
Course Code: W7105GW

Course Title: Deep Learning and Reinforcement Learning

Description:

This course introduces you to two of the most sought-after disciplines in Machine Learning: Deep Learning and Reinforcement Learning. Deep Learning is a subset of Machine Learning that has applications in both Supervised and Unsupervised Learning, and is frequently used to power most of the AI applications that we use on a daily basis. First you will learn about the theory behind Neural Networks, which are the basis of Deep Learning, as well as several modern architectures of Deep Learning. Once you have developed a few Deep Learning models, the course will focus on Reinforcement Learning, a type of Machine Learning that has caught up more attention recently. Although currently Reinforcement Learning has only a few practical applications, it is a promising area of research in AI that might become relevant in the near future. After this course, if you have followed the courses of the IBM Specialization in order, you will have considerable practice and a solid understanding in the main types of Machine Learning which are: Supervised Learning, Unsupervised Learning, Deep Learning, and Reinforcement Learning.

Objectives:

By the end of this course you should be able to:- Explain the kinds of problems suitable for Unsupervised Learning approaches.

- Explain the curse of dimensionality, and how it makes clustering difficult with many features.
- Describe and use common clustering and dimensionality-reduction algorithms.
- Try clustering points where appropriate, compare the performance of per-cluster models.
- Understand metrics relevant for characterizing clusters

Prerequisites:

To make the most out of this course, you should have familiarity with programming on a Python development environment, as well as fundamental understanding of Data Cleaning, Exploratory Data Analysis, Unsupervised Learning, Supervised Learning, Calculus, Linear Algebra, Probability, and Statistics.

Duration:

14.4 Hrs

Topics:

1. Introduction to Neural Networks
2. Neural Network Optimizers and Keras
3. Convolutional Neural Networks
4. Recurrent Neural Networks and Long-Short Term Memory Networks

5. Deep Learning with Autoencoders

6. Deep Learning Applications and Reinforcement Learning

Audience:

This course targets aspiring data scientists interested in acquiring hands-on experience with Deep Learning and Reinforcement Learning.