


On the Road to Artificial General Intelligence

Danny Lange






**Click 'Rate Session'
to rate session
and ask questions.**




On the Road to Artificial General Intelligence


Danny Lange

VP of AI and Machine Learning, Unity





UBER



About Unity Technologies

- Real-time 3D Development Platform
 - Game, AR/VR, film, automotive
- 60% of top 1,000 games
- On over 3 billion unique devices
- 1.5 billion monthly active players
- 2,000 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 '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.



Introducing **Unity ML-Agents** Your Private AI Biodome

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




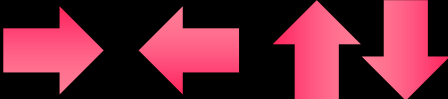
Explore



Exploit

Chicken Crossing the Road: Tabula Rasa



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



Unity ML-Agents

The #1 Development and Research
for Reinforcement Learning

Unity ML-Agents Workflow



Unity ML-Agents Workflow




Unity ML-Agents Workflow




The background of the slide is a dynamic, abstract composition of numerous thin, curved lines. These lines, in shades of teal, blue, green, and yellow, flow from the left side towards the right, creating a sense of movement and depth. The lines are densely packed in some areas and more sparse in others, all set against a solid black background.

Learning Scenarios



2 Agents | 1 Brain - Cooperative Rewards


Goal	Keep ball up as long as possible
Observations	Positions and velocities of racket and ball
Actions	Forward, backward, and upward movement
Rewards	+0.1 when ball sent over net by agent -0.1 when ball falls because of agent



Multi-Stage Soccer Camp

Offense

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



Defense


Train one brain with negative reward for ball entering own goal




Combined

Train both brains together to play against opponent team








Easy


Difficult

Curriculum Learning


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




Curriculum Learning: Progress




Enabling Long Short-Term Memory (LSTM)





The background is black with abstract, multi-colored lines (red, green, blue, yellow) forming a curved, brushstroke-like shape in the top right and bottom left corners.

Control Learning Humanoid Robot



The background is black with abstract, multi-colored lines (red, green, blue, yellow) in the top-right and bottom-left corners, resembling a stylized 'S' or 'Z' shape.

Agency-driven Exploration

Extrinsic + Intrinsic Rewards

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




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!





PUPPO THE CORGI!

START



PUPPO THE CORGI!

START

trial ver.



A stylized, low-poly illustration of a Corgi dog sitting on a purple track. The track has yellow lines and curves through a green landscape with hills in the background. In the foreground, a blue pedestal holds a large, white, pixelated bone. Three green boxes with white text are on the left, and a yellow box with white text is in the bottom right.

1 DOG




10 DOGS

50 DOGS

BACK

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

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

@danny_lange 

dannylange 

Thank You
unity3d.ai



Please

**Remember to
rate this session**

Thank you!



Did you **remember**
to rate the previous
session ?

