

Common Machine Learning Models

Modality: On Demand

Duration: 25 Hours

About this Course:

This course explores the core concepts of machine learning, the models available, and how to train them. We will take a deeper look at what it means to train a machine learning model, as well as the data and methods required to do so. We will also provide an overview of the most common models you're likely to encounter and take a practical approach to understand when and how to use them to solve business problems.

In the second half of this course, you will be guided through a series of case studies that will show you how to apply the concepts covered in this course to real-life examples.

Course Objectives:

- Understand the key concepts and models related to machine learning.
- Learn how to use training data sets with machine learning models.
- Learn how to choose the best machine learning model to suit your requirements.
- Understand how machine learning concepts can be applied to real-world scenarios in property prices, health, animal classification, and marketing activities.

Audience:

This course is intended for anyone who is:

- Interested in understanding machine learning models on a deeper level.
- Looking to enrich their understanding of machine learning and how to use it to solve complex problems.
- Looking to build a foundation for continued learning in the machine learning space and data science in general.

Prerequisites:

To get the most out of this course, you should have a general understanding of data concepts as well as some familiarity with cloud providers and their managed services, especially Amazon or Google. Some experience in data or development is preferable but not essential.

Course Outline:

This Course Includes:

- Course Introduction

- Lecture 1: Regression Trees and Random Forests
- Lecture 2: Support Vector Machines
- Lecture 3: Ensemble Methods and Hyper-Parameter Search
- Lecture 4: Unsupervised Learning: Clustering