Practical Data Science with Amazon SageMaker

AWS Classroom Training

Course description

Artificial intelligence and machine learning (AI/ML) are becoming mainstream. In this course, you will spend a day in the life of a data scientist so that you can collaborate efficiently with data scientists and build applications that integrate with ML. You will learn the basic process data scientists use to develop ML solutions on Amazon Web Services (AWS) with Amazon SageMaker. You will experience the steps to build, train, and deploy an ML model through instructor-led demonstrations and labs.

- Course level: Intermediate
- Duration: 8 hours

Activities

This course includes presentations, hands-on labs, and demonstrations.

Course objectives

In this course, you will learn to:

- Discuss the benefits of different types of machine learning for solving business problems
- Describe the typical processes, roles, and responsibilities on a team that builds and deploys ML systems
- Explain how data scientists use AWS tools and ML to solve a common business problem
- Summarize the steps a data scientist takes to prepare data
- Summarize the steps a data scientist takes to train ML models
- Summarize the steps a data scientist takes to evaluate and tune ML models
- Summarize the steps to deploy a model to an endpoint and generate predictions
- Describe the challenges for operationalizing ML models
- Match AWS tools with their ML function

Intended audience

This course is intended for:

- Development Operations (DevOps) engineers
- Application developers

Prerequisites

We recommend that attendees of this course have:

- AWS Technical Essentials
- Entry-level knowledge of Python programming
- Entry-level knowledge of statistics

Practical Data Science with Amazon SageMaker

AWS Classroom Training

Course outline

Module 1: Introduction to Machine Learning
- Benefits of machine learning (ML)
- Types of ML approaches
- Framing the business problem
- Processes, roles, and responsibilities for ML projects

Module 2: Preparing a Dataset
- Data analysis and preparation
- Data preparation tools
- Amazon CodeWhisperer
- Demonstration: Review Amazon SageMaker Studio and Notebooks
- Hands-On Lab: Data Preparation with SageMaker Data Wrangler

Module 3: Training a Model
- Steps to train a model
- Model training in Amazon SageMaker
- Hands-On Lab: Training a Model with Amazon SageMaker

Module 4: Evaluating and Tuning a Model
- Model evaluation and tuning
- Metrics for supervised learning
- Model tuning and hyperparameter optimization
- Hands-On Lab: Model Tuning and Hyperparameter Optimization with Amazon SageMaker

Module 5: Deploying a Model
- Model deployment
- Hands-On Lab: Deploy a Model with SageMaker Studio

Module 6: Operational Challenges
- MLOps
- Responsible ML
- Automation
- Coordination
- Monitoring
- Updating models
- Hands-On Lab: Integrating a Web Application with an Amazon SageMaker Model Endpoint

Module 7: Other Model-Building Tools
- No-code ML with Amazon SageMaker Canvas
- Demonstration: Overview of Amazon SageMaker Canvas
- SageMaker Studio Lab