

Description

In this course, learn how to solve a real-world use case with machine learning and produce actionable results using Amazon SageMaker. This course teaches you how to use Amazon SageMaker to cover the different stages of the typical data science process, from analyzing and visualizing a data set, to preparing the data and feature engineering, down to the practical aspects of model building, training, tuning and deployment.

Intended Audience

This course is intended for:

- Data science practitioners
- Machine learning practitioners
- Developers and engineers
- Systems architects

Course Objectives

In this course, you will learn how to:

- Apply Amazon SageMaker to a specific use case and dataset
- Practice all the steps of the typical data science process
- Visualize and understand the dataset
- Explore how the attributes of the dataset relate to each other
- Prepare the dataset for training
- Use built-in algorithms
- Train models with Amazon SageMaker using built-in algorithms
- Explore results and performance of the model, and demonstrate how it can be tuned and executed outside of SageMaker
- Run predictions on a batch of data with Amazon SageMaker
- Deploy a model to an endpoint in Amazon SageMaker for real-time predictions
- Learn how to configure an endpoint for serving predictions at scale
- Understand Hyperparameter Optimization (HPO) with Amazon SageMaker to find optimal model parameters
- Understand how to perform A/B model testing using Amazon SageMaker
- Perform the domain-specific cost of errors analysis to further tune the model threshold in order to maximize model utility expressed in financial terms

Prerequisites

We recommend that attendees of this course have the following prerequisites:

- Experience with Python programming language
- Familiarity with NumPy and Pandas Python libraries is a plus
- Familiarity with fundamental machine learning algorithms
- Familiarity with productionizing machine learning models

Delivery Method

This course is delivered through [a mix of]:

- Hands-on labs

Hands-On Activity

This course allows you to test new skills and apply knowledge to your working environment through a variety of practical exercises.

Duration

1 day

Course Outline

- Analyzing and visualizing a dataset
- Preparing the data and feature engineering
- Model building, training, tuning and deployment