

The Machine Learning Pipeline on AWS (AWS-ML-PL) (5 Day)

Overview

Taking ML models from conceptualization to production is typically complex and time-consuming. You have to manage large amounts of data to train the model, choose the best algorithm for training it, manage the compute capacity while training it, and then deploy the model into a production environment. SageMaker reduces this complexity by making it much easier to build 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. Students will learn about each phase of the pipeline from instructor presentations and demonstrations. They will then apply that knowledge to complete the project. By the end of the course, you will have successfully built, trained, evaluated, tuned, and deployed an ML model.

Learners with little to no machine learning experience or knowledge will benefit from this course. Basic knowledge of Statistics will be helpful.

Skills Gained

- 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 using 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

Who Can Benefit

- Developers
- Solutions Architects
- Data Engineers
- Anyone with little to no experience with ML and wants to learn about the ML pipeline using Amazon SageMaker

Prerequisites

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

Course Outline

Day 1

- Module 0: Introduction
- Module 1: Introduction to Machine Learning and the ML Pipeline
- Module 2: Introduction to Amazon SageMaker
- Module 3: Problem Formulation

Day 2

- Module 4: Preprocessing

Day 3

- Module 5: Model Training
- Module 6: Model Evaluation

Day 4

- Module 7: Feature Engineering and Model Tuning
- Module 8: Deployment