

Amazon SageMaker Studio for Data Scientists

Course description

Amazon SageMaker Studio helps data scientists prepare, build, train, deploy, and monitor machine learning (ML) models quickly by bringing together a broad set of capabilities purpose-built for ML. This course prepares experienced data scientists to use the tools that are part of SageMaker Studio to improve productivity at every step of the ML lifecycle.

- Course level: Advanced
- Duration: 3 days

Activities

This course includes presentations, demonstrations, practice labs, discussions, and a capstone project.

Course objectives

In this course, you will learn to:

- Accelerate the preparation, building, training, deployment, and monitoring of ML solutions by using Amazon SageMaker Studio.

Intended audience

This course is intended for:

- Experienced data scientists who are proficient in ML and deep learning fundamentals. Relevant experience includes using ML frameworks, Python programming, and the process of building, training, tuning, and deploying models.

Prerequisites

We recommend that all students complete the following AWS course prior to attending this course:

- AWS Tech Essentials (1-day AWS instructor led course)

We recommend students who are not experienced data scientists complete the following two courses followed by 1-year on-the-job experience building models prior to taking this course:

- The Machine Learning Pipeline on AWS (4-day AWS instructor led course)
- Deep Learning on AWS (1-day AWS instructor led course)

Enroll today

Visit aws.training to find a class today.

Course outline

Module 1: Amazon SageMaker Setup and Navigation

- Launch SageMaker Studio from the AWS Service Catalog.
- Navigate the SageMaker Studio UI.
- Demo 1: SageMaker UI Walkthrough
- Lab 1: Launch SageMaker Studio from AWS Service Catalog

Module 2: Data Processing

- Use Amazon SageMaker Studio to collect, clean, visualize, analyze, and transform data.
- Set up a repeatable process for data processing.
- Use SageMaker to validate that collected data is ML ready.
- Detect bias in collected data and estimate baseline model accuracy.
- Lab 2: Analyze and Prepare Data Using SageMaker Data Wrangler
- Lab 3: Analyze and Prepare Data at Scale Using Amazon EMR
- Lab 4: Data Processing Using SageMaker Processing and the SageMaker Python SDK
- Lab 5: Feature Engineering Using SageMaker Feature Store

Module 3: Model Development

- Use Amazon SageMaker Studio to develop, tune, and evaluate an ML model against business objectives and fairness and explainability best practices.
- Fine-tune ML models using automatic hyperparameter optimization capability.
- Use SageMaker Debugger to surface issues during model development.
- Demo 2: Autopilot
- Lab 6: Track Iterations of Training and Tuning Models Using SageMaker Experiments
- Lab 7: Analyze, Detect, and Set Alerts Using SageMaker Debugger
- Lab 8: Identify Bias Using SageMaker Clarify

Module 4: Deployment and Inference

- Use Model Registry to create a model group; register, view, and manage model versions; modify model approval status; and deploy a model.
- Design and implement a deployment solution that meets inference use case requirements.
- Create, automate, and manage end-to-end ML workflows using Amazon SageMaker Pipelines.
- Lab 9: Inferencing with SageMaker Studio
- Lab 10: Using SageMaker Pipelines and the SageMaker Model Registry with SageMaker Studio

Module 5: Monitoring

- Configure a SageMaker Model Monitor solution to detect issues and initiate alerts for changes in data quality, model quality, bias drift, and feature attribution (explainability) drift.
- Create a monitoring schedule with a predefined interval.
- Demo 3: Model Monitoring

Module 6: Managing SageMaker Studio Resources and Updates

- List resources that accrue charges.
- Recall when to shut down instances.
- Explain how to shut down instances, notebooks, terminals, and kernels.
- Understand the process to update SageMaker Studio.

Capstone

- The Capstone lab will bring together the various capabilities of SageMaker Studio discussed in this course. Students will be given the opportunity to prepare, build, train, and deploy a model using a tabular dataset not seen in earlier labs. Students can choose among basic, intermediate, and advanced versions of the instructions.
- Capstone Lab: Build an End-to-End Tabular Data ML Project Using SageMaker Studio and the SageMaker Python SDK