PEDAGOGICAL TOOLS TO START REINFORCEMENT LEARNING.

Salma Samiei – David Rousseau

BOOKS

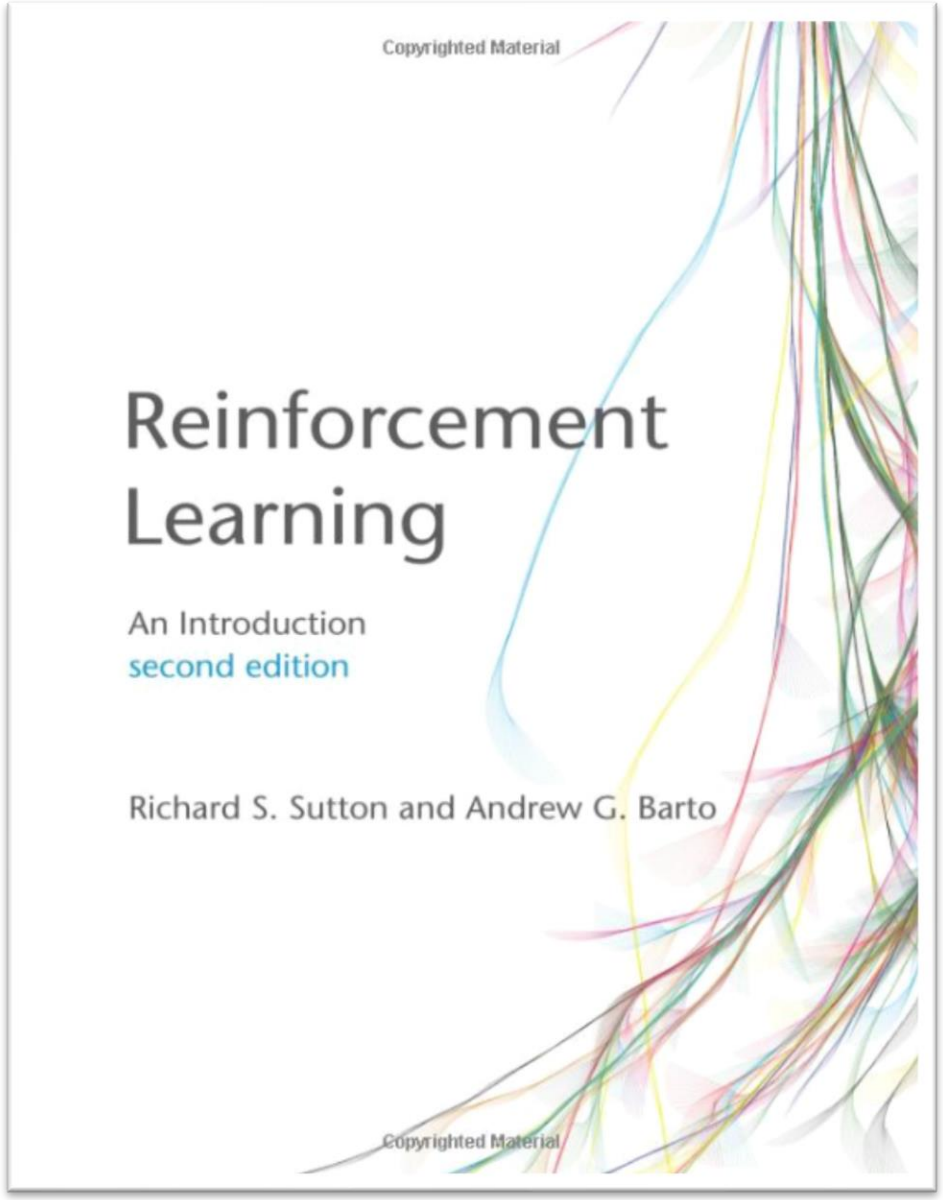

Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series)

By Richard S. Sutton , Andrew G. Barto, Francis Bach

Available free online
https://web.stanford.edu/class/psych209/Readings/SuttonBartoIPRLBook2ndEd.pdf

BOOKS

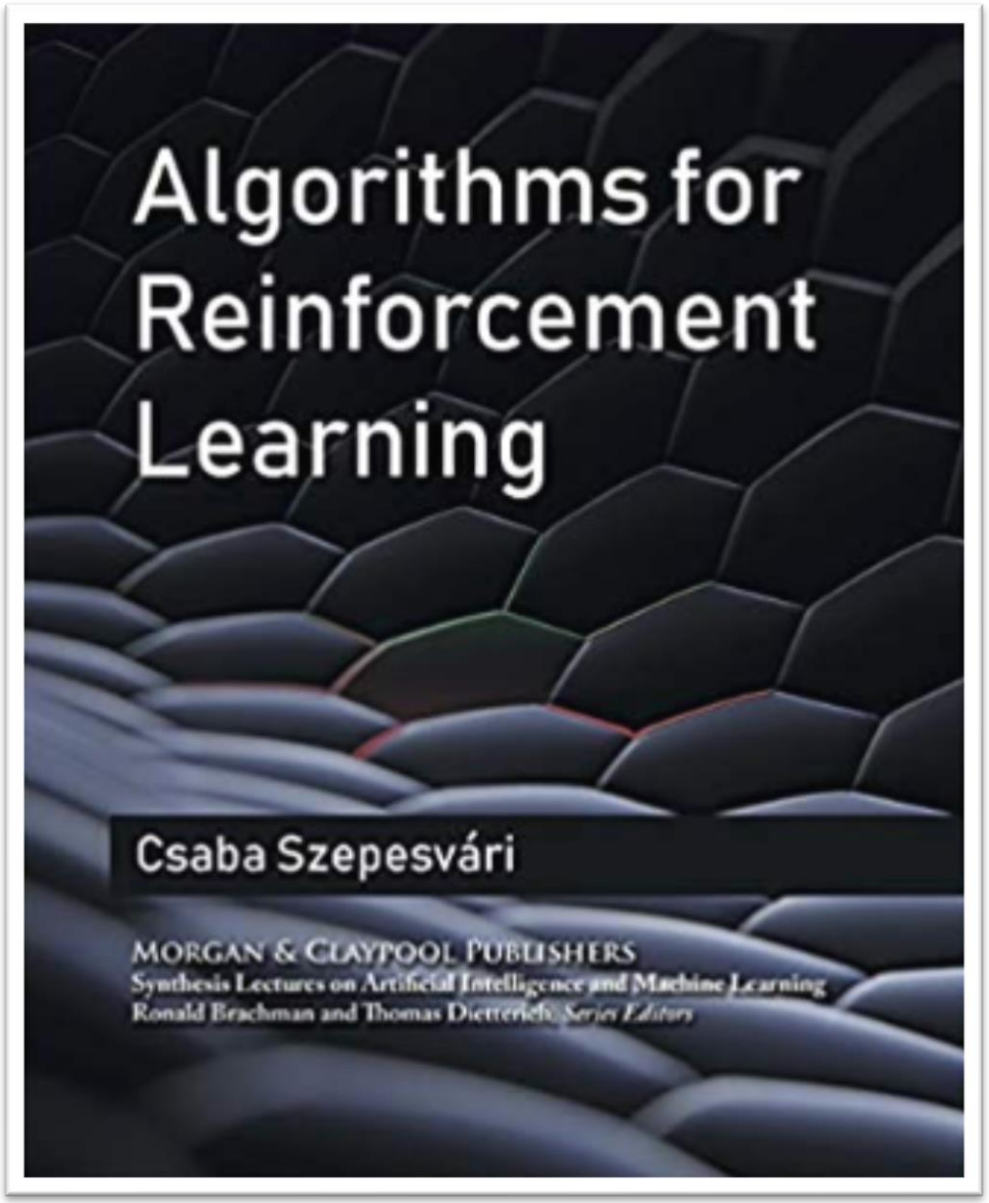

Algorithms for Reinforcement Learning (Synthesis Lectures on Artificial Intelligence and Machine Learning)

by Csaba Szepesvari.

Available free online
https://sites.ualberta.ca/~szepesva/papers/RLAlgsInMDPs.pdf

BOOKS

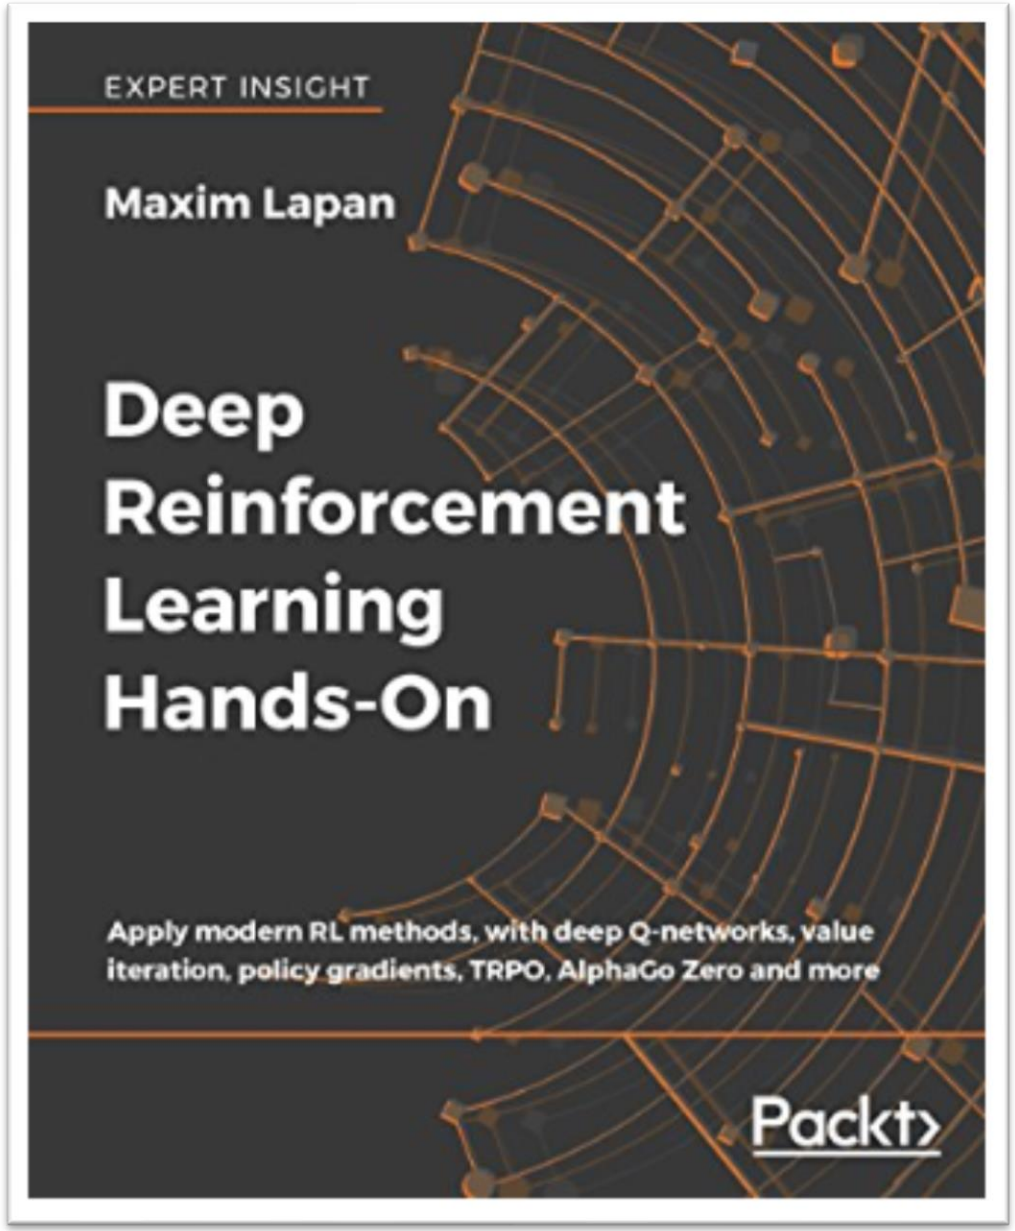

Deep Reinforcement Learning Hands-On: Apply modern RL methods, with deep Q-networks, value iteration, policy gradients, TRPO, AlphaGo Zero and more.

By Maxim Lapan.

MOOCS

UDACITY

FREE COURSE

Reinforcement Learning

by Georgia Tech

Offered at Georgia Tech as CS 8803

START FREE COURSE

COURSE COST	TIMELINE	SKILL LEVEL
Free	Approx. 4 Months	Advanced

https://www.udacity.com/course/reinforcement-learning--ud600

MOOCS

coursera

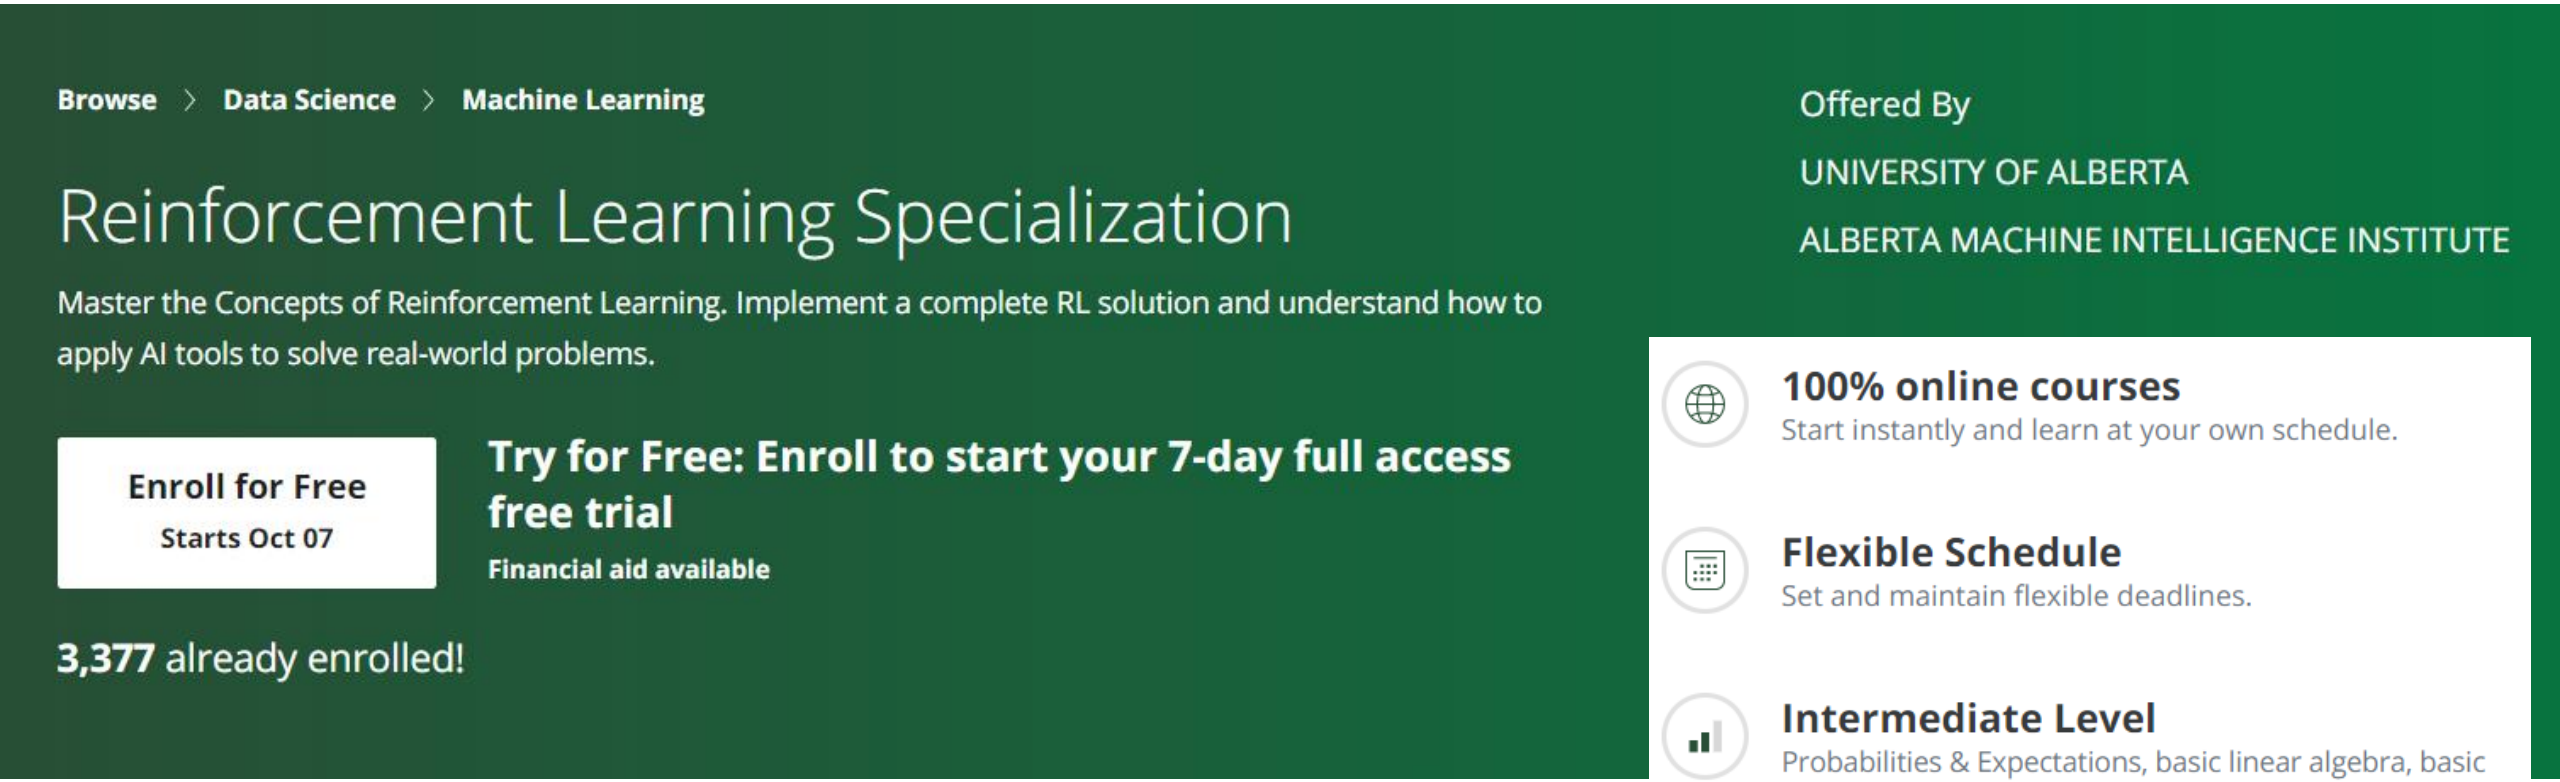

Course 1: Fundamentals of Reinforcement Learning.
Course 2: Sample-based Learning Methods.
Course 3: Prediction and Control with Function Approximation.
Course 4: A Complete Reinforcement Learning System (Capstone).

https://www.coursera.org/specializations/reinforcement-learning?action=enroll

MOOCS

coursera

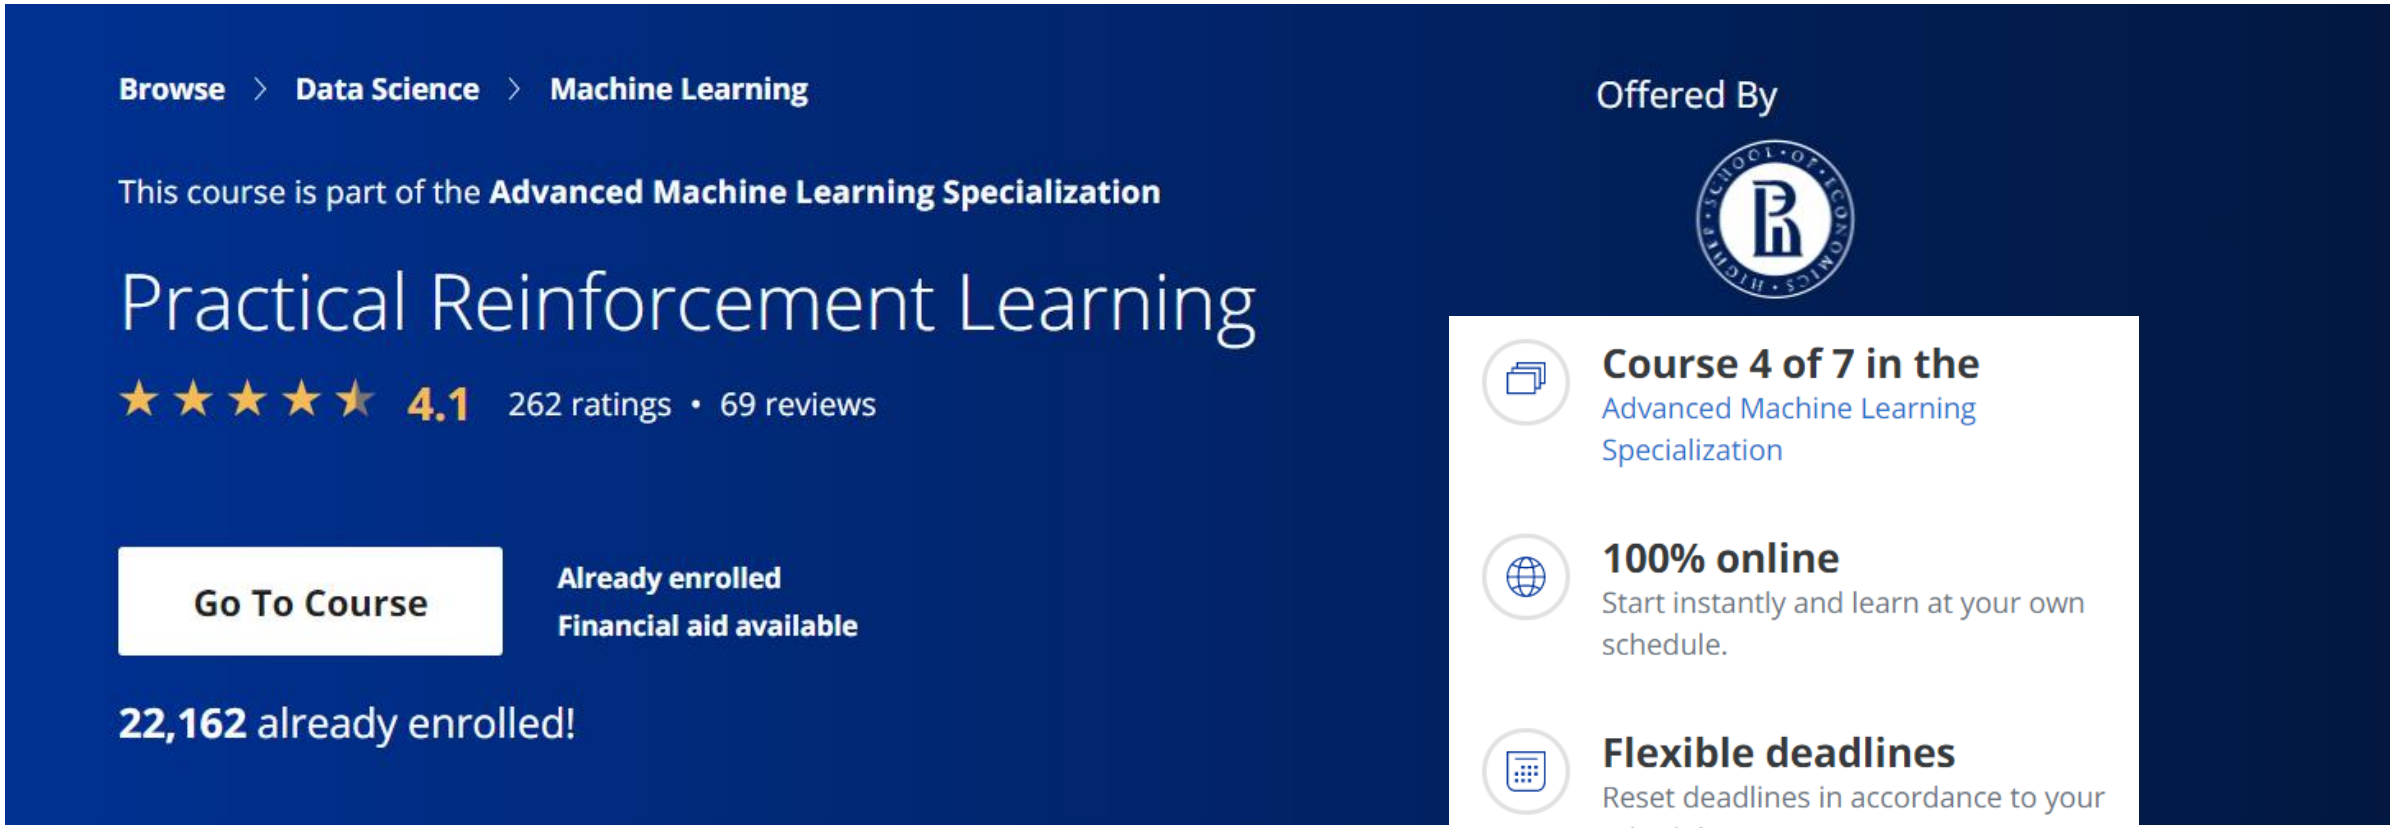

https://www.coursera.org/learn/practical-rl?specialization=aml

MOOCS

Reinforcement Learning Explained

Learn how to frame reinforcement learning problems, tackle classic examples, explore basic algorithms from dynamic programming, temporal difference learning, and progress towards larger state space using function approximation and DQN (Deep Q Network).

Microsoft

16,941 already enrolled!

Enroll
Starts Oct 7

☑ I would like to receive email from Microsoft and learn about other offerings related to Reinforcement Learning Explained.

Length:	6 Weeks
Effort:	4–8 hours per week
Price:	FREE Add a Verified Certificate for €90
Institution	Microsoft
Subject:	Computer Science
Level:	Advanced
Language:	English
Video Transcript:	English

https://www.edx.org/course/reinforcement-learning-explained-4

ONLINE LECTURES

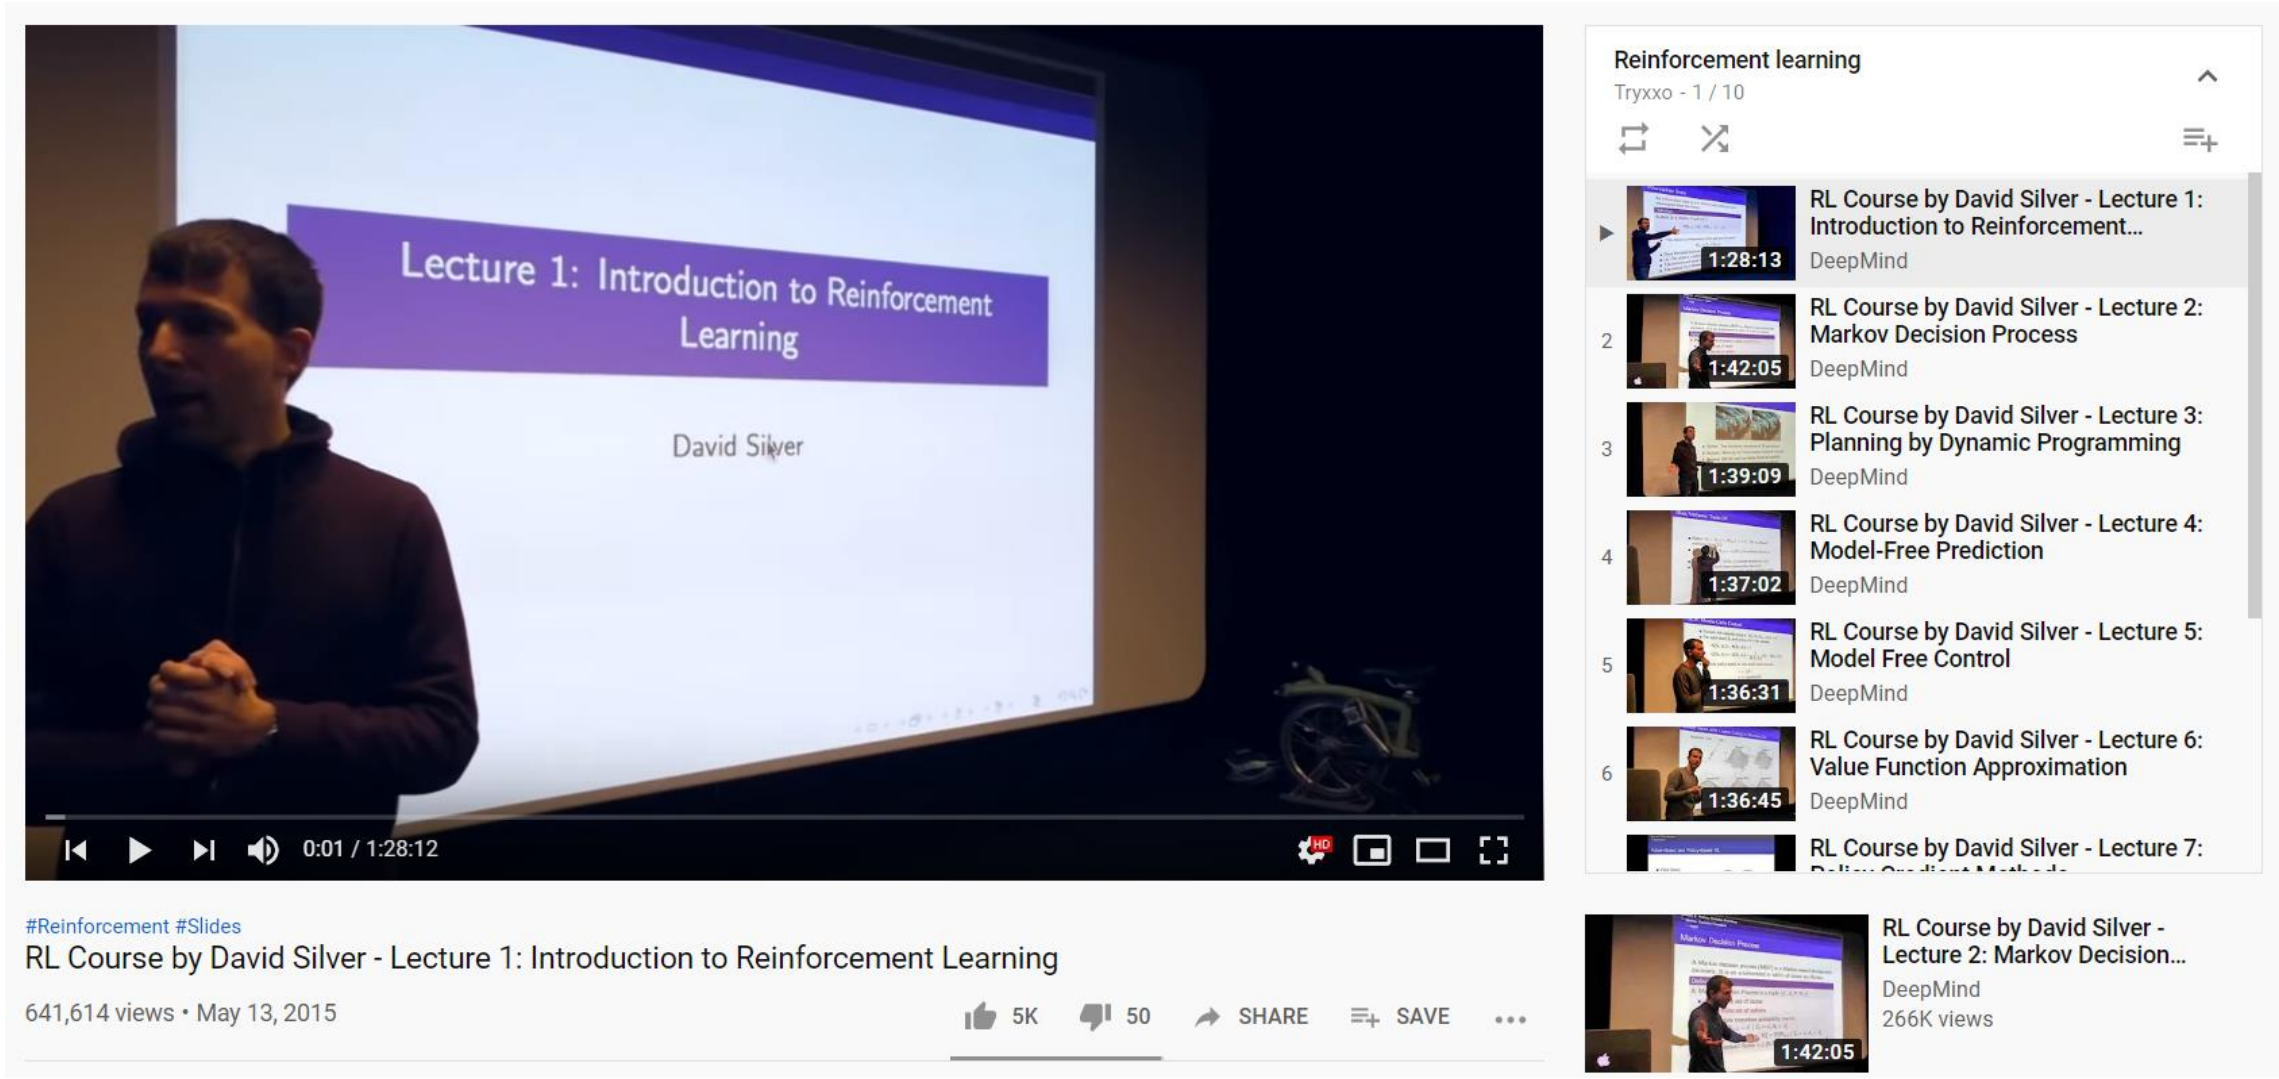

http://www0.cs.ucl.ac.uk/staff/d.silver/web/Teaching.html
https://www.youtube.com/watch?list=PL7-jPKtc4r78-wCZcQn5IqyuWhBZ8fOxT&v=2pWv7GOvuf0

ONLINE LECTURES

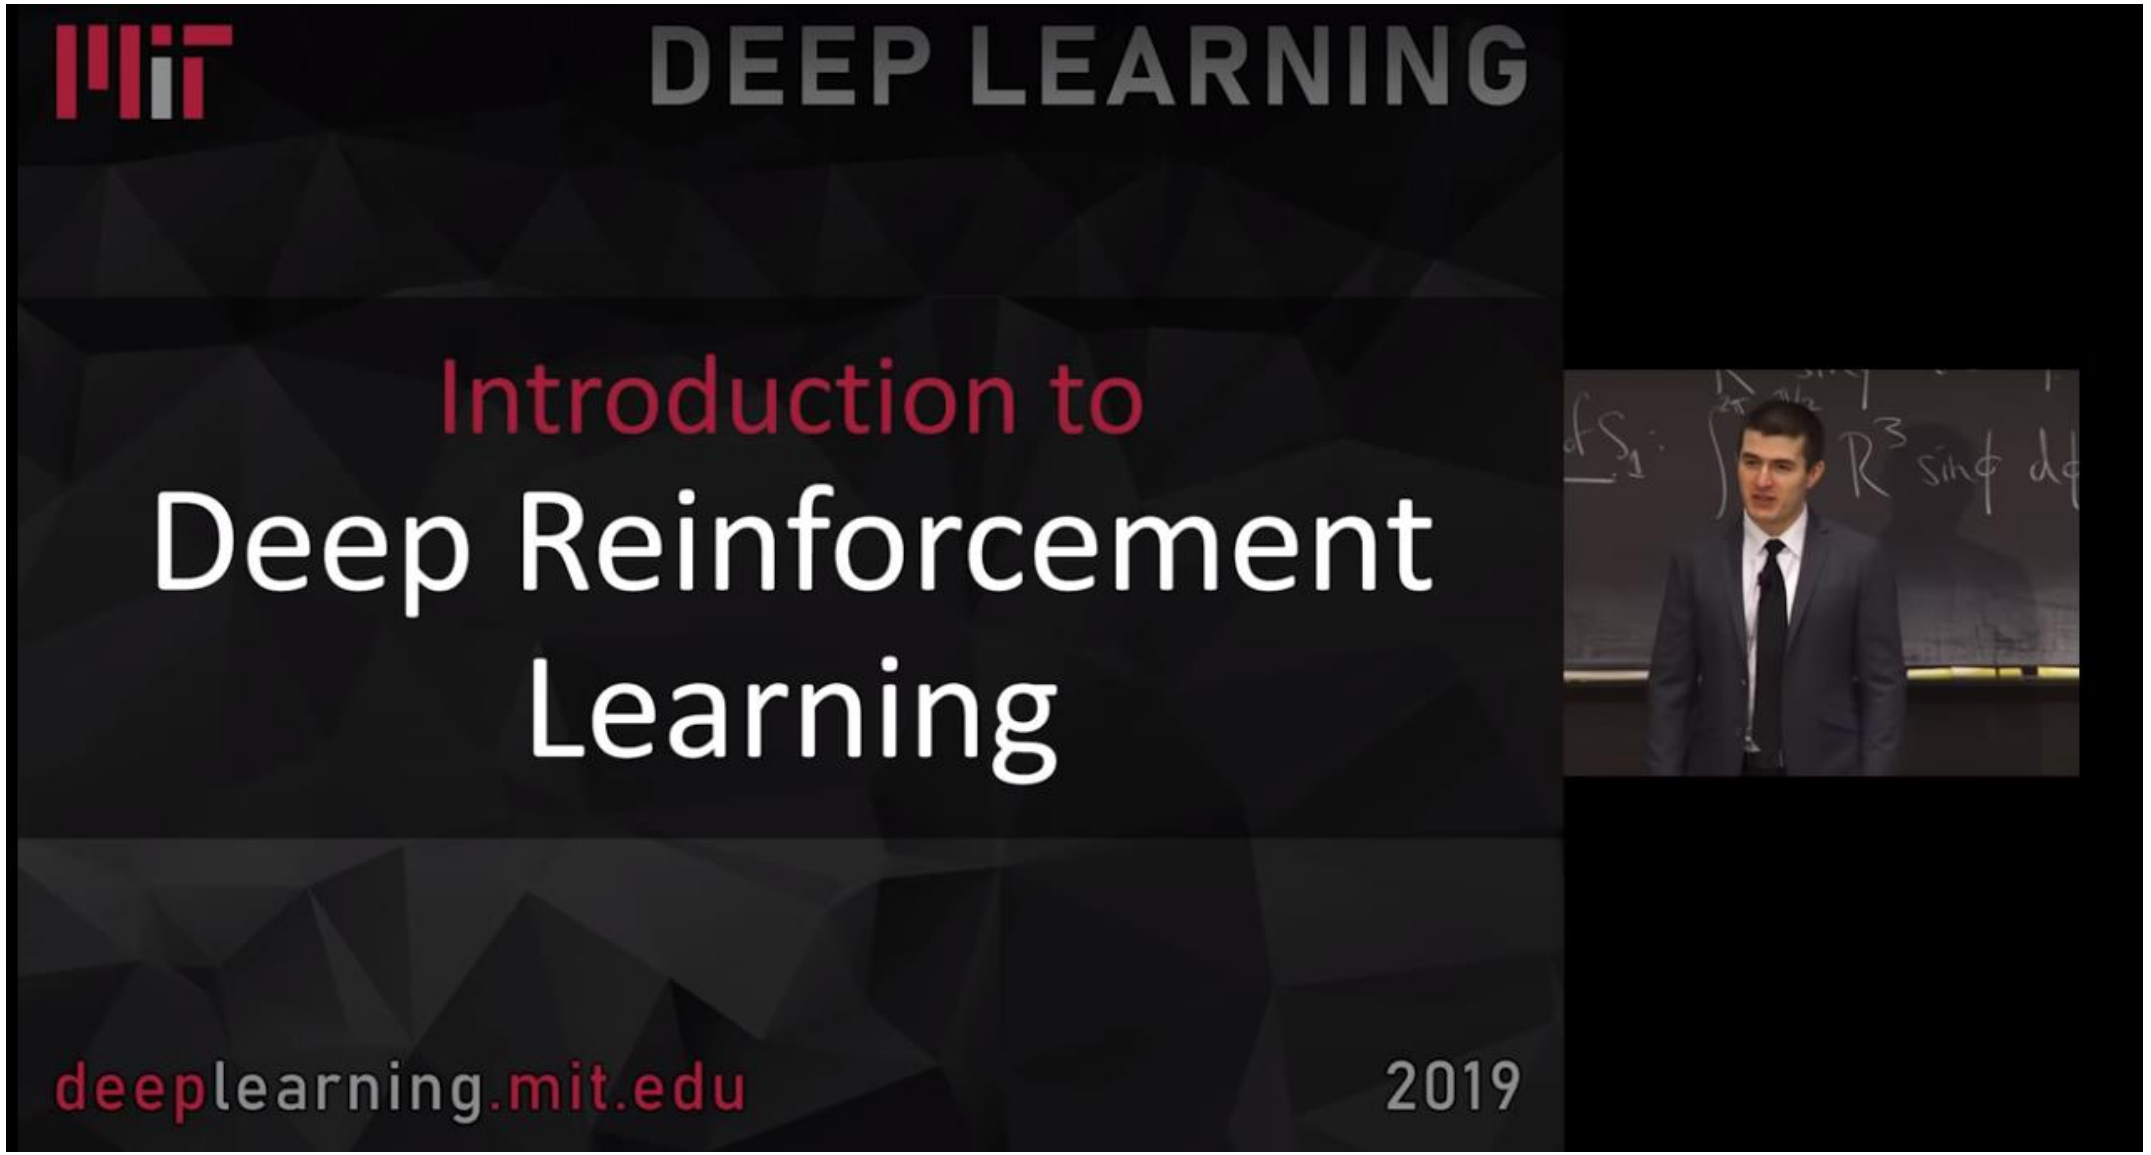

https://deeplearning.mit.edu/

ONLINE LECTURES

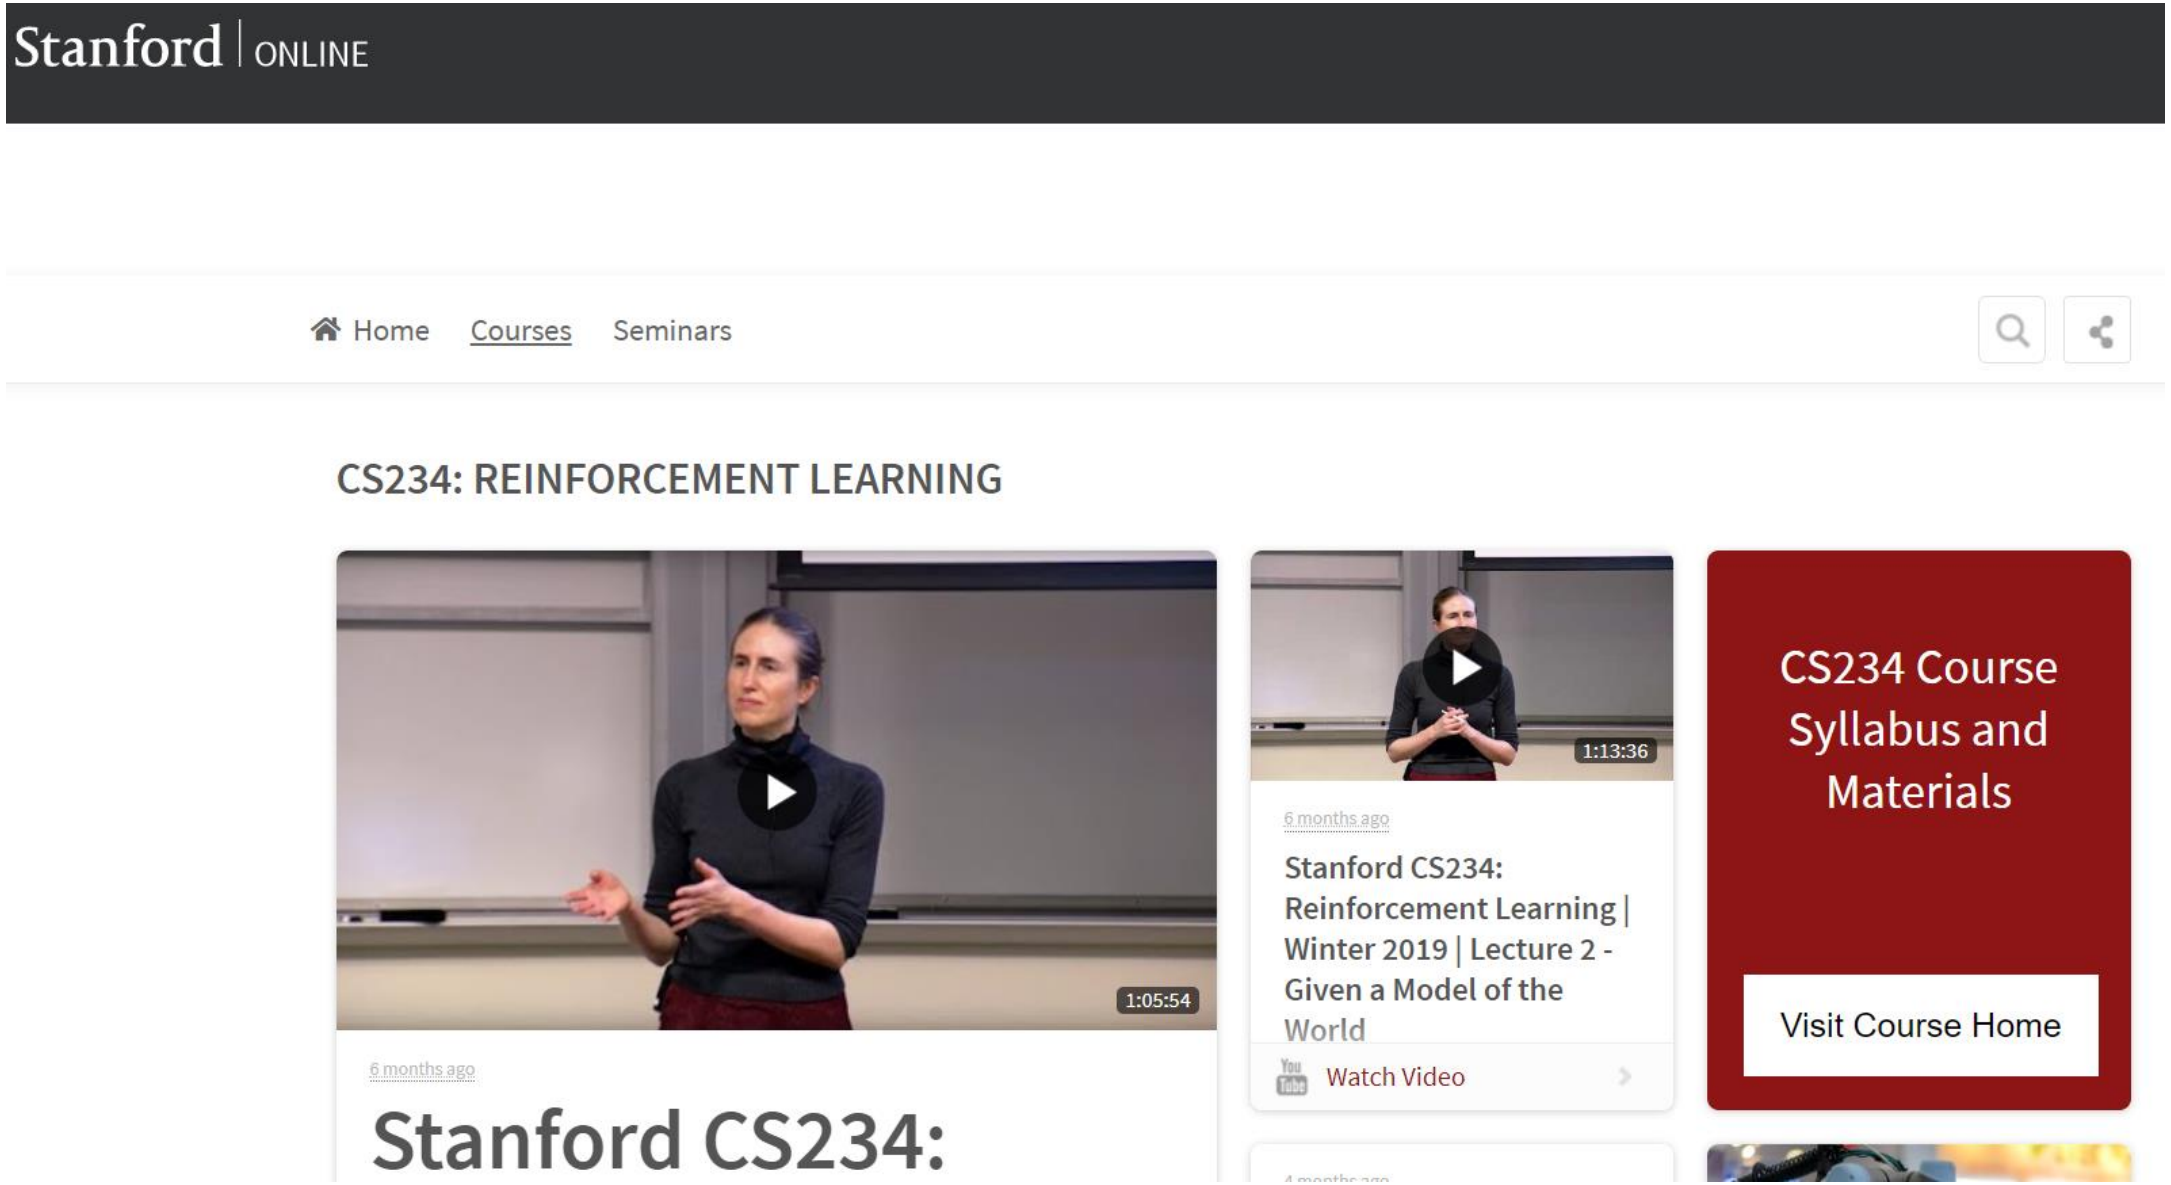

http://onlinehub.stanford.edu/cs234

ONLINE LECTURES

https://www.cs.ox.ac.uk/people/nando.defreitas/machinelearning/

ONLINE LECTURES

Deep RL Bootcamp

26-27 August 2017 | Berkeley CA

Lectures

- Core Lecture 1 Intro to MDPs and Exact Solution Methods -- Pieter Abbeel (video | slides)
- Core Lecture 2 Sample-based Approximations and Fitted Learning -- Rocky Duan (video | slides)
- Core Lecture 3 DQN + Variants -- Vlad Mnih (video | slides)
- Core Lecture 4a Policy Gradients and Actor Critic -- Pieter Abbeel (video | slides)
- Core Lecture 4b Pong from Pixels -- Andrej Karpathy (video | slides)
- Core Lecture 5 Natural Policy Gradients, TRPO, and PPO -- John Schulman (video | slides)
- Core Lecture 6 Nuts and Bolts of Deep RL Experimentation -- John Schulman (video | slides)
- Core Lecture 7 SVG, DDPG, and Stochastic Computation Graphs -- John Schulman (video | slides)
- Core Lecture 8 Derivative-free Methods -- Peter Chen (video | slides)
- Core Lecture 9 Model-based RL -- Chelsea Finn (video | slides)
- Core Lecture 10a Utilities -- Pieter Abbeel (video | slides)
- Core Lecture 10b Inverse RL -- Chelsea Finn (video | slides)
- Frontiers Lecture I: Recent Advances, Frontiers and Future of Deep RL -- Vlad Mnih (video | slides)
- Frontiers Lecture II: Recent Advances, Frontiers and Future of Deep RL -- Sergey Levine (video | slides)
- TAs Research Overviews (video | slides)

https://sites.google.com/view/deep-rl-bootcamp/lectures

ONLINE LECTURES

https://www.youtube.com/watch?v=OYhFoMySoVs

SCIENTIFIC ARTICLES

Semantic Scholar | All Fields | Search 177,623,775 papers from all fields of science

Reinforcement Learning: A Survey

Leslie Pack Kaelbling, Michael L. Littman, Andrew W. Moore • J. Artif. Intell. Res. • 1996

This paper **surveys** the field of **reinforcement learning** from a computer-science perspective. It is written to be accessible to researchers familiar with machine **learning**. Both the historical basis of... (More)

311 229 View PDF on ArXiv Cite Save

Introduction to Reinforcement Learning

Richard S. Sutton, Andrew G. Barto • 1998

From the Publisher: In **Reinforcement Learning**, Richard Sutton and Andrew Barto provide a clear and simple account of the key ideas and algorithms of **reinforcement learning**. Their discussion ranges... (More)

325 300 2 View Slides Cite Save

Reinforcement learning in robotics: A survey

Jens Kober, J. Andrew Bagnell, Jan Peters • I. J. Robotics Res. • 2012

Reinforcement learning offers to robotics a framework and set of tools for the design of sophisticated and hard-to-engineer behaviors. Conversely, the challenges of robotic problems provide both... (More)

48 184 View PDF View Slides Cite Save

Reinforcement Learning: An Introduction

Richard S. Sutton, Andrew G. Barto • IEEE Transactions on Neural Networks • 1988

Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to **learning** whereby an agent tries to maximize the total amount of reward it... (More)

3263 1,637 View on IEEE Cite Save

WEB ARTICLES

Applications of Reinforcement Learning in Real World

There is no reasoning, no process of inference or comparison; there is no thinking about things, no putting two and two together; there are no ideas — the animal does not think of the box or of the food or of the act he is to perform. — — Edward Thorndike(1874–1949), the psychologist who proposed Law of effect.

Gary Follow

Aug 2, 2018 · 13 min read

Levine, Sergey, et al. "End-to-end training of deep visuomotor policies." *The Journal of Machine Learning Research* 17.1 (2016): 1334-1373.

https://towardsdatascience.com/applications-of-reinforcement-learning-in-real-world-1a94955bcd12

WEB ARTICLES

A Free course in Deep Reinforcement Learning from beginner to expert.

Some of the agents you'll implement during this course:

This course is a series of articles and videos where you'll master the skills and architectures you need, to become a deep reinforcement learning expert.

https://simoninithomas.github.io/Deep_reinforcement_learning_Course/ **Code is available**

Reinforcement Learning - A Simple Python Example and A Step Closer to AI with Assisted Q-Learning

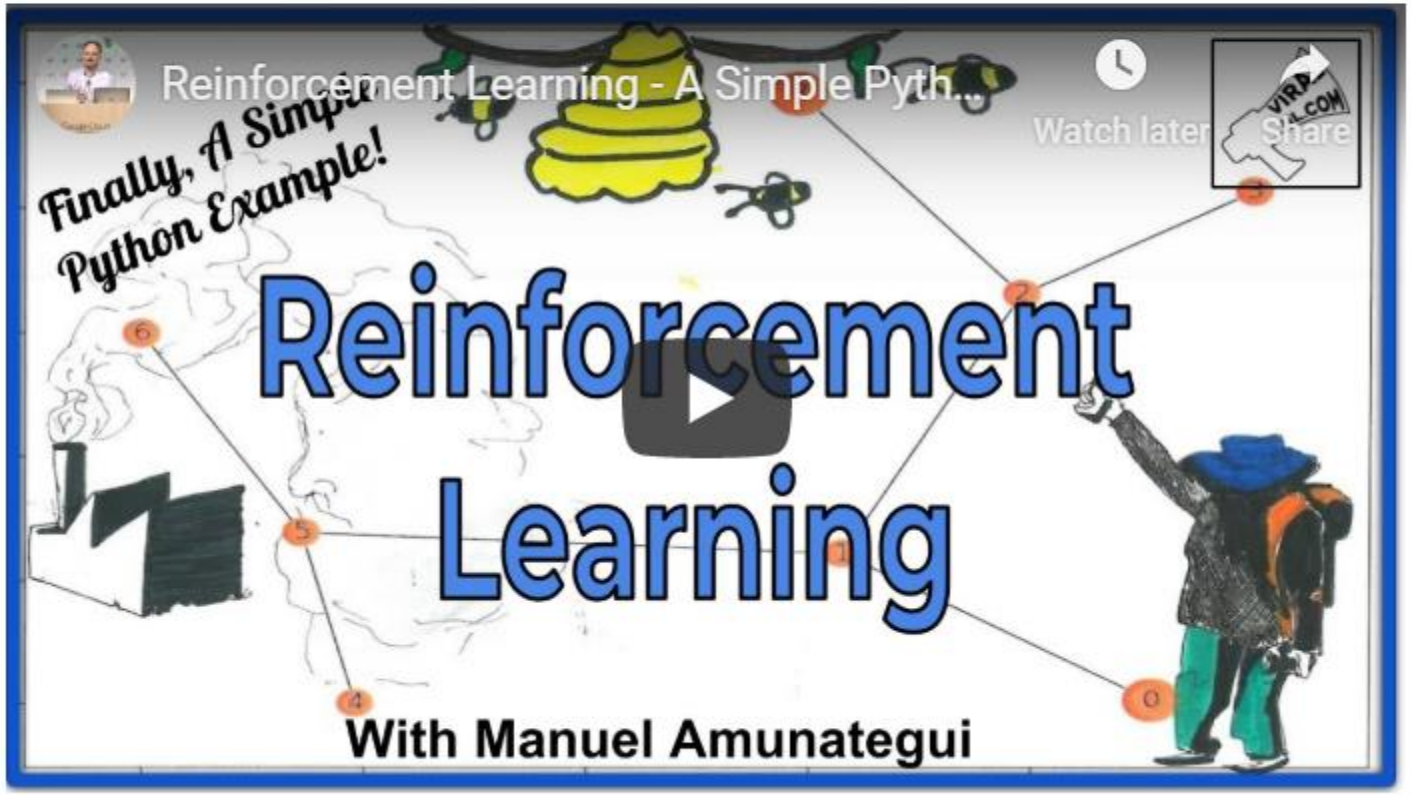

https://amunategui.github.io/reinforcement-learning/index.html

http://www.viralml.com/video-content.html?v=nSxaG_Kjw_w

Code is available

Q-Learning introduction and Q Table - Reinforcement Learning w/ Python Tutorial p.1

Code is available

https://pythonprogramming.net/q-learning-reinforcement-learning-python-tutorial/

WEB ARTICLES

Andrej Karpathy blog

About Hacker's guide to Neural Networks

Deep Reinforcement Learning: Pong from Pixels

May 31, 2016

This is a long overdue blog post on Reinforcement Learning (RL). RL is hot! You may have noticed that computers can now automatically learn to play ATARI games (from raw game pixels!), they are beating world champions at Go, simulated quadrupeds are learning to run and leap, and robots are learning how to perform complex manipulation tasks that defy explicit programming. It turns out that all of these advances fall under the umbrella of RL research. I also became interested in RL myself over the last ~year: I worked through Richard Sutton's book, read through David Silver's course, watched John Schulmann's lectures, wrote an RL library in Javascript, over the summer interned at DeepMind working in the DeepRL group, and most recently pitched in a little with the design/development of OpenAI Gym, a new RL benchmarking toolkit. So I've certainly been on this funwagon for at least a year but until now I haven't gotten around to writing up a short post on why RL is a big deal, what it's about, how it all developed and where it might be going.

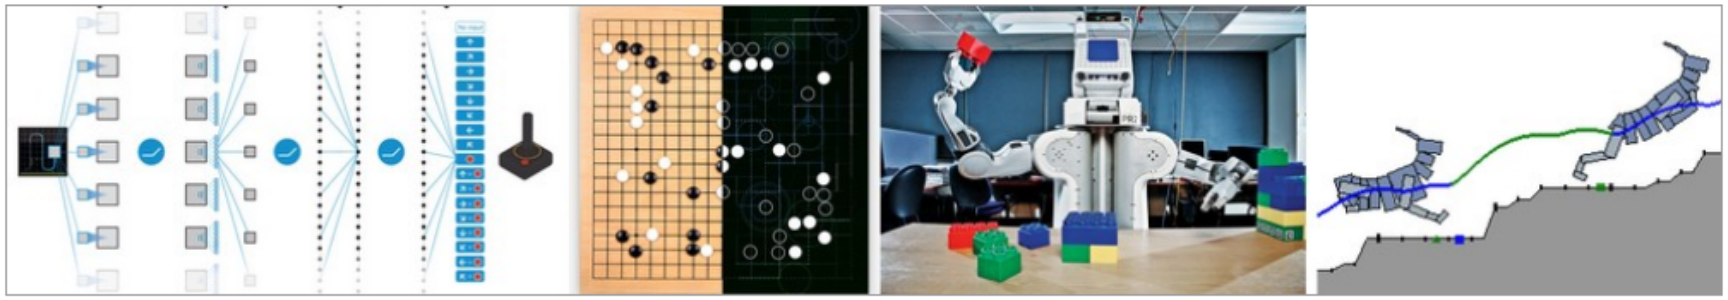

Examples of RL in the wild. **From left to right**: Deep Q Learning network playing ATARI, AlphaGo, Berkeley robot stacking Legos, physically-simulated quadruped leaping over terrain.

http://karpathy.github.io/2016/05/31/rl/ **Code is available**

Reinforcement Learning Tutorial with Open AI Gym

Abhinav Sagar Follow
Jul 12 · 4 min read

abhinavsagar / Reinforcement-Learning-Tutorial

Watch 1 | Star 3 | Fork 0

Code | Issues 0 | Pull requests 0 | Projects 0 | Security | Insights

Sample reinforcement learning tutorial notebooks

reinforcement-learning | reinforcement-learning-algorithms | reinforcement-learning-environments | reinforcement-learning-tutorials | machine-learning | pytorch

ddpg | dqn | dqn-pytorch | ddpg-pytorch | bipedalwalker | mountaincar | lunar-lander | pendulum

28 commits | 1 branch | 0 releases | 1 contributor | MIT

Branch: master | New pull request | Find File | Clone or download

abhinavsagar Update README.md — Latest commit 2776cf8 2 days ago

cross entropy	Add files via upload	3 months ago
ddpg walker	Add files via upload	3 months ago
ddqn space	Update README.md	2 months ago
.gitignore	Initial commit	3 months ago
LICENSE	Initial commit	3 months ago
README.md	Update README.md	2 days ago
ct1.gif	Add files via upload	2 months ago

https://towardsdatascience.com/reinforcement-learning-tutorial-with-open-ai-gym-9b11f4e3c204

https://github.com/abhinavsagar/ **Code is available**

WEB ARTICLES

Beat Atari with Deep Reinforcement Learning! (Part 0: Intro to RL)

Adrien Lucas Ecoffet Follow
Aug 22, 2017 · 8 min read

https://becominghuman.ai/lets-build-an-atari-ai-part-0-intro-to-rl-9b2c5336e0ec

http://www.arcadepunk

Beat Atari with Deep Reinforcement Learning! (Part 1: DQN)

Adrien Lucas Ecoffet Follow
Aug 22, 2017 · 13 min read

https://becominghuman.ai/lets-build-an-atari-ai-part-1-dqn-df57e8ff3b26

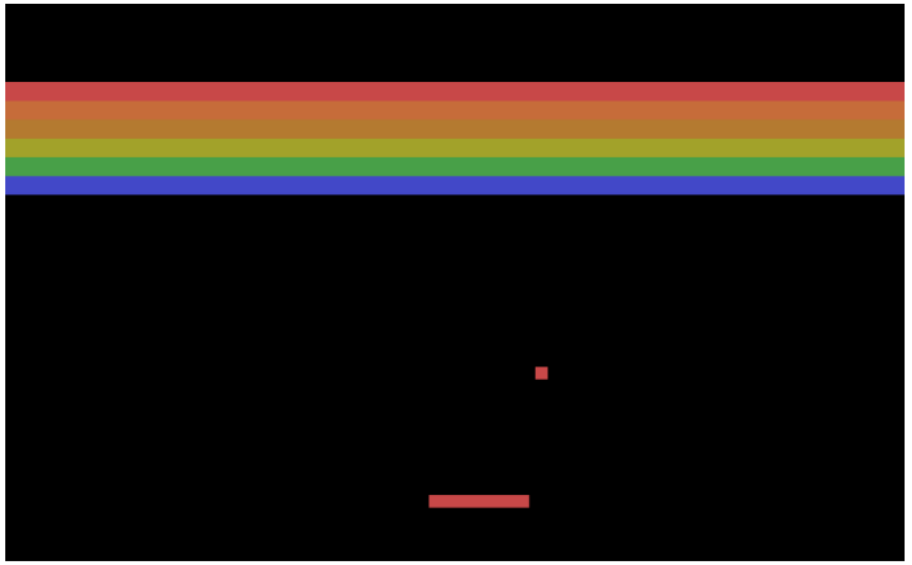

https://s-media-cache-ak0.pinimg.com/originals/4a/1c/2a/4a1c2a0dc943ab301a4b8b9ed1d27d00.jpg

Beat Atari with Deep Reinforcement Learning! (Part 2: DQN improvements)

Adrien Lucas Ecoffet Follow
Oct 31, 2017 · 8 min read

https://becominghuman.ai/beat-atari-with-deep-reinforcement-learning-part-2-dqn-improvements-d3563f665a2c

https://wallpapercave.com/wp/7PEFO6w.png

Codes are available

Spinning Up a Pong AI With Deep Reinforcement Learning

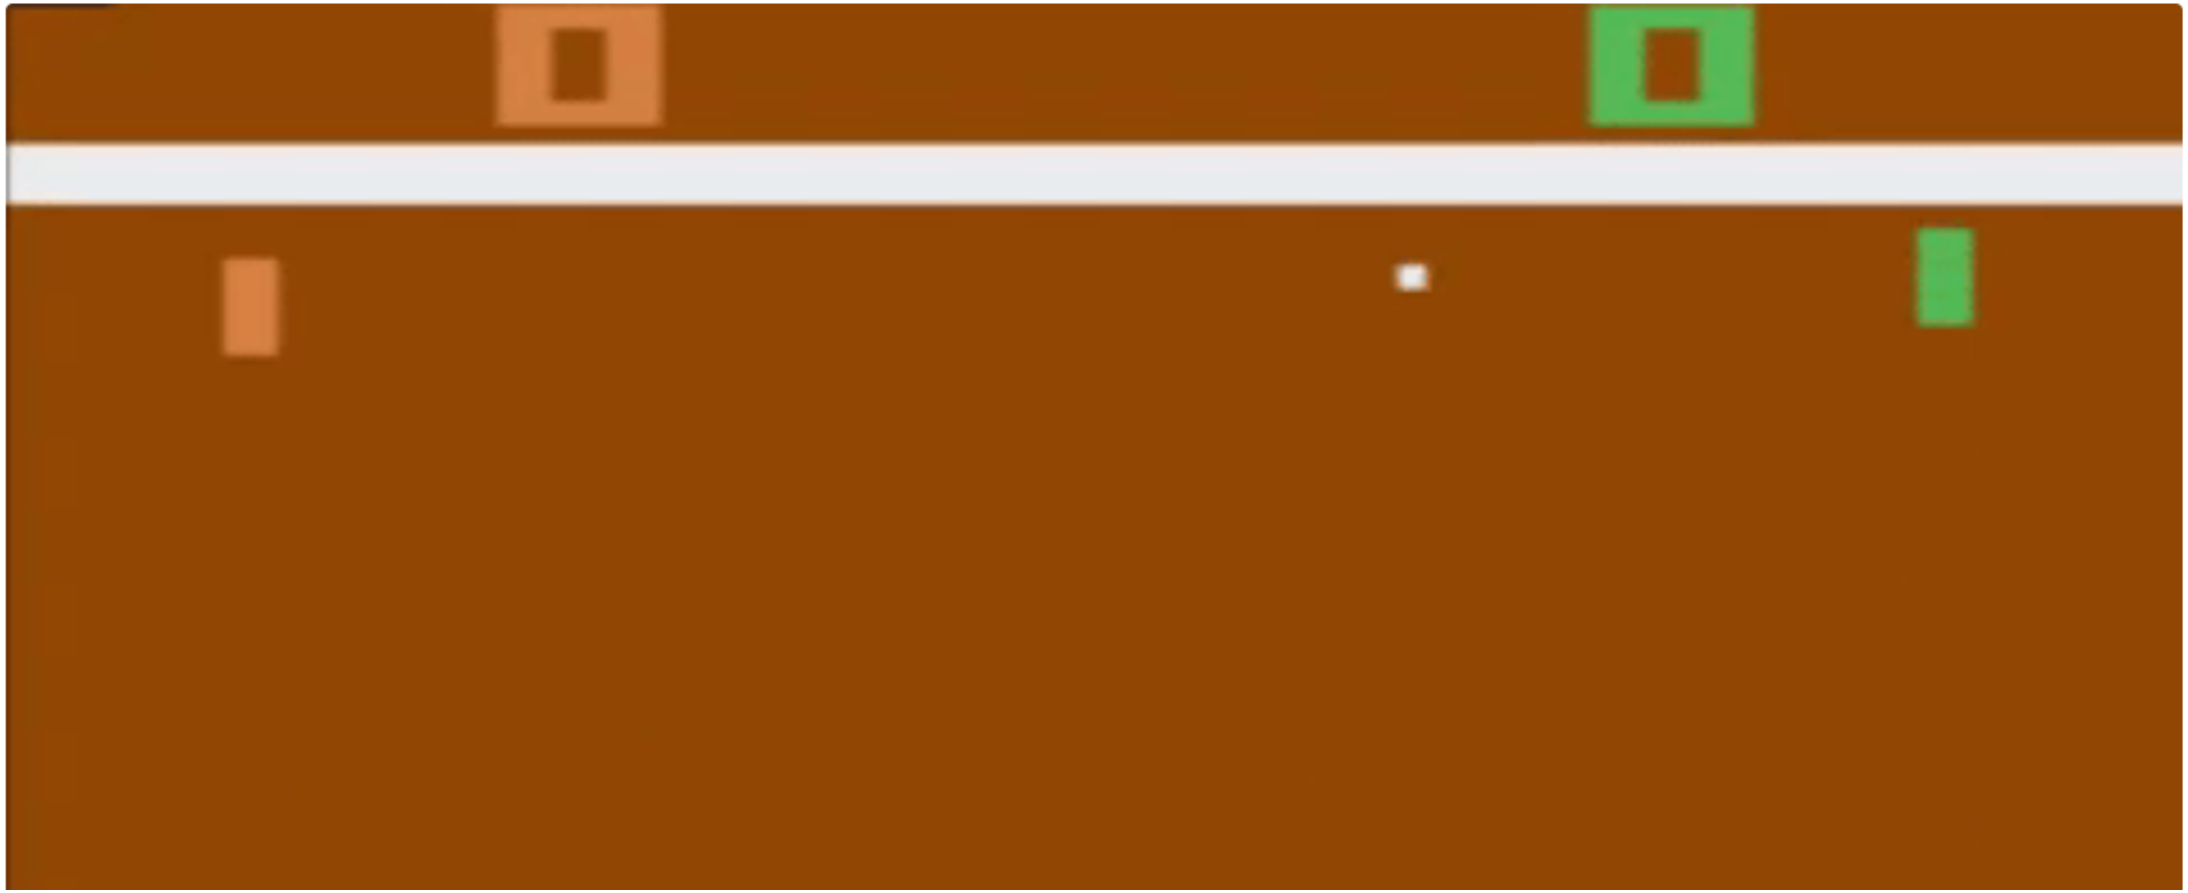

https://blog.floydhub.com/spinning-up-with-deep-reinforcement-learning/ **Code is available**

WEB ARTICLES

PyTorch

Get Started | Features | Ecosystem | Blog | Tutorials | Docs | Resources | Github

1.2.0

Search Tutorials

Getting Started

Deep Learning with PyTorch: A 60 Minute Blitz

Writing Custom Datasets, DataLoaders and Transforms

Visualizing Models, Data, and Training with TensorBoard

Image

TorchVision Object Detection Finetuning Tutorial

Transfer Learning for Computer Vision Tutorial

Spatial Transformer Networks Tutorial

Neural Transfer Using PyTorch

Adversarial Example Generation

DCGAN Tutorial

Tutorials > Reinforcement Learning (DQN) Tutorial

Shortcuts

Run in Google Colab | Download Notebook | View on GitHub

REINFORCEMENT LEARNING (DQN) TUTORIAL

Author: Adam Paszke

This tutorial shows how to use PyTorch to train a Deep Q Learning (DQN) agent on the CartPole-v0 task from the OpenAI Gym.

Task

The agent has to decide between two actions - moving the cart left or right - so that the pole attached to it stays upright. You can find an official leaderboard with various algorithms and visualizations at the Gym website.

Reinforcement Learning (DQN) Tutorial

Replay Memory

+ DQN algorithm

+ Training

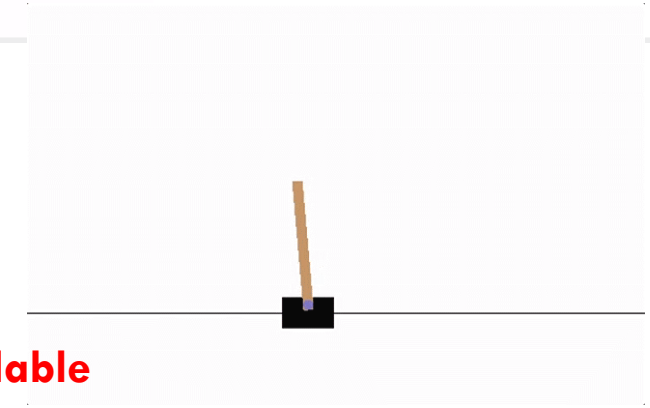

https://pytorch.org/tutorials/intermediate/reinforcement_q_learning.html

Code is available

WEB ARTICLES

Reinforcement Learning: Introduction to Monte Carlo Learning using the OpenAI Gym Toolkit

ANKIT CHOUDHARY, NOVEMBER 19, 2018 LOGIN TO BOOKMARK THIS ARTICLE

Monte Carlo Implementation in Python

Frozen Lake Environment

The agent controls the movement of a character in a grid world. Some tiles of the grid are walkable, and others lead to the agent falling into the water. Additionally, the movement direction of the agent is uncertain and only partially depends on the chosen direction. The agent is rewarded for finding a walkable path to a goal tile.

The surface is described using a grid like the following:

S	F	F	F
F	H	F	H
F	F	F	H
H	F	F	G

(S: starting point, safe), (F: frozen surface, safe), (H: hole, fall to your doom), (G: goal)

Monte Carlo Methods – An Example

Any method which solves a problem by generating suitable random numbers, and observing that fraction of numbers obeying some property or properties, can be classified as a Monte Carlo method.

Let's do a fun exercise where we will try to find out the value of pi using pen and paper. Let's draw a square of unit length and draw a quarter circle with unit length radius. Now, we have a helper bot C3PO with us. It is tasked with putting as many dots as possible on the square randomly 3,000 times, resulting in the following figure:

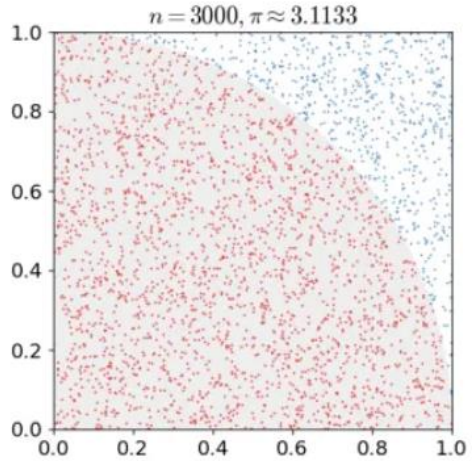

C3PO needs to count each time it puts a dot inside a circle. So, the value of pi will be given by:

$$pi = 4 * \frac{N}{3000}$$

Code is available

https://www.analyticsvidhya.com/blog/2018/11/reinforcement-learning-introduction-monte-carlo-learning-openai-gym/

Q-LEARNING ALGORITHM

Q-learning is a model-free reinforcement learning algorithm. The goal of Q-learning is to learn a policy, which tells an agent what action to take under what circumstances.

```
define epsilon, alpha, gamma
initiate Q-table with zeros
observe initial state s
repeat:
        select an action a
            with probability ε select random action a
            else select action a = argmax(Q(s,a'))
        carry out action a
        observe next state s' and reward r
        calculate Q_target = r + gamma*max[Q(s',A)]
        Calculate Q_delta = Q_target - Q(s,a)
        add alpha*Q_delta to Q(s,a)
        s = s'
until terminated
```

GRID WORLD EXAMPLE

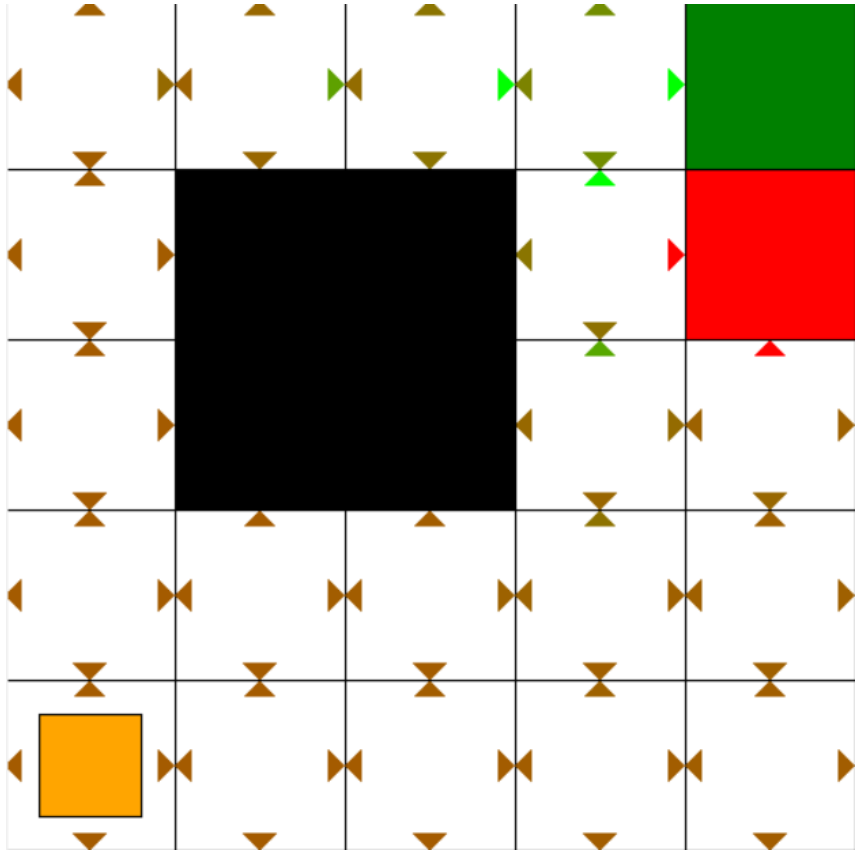

Environment 5*5 grid - Each cell represent the state the agent can be in.

Agent

4 actions: Moving step will give us a reward of -0.04.

Trap: Red square will give us a reward of -1.

Target: Green square will give us a reward of +1.

LET SEE THE CODE

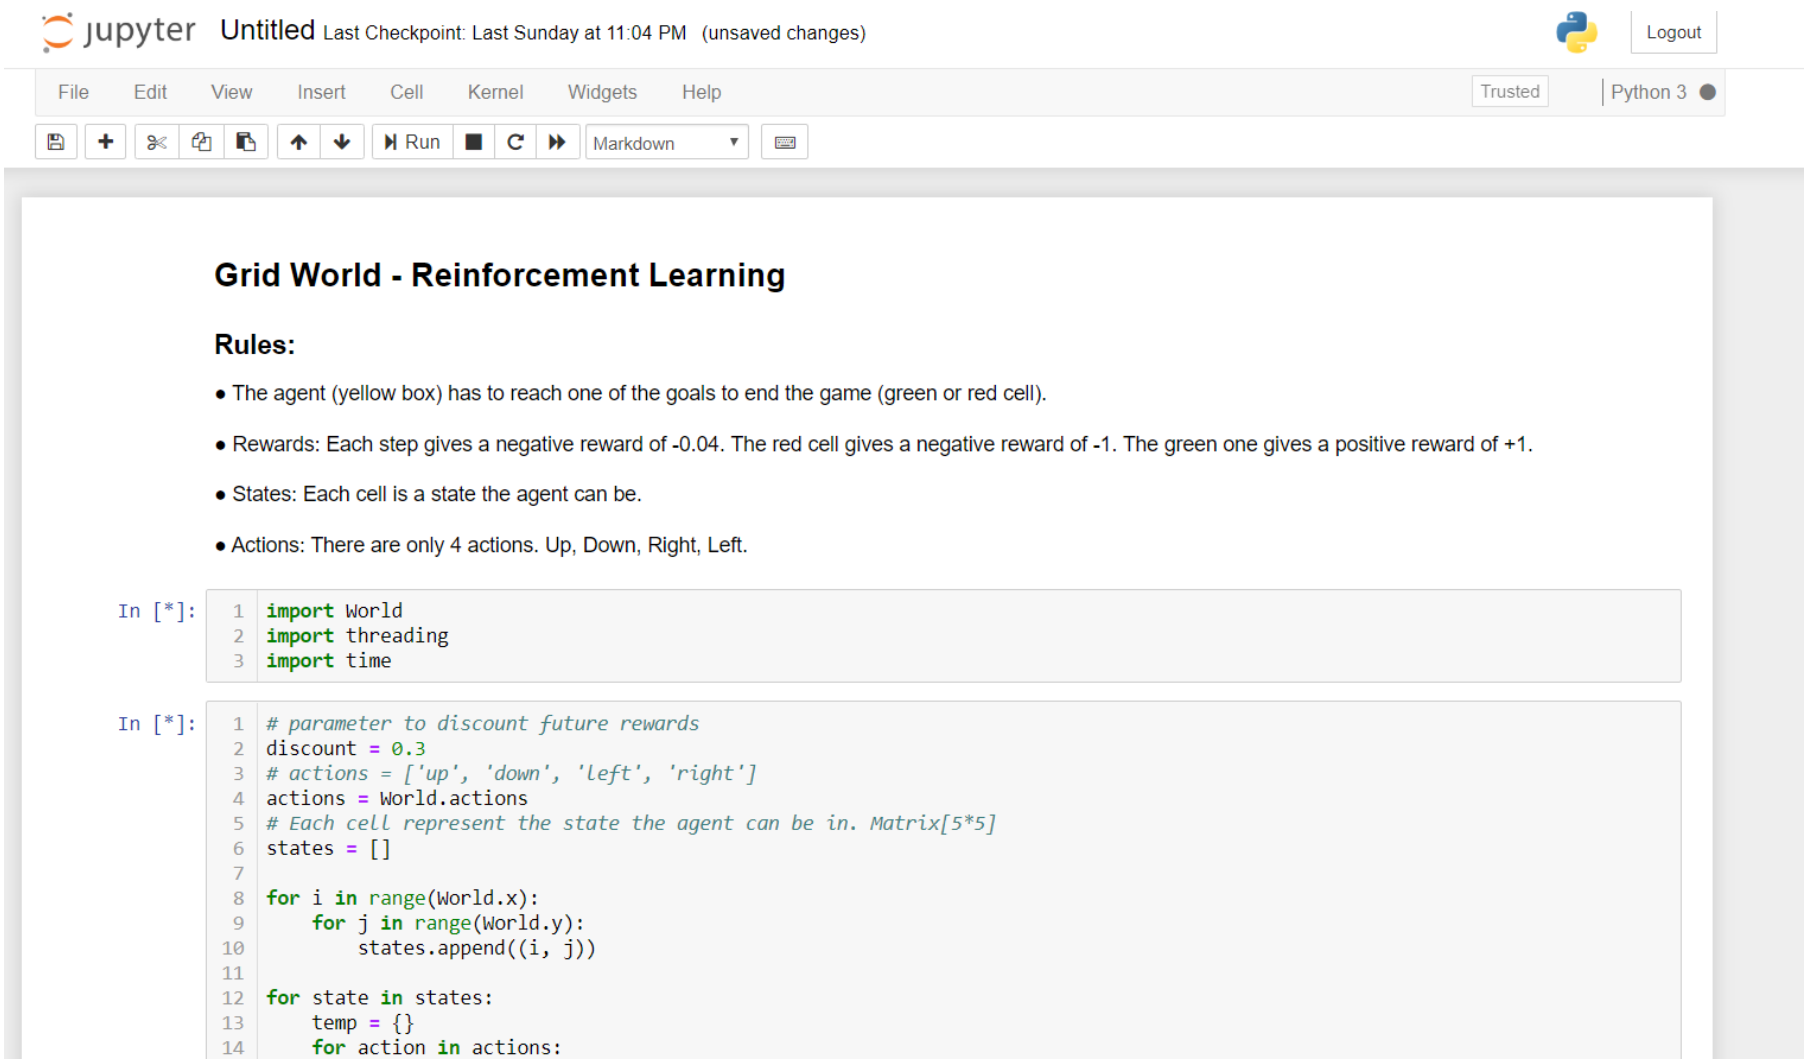

THANK YOU!