Reinforcement Learning Explained

Modality: On Demand

Duration: 42 Hours

About this course:

Reinforcement Learning (RL) is a machine learning region, where a specialist learns by cooperating with its condition to accomplish an objective.

With this course, we will be acquainted with the universe of reinforcement learning. You will figure out how to outline problems of reinforcement learning and begin handling great models like learning to navigate in a grid-world, news recommendation, and balancing a cart-pole.

You will investigate the essential calculations from dynamic programming, multi-equipped bandits, TD (temporal difference) learning, and progress towards bigger state-space utilizing the approximation of function, specifically utilizing profound learning. Also, you will find out calculations that emphasis on looking through the best strategy with the methods of policy gradient and actor-critic. In this way, you will get acquainted with Project Malmo, a stage for Artificial Intelligence experimentation and research-based over the Minecraft game.

Course Objective:

- Markov Decision Process
- Reinforcement Learning Problem
- Dynamic Programming
- Approximate Solution Methods
- Policy Gradient and Actor-Critic
- Bandits
- Temporal Difference Learning
- RL that Works

Audience:

Programmers

Data Analyst

Prerequisite:

No prerequisite required for this course

Course Outline:

Introduction to Reinforcement Learning

@Monto

- Comparisons
- Elements of RL
- Lab
- Knowledge Checks

Bandits

- Bandits Framework
- Regret Minimization
- Bridge to Reinforcement Learning
- Lab
- Knowledge Checks

The Reinforcement Learning Problem

- Agent and Environment Interface
- Markov Decision Process
- Lab
- Knowledge Checks

Dynamic Programming

- Basics of Dynamic Programming
- DP Observations
- Lab
- Knowledge Checks

Temporal Difference Learning

- Policy Evaluation
- Policy Optimization
- Lab
- Knowledge Checks

Function Approximation

- Why Use Function Approximation
- Linear Function Approximation
- Lab
- RL with Deep Neural Networks
- Lab
- Extensions to Deep Q-Learning
- Knowledge Checks

Policy Gradient and Actor Critic

• Introduction to Policy Optimization

@Morro

- Likelihood Ratio Methods
- Lab
- Variance Reduction
- Lab
- Actor Critic
- Lab
- Knowledge Checks?