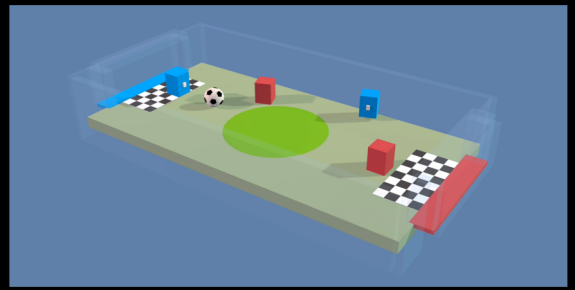

#1 in Research and Development
of Reinforcement Learning

ML-Agents Training Environments

- Basic
- 3D Balance Ball
- Grid World
- Tennis
- Push Block
- Wall Jump
- Reacher
- Crawler
- Banana Collector
- Hallway
- Bouncer
- Soccer Twos









Get ML-Agents at GitHub Now

github.com/Unity-Technologies/ml-agents

Contact Us

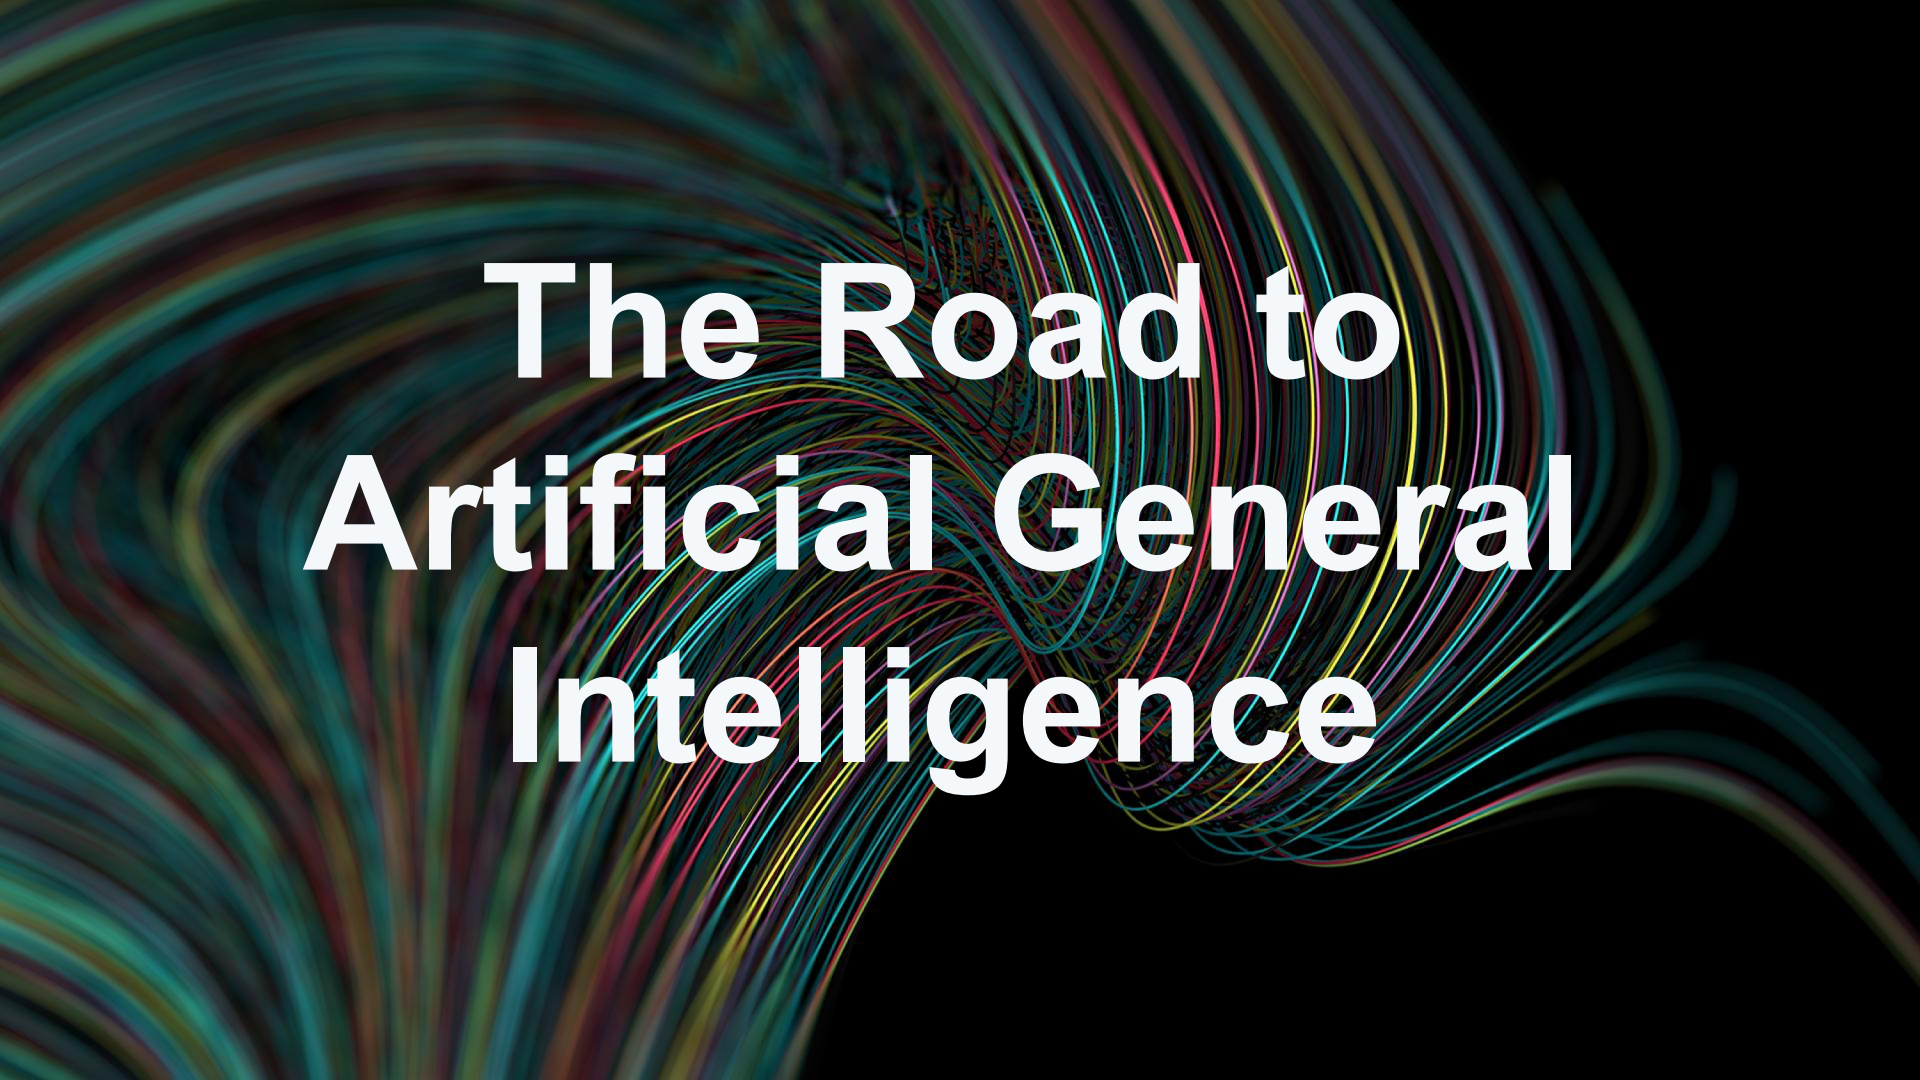
ML-Agents@Unity3d.com

ML researchers


http://unity3d.com/OTC
unity
Obstacle
Tower Challenge

Unity: A General Platform for Intelligent Agents
https://arxiv.org/abs/1809.02627

Obstacle Tower: A Generalization Challenge in Vision, Control, and Planning
https://arxiv.org/abs/1902.01378


The Road to
Artificial General
Intelligence

The Road to Artificial General Intelligence

Biology has a great trip in store for us so let’s start here

- Attention: Representation Learning & Deep Embeddings
- Episodic Memory: One-shot Learning
- Working Memory: Long Short-term Memory
- Continuous Learning: Elastic Weight Consolidation
- Imagination: Simulation & Monte Carlo Tree Search
- Agency: Intrinsic Values
- Decomposition: Hierarchical Learning

From Biological to Cultural Evolution

Human Culture has Evolved at its own timescale

- 200-300,000 Years Ago: Homo Sapiens
- 100,000 Sporadic Migration (failed)
- 45-70,000 Global Migration (Australia)
- 30-70,000 Cognitive Revolution
- But no Moore's Law for the Human Brain with stable DNA for the last 100,000 years
- Emergence, Language, Collaboration, and Anticipation




@danny_lange
dannylange
Thank You
unity3d.ai

goto; chicago



Please

Remember to rate this session

Thank you!

follow us @gotochgo

goto; chicago





Click ‘Rate Session’

Rate **5** sessions to get the supercool GOTO reward

follow us @gotochgo