

The Impact of Simulation on Advances in Reinforcement Learning

Danny Lange
VP of AI, Unity Technologies



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 'teləjəns*] –
the ability to acquire and
apply knowledge and skills.**

What is *Real* Intelligence?




**What is the only Real Intelligence
you Truly Know?**



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



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.




Unity ML-Agents

Practice Deep RL with Unity ML-Agents



Build





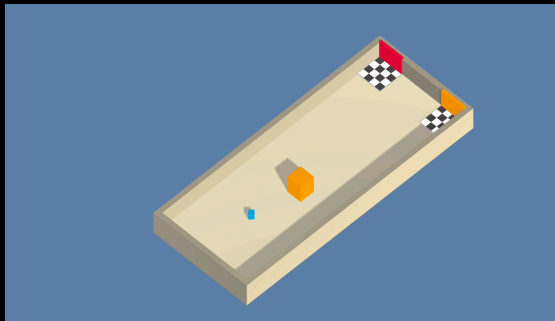
Train



Embed

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





Learning Scenarios

Control Learning

The background is black with abstract, colorful, curved lines in the corners. The lines are composed of many thin, overlapping lines in various colors including red, green, blue, yellow, and purple, creating a sense of motion and depth. The lines curve from the corners towards the center of the page.








Curriculum Learning


The image features a black background with the text "Curriculum Learning" centered in a white, sans-serif font. In the top-right and bottom-left corners, there are decorative elements consisting of numerous thin, overlapping lines in various colors (including red, green, blue, and yellow) that curve and fan out, creating a sense of motion or data flow.




Easy

Difficult

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

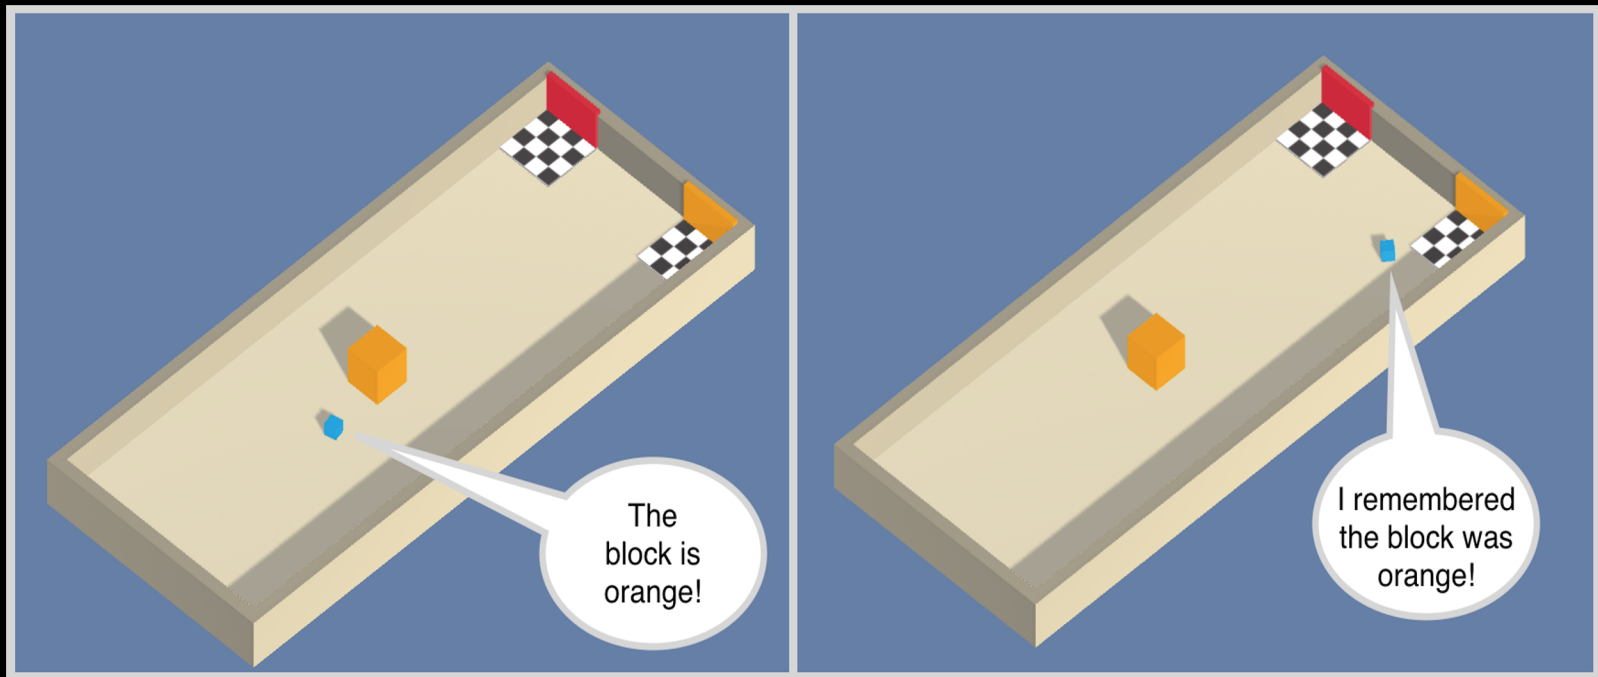





Memory in Learning

The image features a solid black background. In the top-right and bottom-left corners, there are decorative elements consisting of numerous thin, overlapping lines in various colors (including red, green, blue, yellow, and purple) that curve and fan out, creating a sense of motion or data flow.

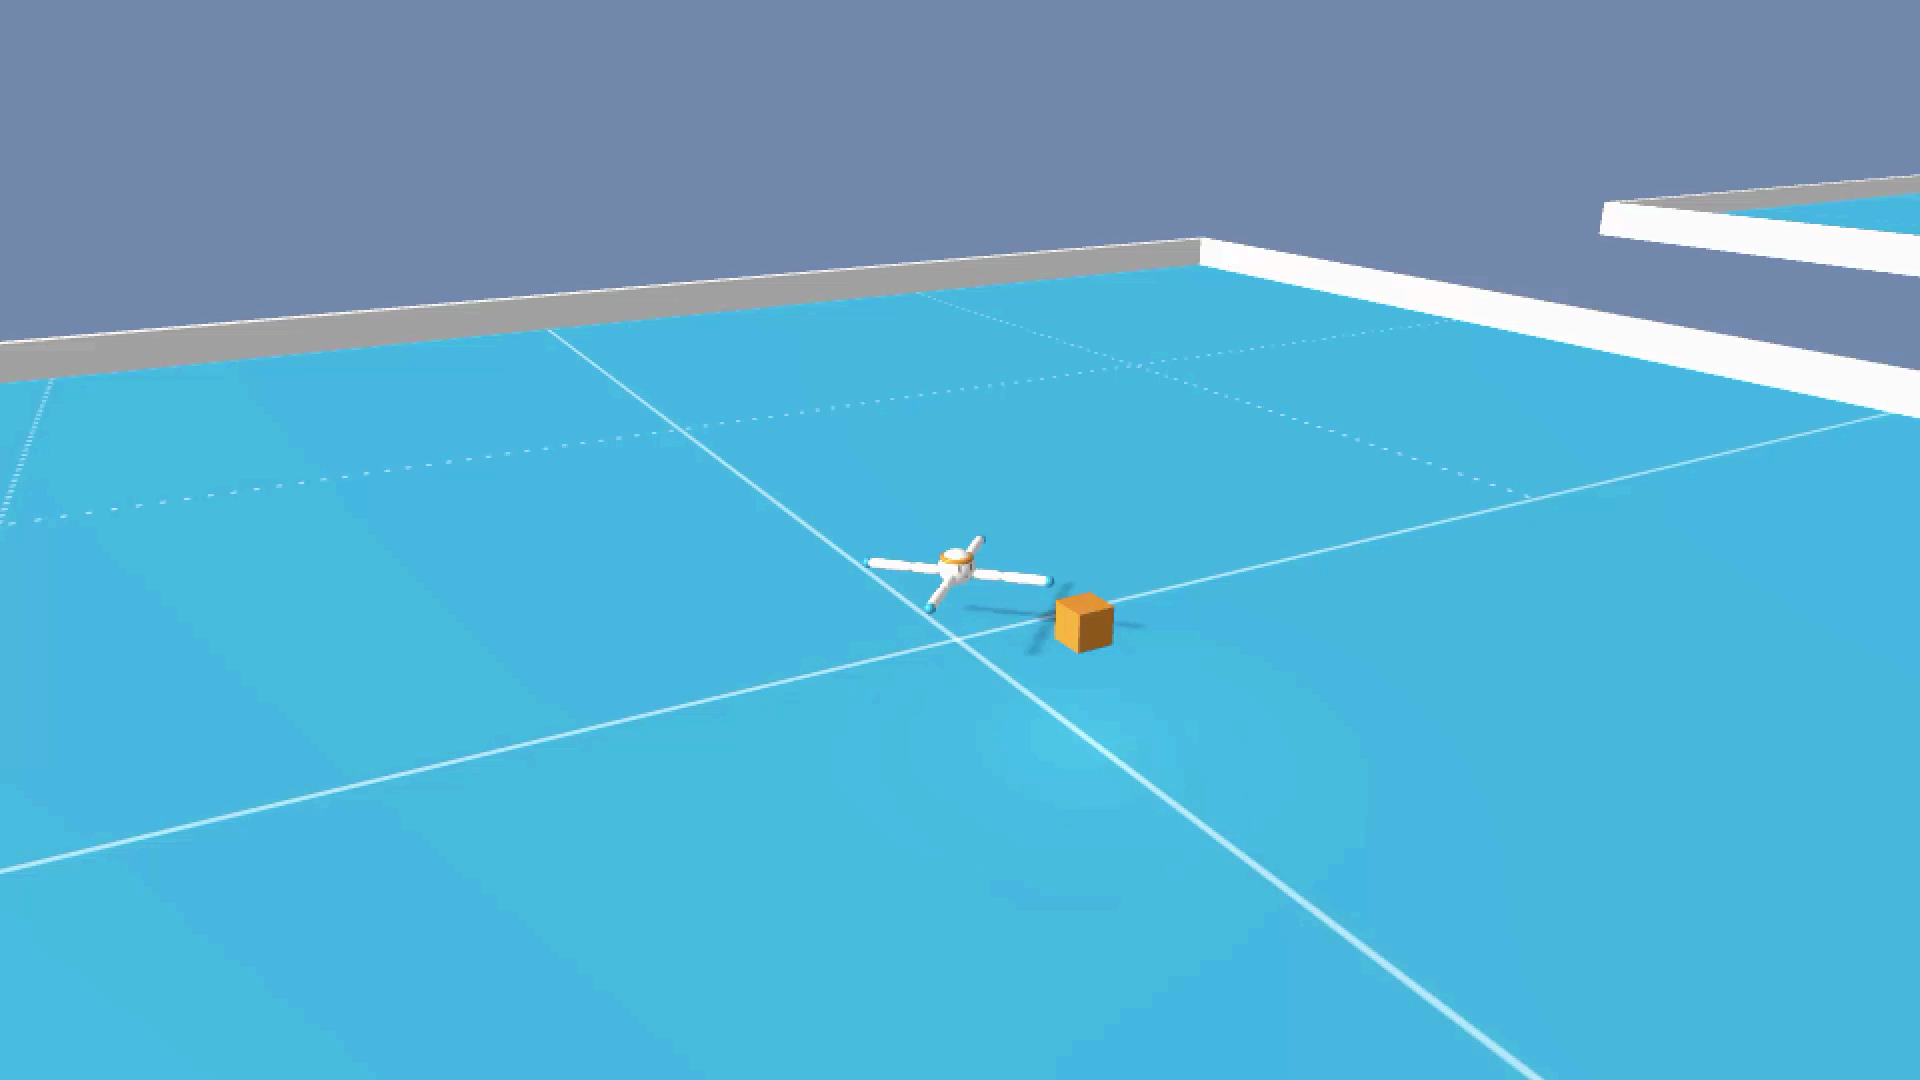
Enabling Long Short-Term Memory (LSTM)






The background is black with abstract, multi-colored lines in shades of blue, green, and purple, curving from the corners towards the center.

Hierarchical Learning Navigation + Control



Multi-agent Learning

The image features a black background with the text "Multi-agent Learning" centered in a white, bold, sans-serif font. In the top-right and bottom-left corners, there are decorative elements consisting of numerous thin, overlapping lines in various colors (including red, green, blue, and yellow) that curve and fan out, creating a dynamic, abstract effect.



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

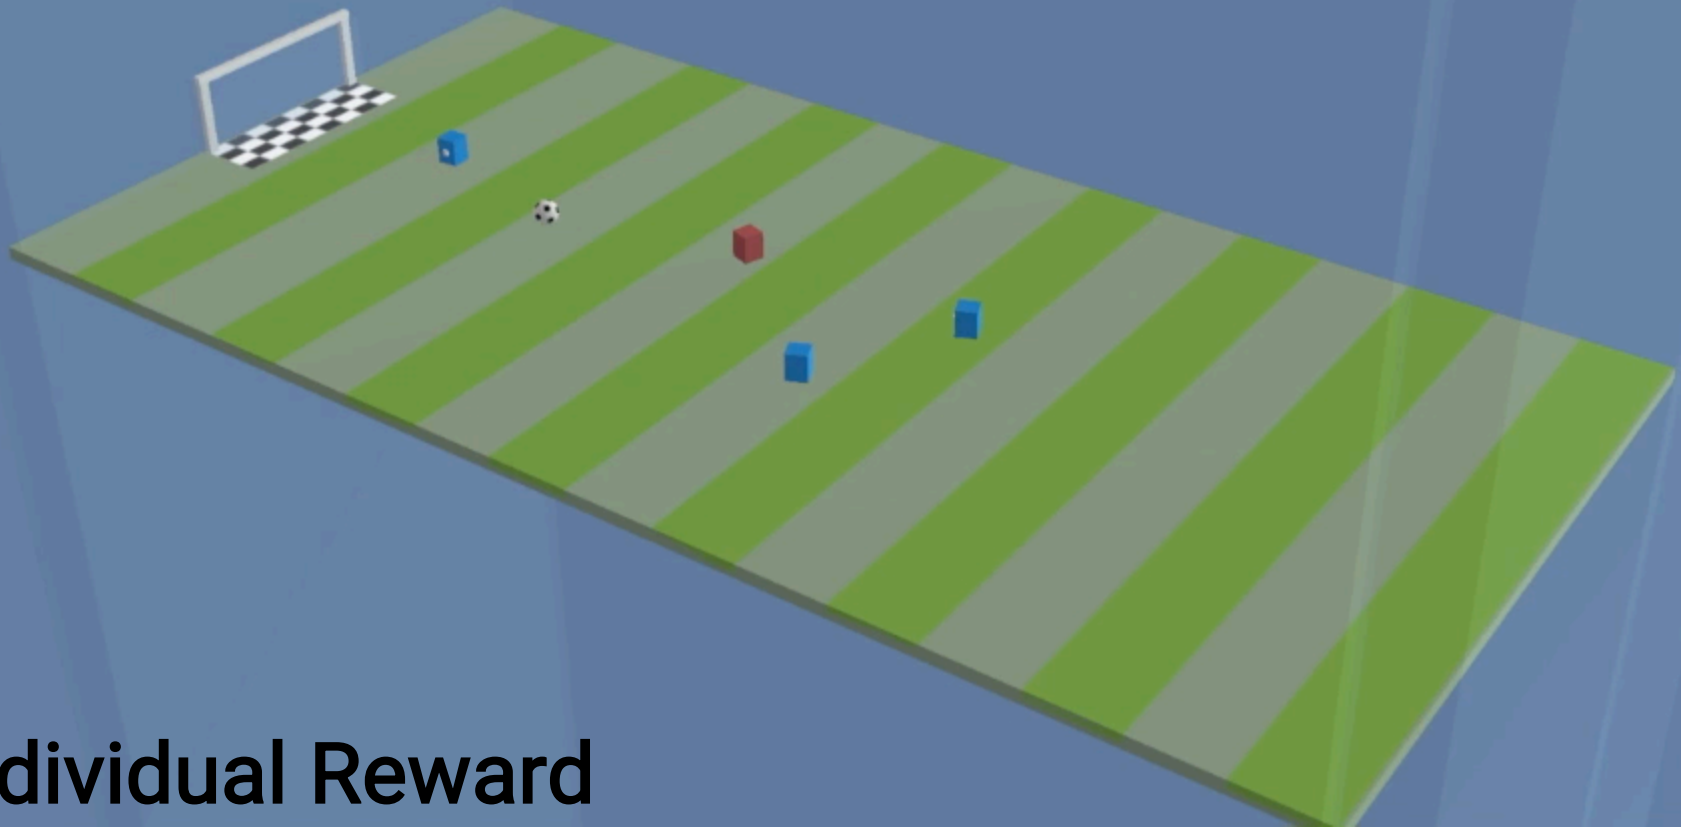
Train both brains
together to play
against opponent
team






More Soccer-playing Agents

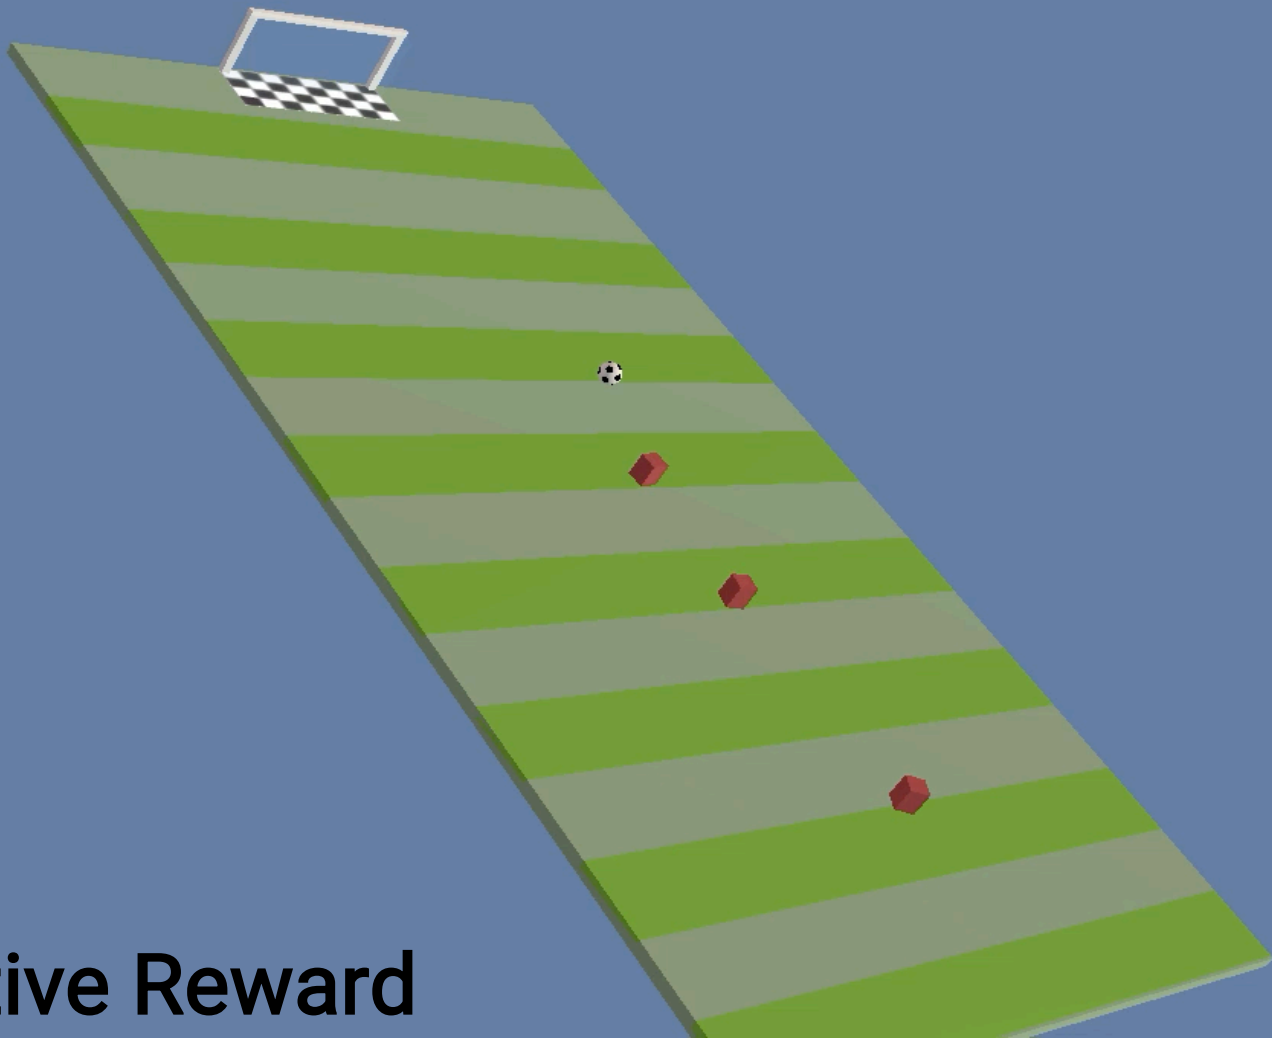
- Blue Agent: Free to move around
- Red Agent: Two seconds penalty after kick (cost)
- Individual Reward versus Collective Reward



Individual Reward



Inclined Field



Collective Reward

Traits in Learning

The background is black with abstract, multi-colored lines in the corners. The lines are thin and densely packed, creating a sense of motion and energy. The colors include shades of blue, green, yellow, and red. The lines in the top right corner curve downwards and to the left, while the lines in the bottom left corner curve upwards and to the right.

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

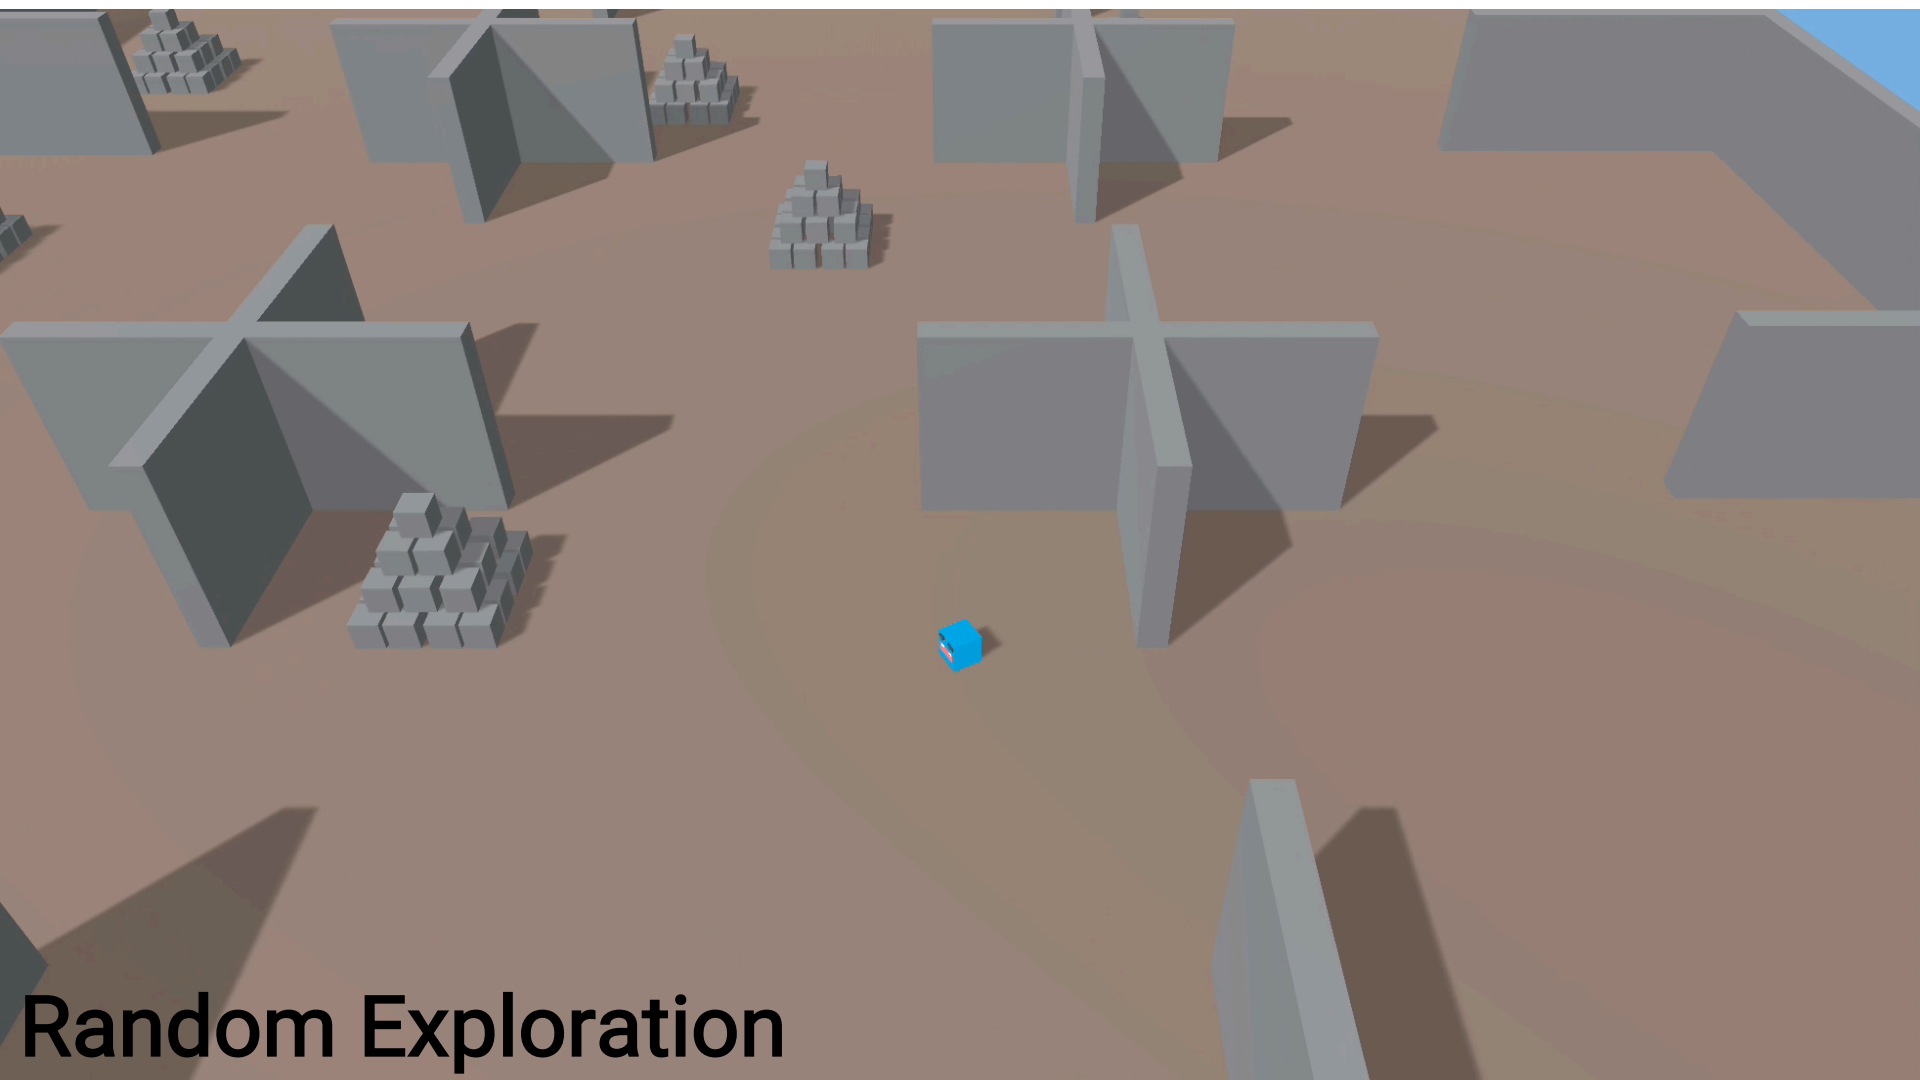
- Agatha Christie's event chains
- Ridiculous improbable scenarios
- Agent → Rooms → Button → Pyramid → Tumble → Collect

Solution is to add Strategy to Exploration


- Favor Agency over Randomness
- Key Innovation: Use Intrinsic Rewards (Traits)

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 $p(\phi(x_{t+1}) | x_t, a_t)$
- **Reward $r_t = -\log p(\phi(x_{t+1}) | x_t, a_t)$** 
- Train to maximize r_t
- Agent now favors transitions with high prediction error



Random Exploration



Curiosity Only



Curiosity + Extrinsic Reward

The TV Trap



The background is black with abstract, multi-colored lines in shades of blue, green, and purple, curving from the corners towards the center.

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
- 3D plus Physics
- Reinforcement Learning
- The Reward: Return the stick!





PUPPO THE CORGI!

START





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

Contact Us
ML-Agents@Unity3d.com

@danny_lange 

dannylange 

Thank You

unity3d.ai