

The Machine Learning Pipeline on AWS (AWS-ML-PL)

Course Summary

Description

Learn how to use the machine learning (ML) pipeline with Amazon SageMaker with hands-on exercises and four days of instruction. You will learn how to frame your business problems as ML problems and use Amazon SageMaker to train, evaluate, tune, and deploy ML models. Hands-on learning is a key component of this course, so you'll choose a project to work on, and then apply the knowledge and skills you learn to your chosen project in each phase of the pipeline. You'll have a choice of projects: fraud detection, recommendation engines, or flight delays.

Objectives

At the end of this course, students will be able to:

- Select and justify the appropriate ML approach for a given business problem
- Use the ML pipeline to solve a specific business problem
- Train, evaluate, deploy, and tune an ML model in Amazon SageMaker
- Describe some of the best practices for designing scalable, cost-optimized, and secure ML pipelines in AWS
- Apply machine learning to a real-life business problem after the course is complete

Topics

- Introduction
- Introduction to Machine Learning and the ML Pipeline
- Introduction to Amazon SageMaker
- Problem Formulation
- Problem Formulation (continued)
- Preprocessing
- Model Training
- Model Evaluation
- Feature Engineering and Model Tuning
- Deployment

Audience

This course is designed for:

- Developers
- Solutions architects
- Data engineers
- Anyone who wants to learn about the ML pipeline via Amazon SageMaker, even if you have little to no experience with machine learning

Prerequisites

Before taking this course, students should have:

- Basic knowledge of Python
- Basic understanding of working in a Jupyter notebook environment
- Basic understanding of AWS Cloud infrastructure (Amazon S3 and Amazon CloudWatch)

Duration

Four days