Lecture 1: Introduction

Reinforcement Learning with TensorFlow&OpenAI Gym

Sung Kim <hunkim+ml@gmail.com>

What is Positive Reinforcement Dog Training?

- Teaching dogs desirable behaviors using SCIENCE-based & REWARD-based methods.
- Helping dogs learn and succeed step by step.
- Motivating dogs with fun exercises and games. No force! No pain!
- Encouraging dogs to think more for themselves.
- Valuing dogs' voluntary behaviors.
- Understanding dogs' feelings from their body language.
- Understanding how dogs learn, their needs and wants.
- Using methods that work humanely with ANY dog. Big dogs, small dogs, puppies, senior dogs, disabled dogs, fearful dogs, reactive dogs... can all learn and have fun!

1. develop dog's self-control

2. develop a trust relationship

3. develop dog's self-confidence

http://angelpawstherapy.org/positive-reinforcement-dog-training.html

Nature of Learning

- We learn from past experiences.
  - When an infant plays, waves its arms, or looks about, it has no explicit teacher
  - But it does have direct interaction to its environment.
- Years of positive compliments as well as negative criticism have all helped shape who we are today.
- Reinforcement learning: computational approach to learning from interaction.

Richard Sutton and Andrew Barto, Reinforcement Learning: An Introduction
Nishant Shukla , Machine Learning with TensorFlow

Reinforcement Learning

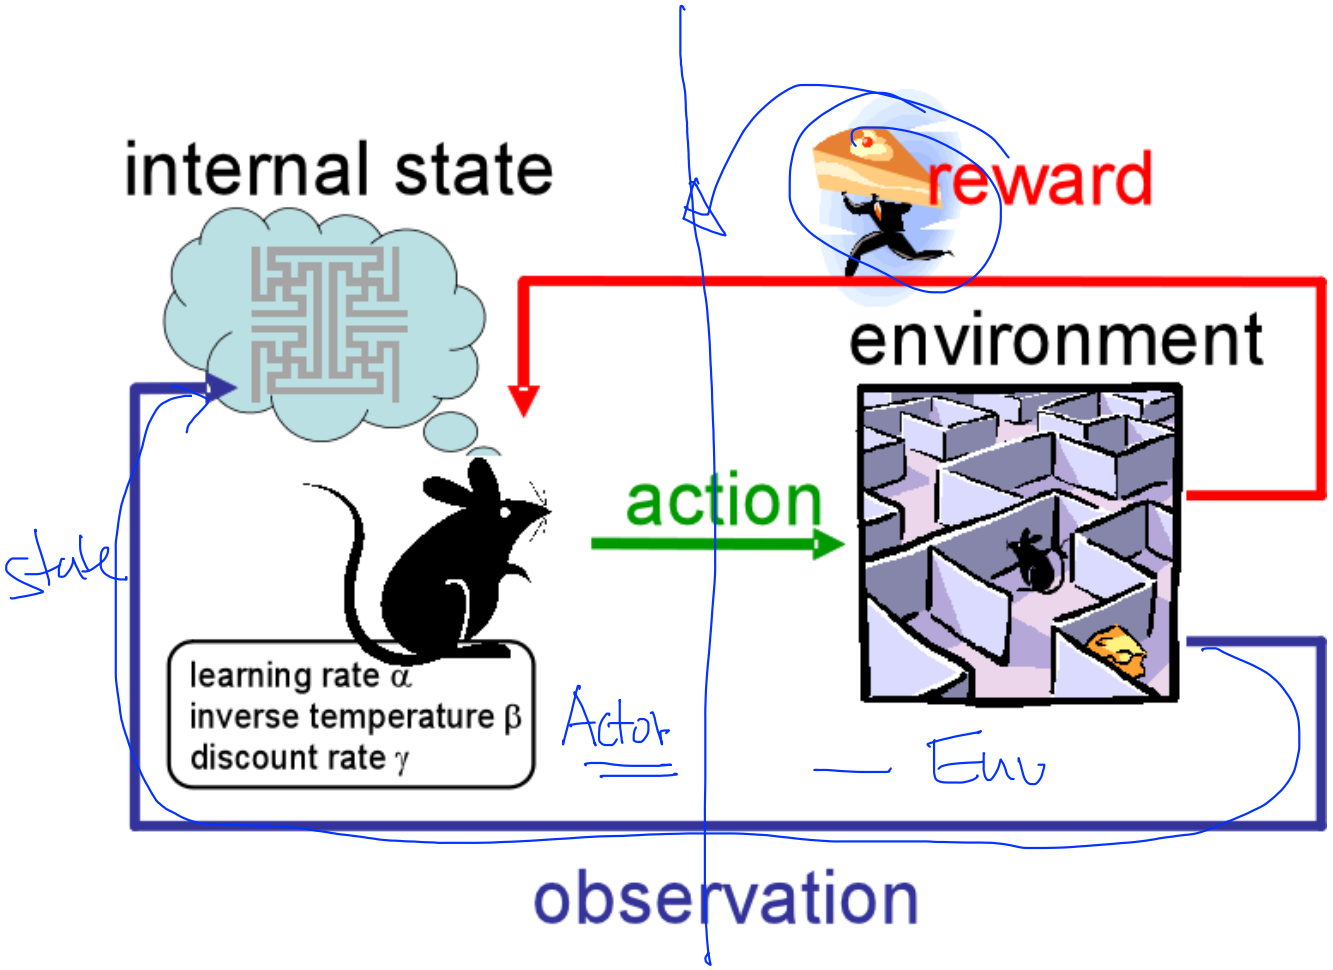

https://www.cs.utexas.edu/~eladlieb/RLRG.html

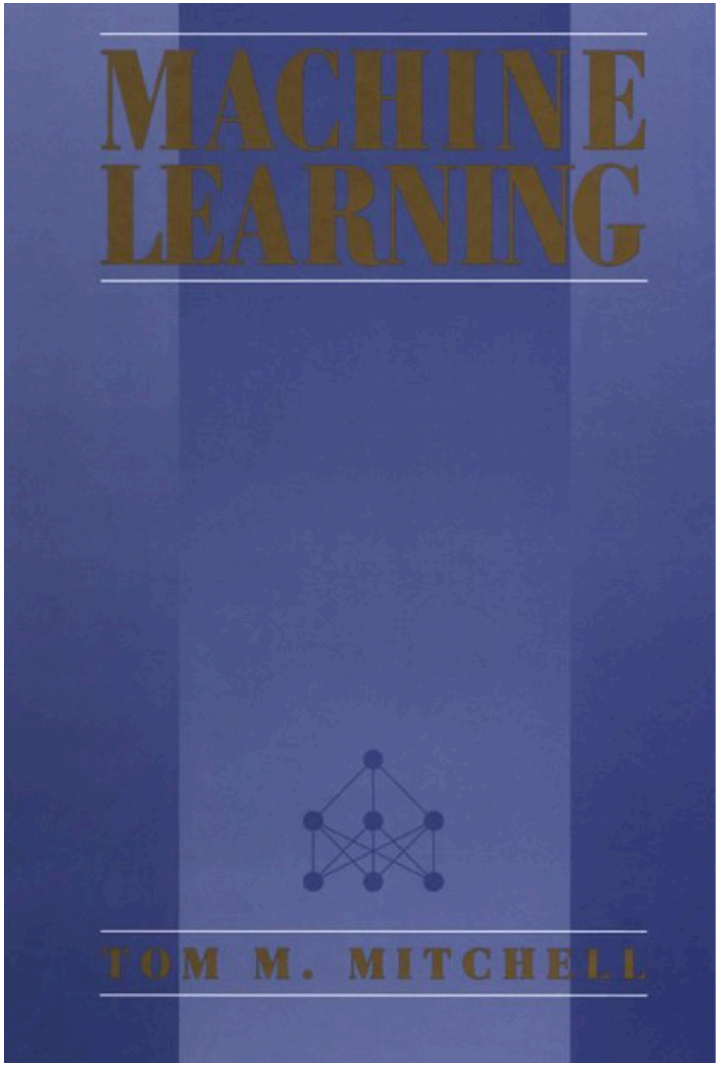

Machine Learning, Tom Mitchell, 1997

Atari Breakout Game (2013, 2015)

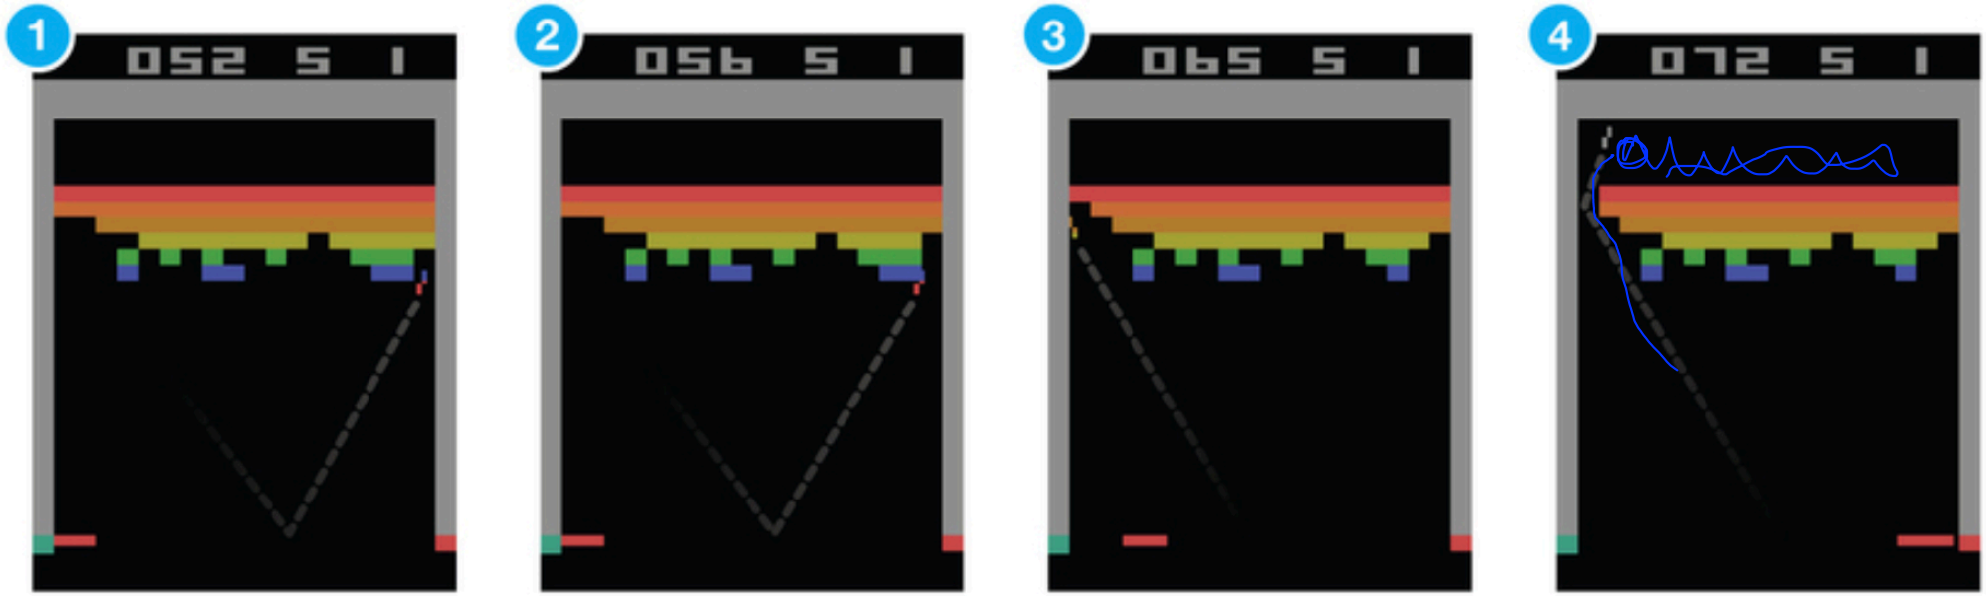

Figure 1: Atari Breakout game. Image credit: DeepMind.

Atari Games

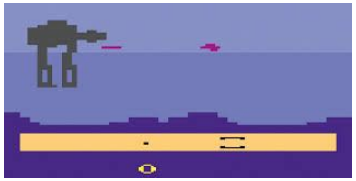

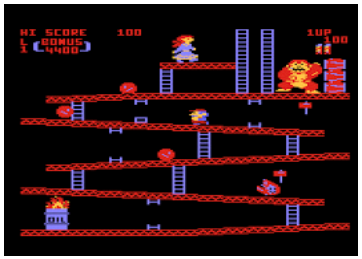

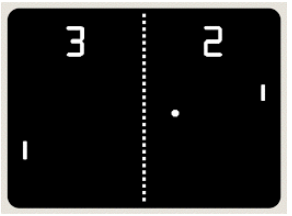

Nature : Human-level control through deep reinforcement learning

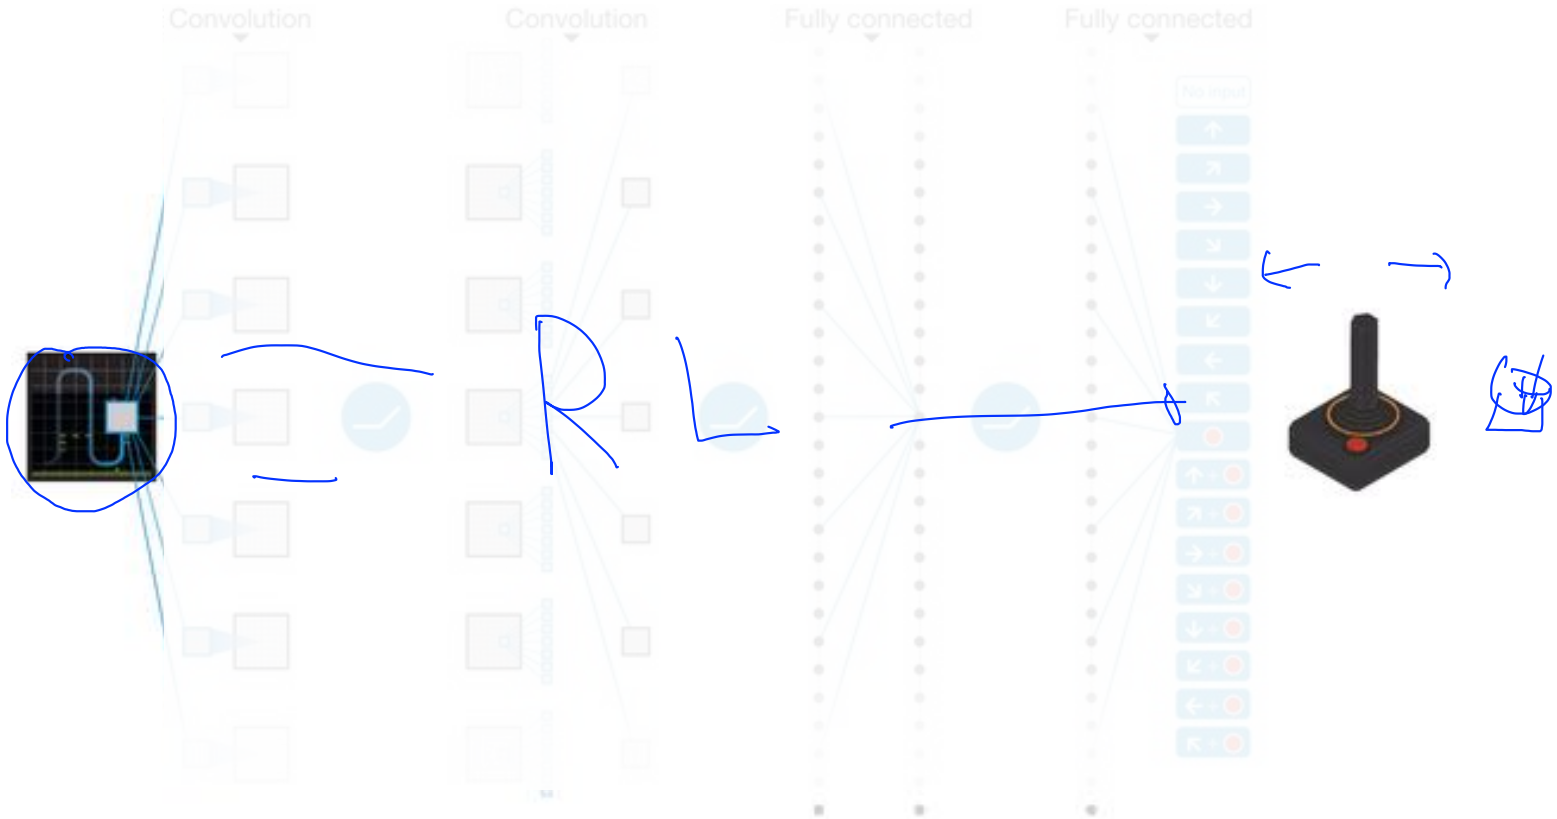

Human-level control through deep reinforcement learning, Nature
http://www.nature.com/nature/journal/v518/n7540/full/nature14236.html

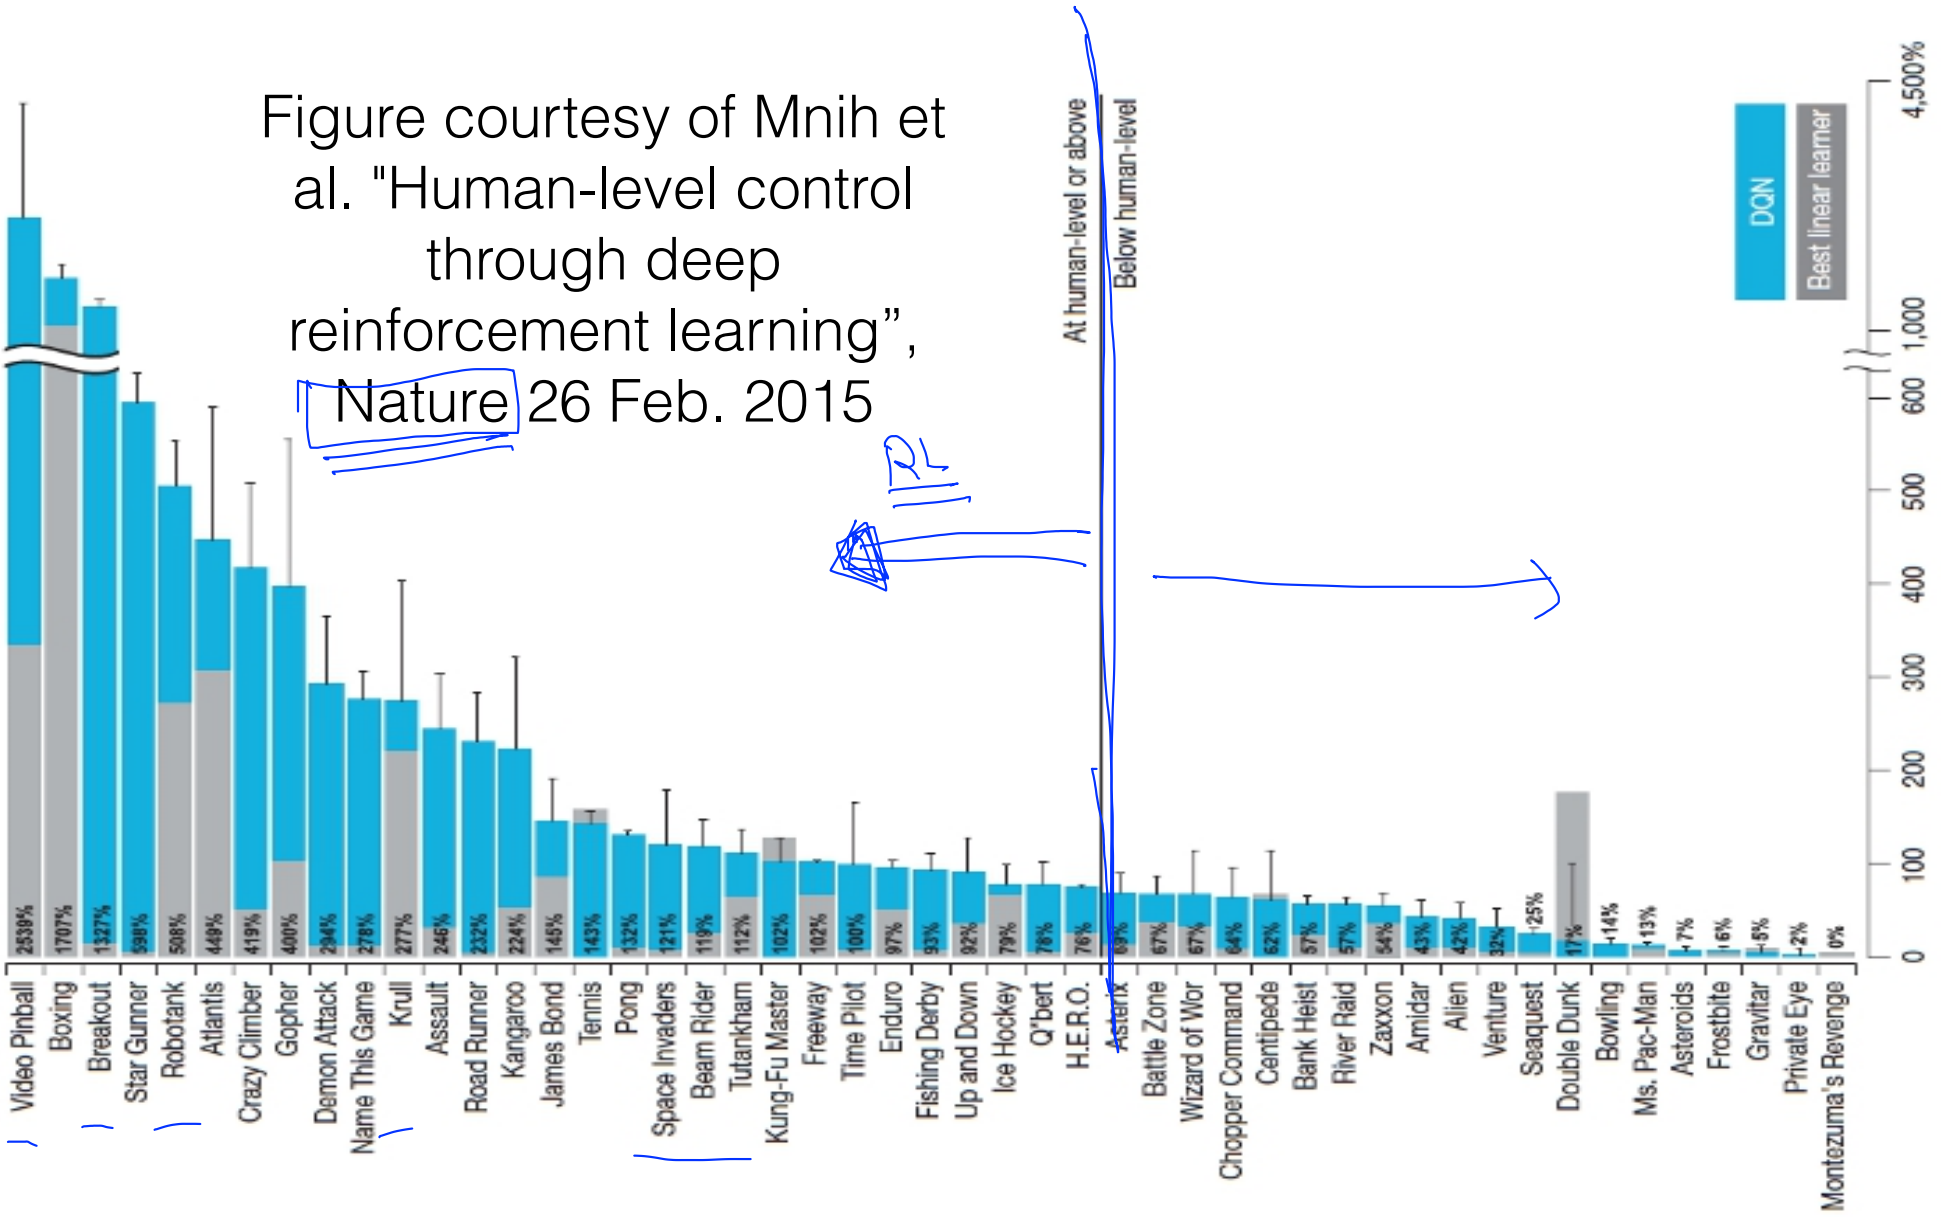

Figure courtesy of Mnih et al. "Human-level control through deep reinforcement learning", Nature 26 Feb. 2015


KBA
Korea Baduk Association
Google DeepMind
Challenge Match
8 - 15 March 2016
AlphaGo
RL
ALPHAGO
01:59:50
LEE SEDOL
AlphaGo
Lee Sedol
Google DeepMind
Challenge Match
https://deepmind.com/blog/deep-reinforcement-learning/

DeepMind AI Reduces Google Data Centre Cooling Bill by 40%

https://deepmind.com/applied/deepmind-for-google/

Reinforcement Learning Applications

- Robotics: torque at joints
- Business operations
  - Inventory management: how much to purchase of inventory, spare parts
  - Resource allocation: e.g. in call center, who to service first
- Finance: Investment decisions, portfolio design
- E-commerce/media
  - What content to present to users (using click-through / visit time as reward)
  - What ads to present to users (avoiding ad fatigue)

Audience

- Want to understand basic reinforcement learning (RL)
- No/weak math/computer science background
  - Q = r + Q
- Want to use RL as black-box with basic understanding
- Want to use TensorFlow and Python (optional labs)

Schedule

1. Introduction
2. Playing Games, OpenAI Gym Introduction & Lab
3. Q-learning with Tables & Lab
4. Q learning on nondeterministic Rewards and Actions & Lab
5. Q-learning with Networks (DQN) & Lab
6. Policy Gradients & Lab
7. Further Topics

References

- Awesome Reinforcement Learning http://aikorea.org/awesome-rl/
- Simple Reinforcement Learning with TensorFlow, https://medium.com/emergent-future/
- http://kvfrans.com/simple-algoritms-for-solving-cartpole/ (written by a high school student)
- Deep Reinforcement Learning: Pong from Pixels - Andrej Karpathy blog http://karpathy.github.io/2016/05/31/rl/
- Machine Learning, Tom Mitchell, 1997
- CS 294: Deep Reinforcement Learning, Spring 2017, http://rll.berkeley.edu/
- Fundamental of Reinforcement Learning, https://www.gitbook.com/book/dnddnjs/rl/details (Korean Book)

Online video lectures

- A Tutorial on Reinforcement Learning, https://simons.berkeley.edu/talks/tutorial-reinforcement-learning 2017
- Berkeley CS 294: Deep Reinforcement Learning, Spring 2017 http://rll.berkeley.edu/deeprlcourse/, 2017
- MIT 6.S094: Deep Learning for Self-Driving Cars (Lecture 2) http://selfdrivingcars.mit.edu/, 2017
- Deep Reinforcement Learning (John Schulman, OpenAI) https://www.youtube.com/watch?v=PtAIh9KSnjo&t=2457s (summary) and https://www.youtube.com/watch?v=aUrX-rP_ss4&list=PLjKEIQlKCTZYN3CYBlj8r58SbNorobqcp (4 lectures)
- UCL, David Silver, Reinforcement Learning http://www0.cs.ucl.ac.uk/staff/d.silver/web/Teaching.html, 2015
- Stanford Andrew Ng CS229 Lecture 16 https://www.youtube.com/watch?v=RtxI449ZjSc, 2008

Prerequisite: http://hunkim.github.io/ml/ or https://www.inflearn.com/course/기본적인-머신러닝-딥러닝-강좌/

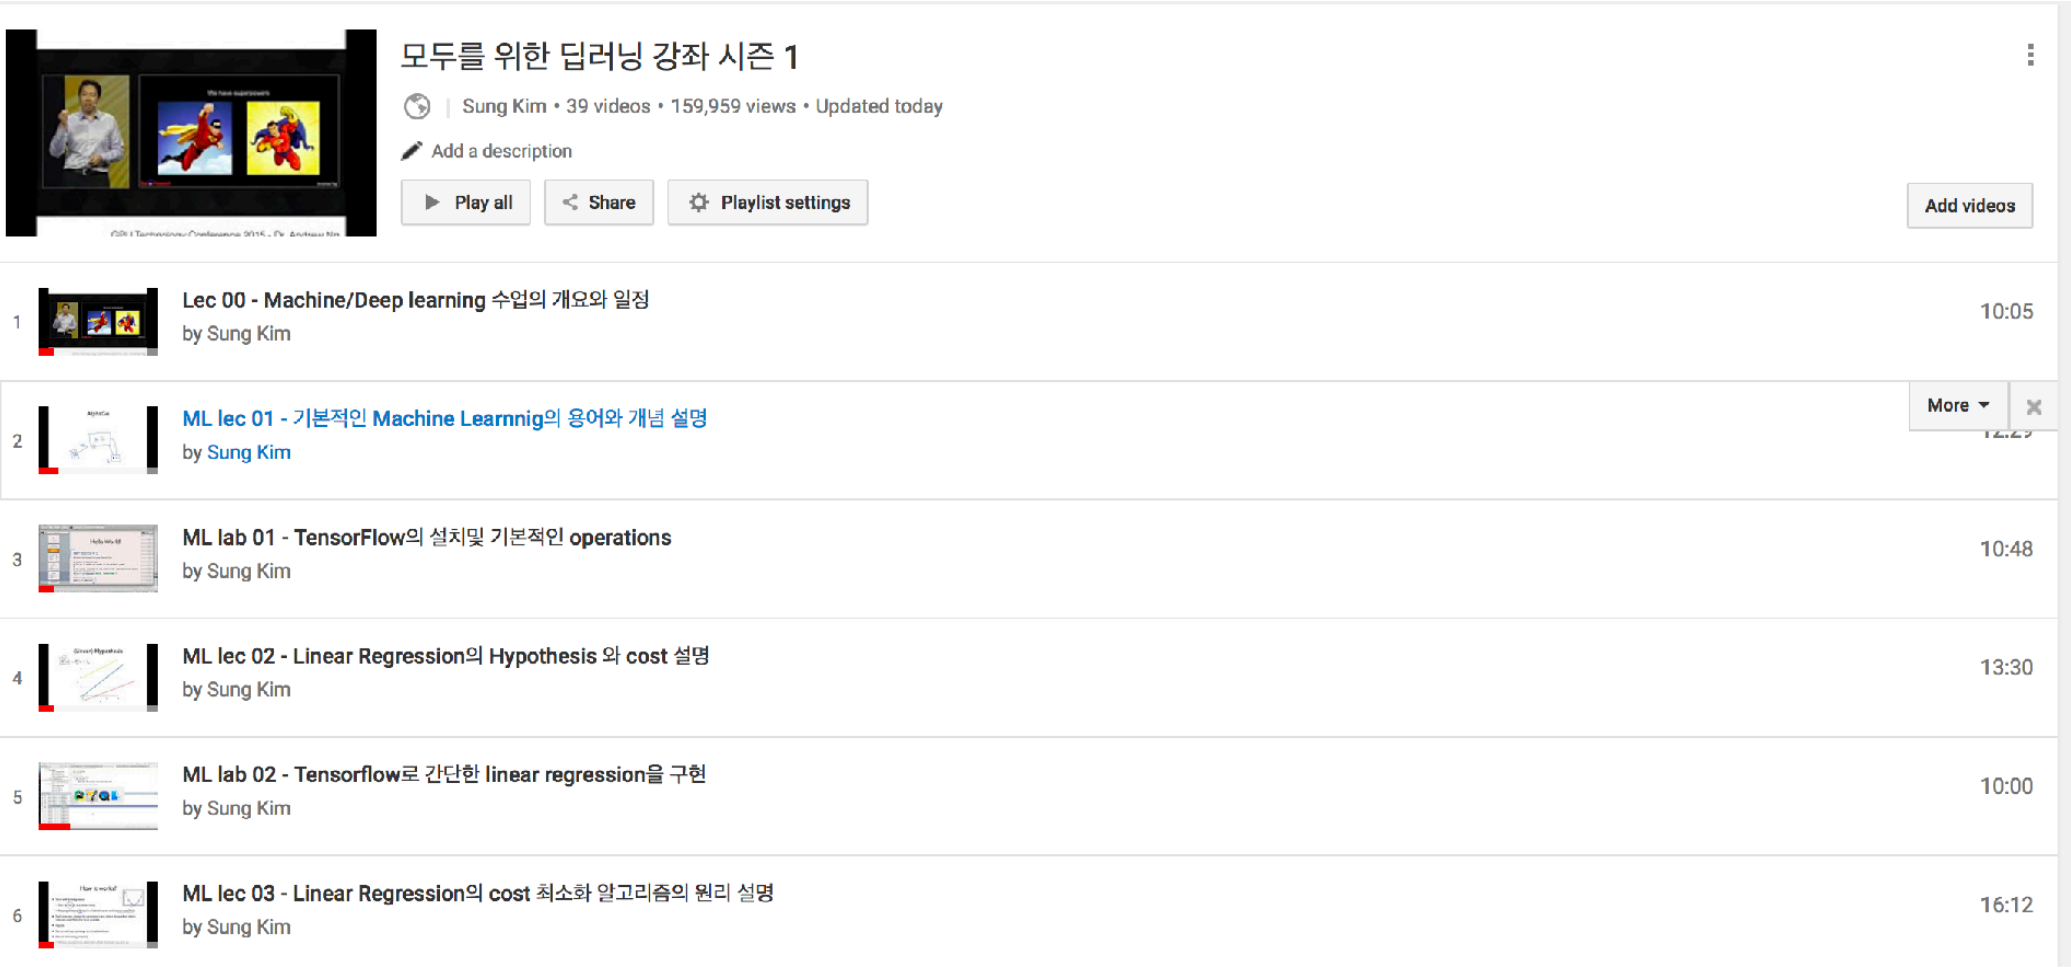

Next

Playing
OpenAI Gym games