

Amazon SageMaker Studio for Data Scientists

AWS Classroom Training

Course description

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

- Course level: Advanced
- Duration: 3 days

Activities

This course includes presentations, hands-on labs, demonstrations, discussions, and a capstone project.

Course objectives

In this course, you will learn to:

- Accelerate the process to prepare, build, train, deploy, and monitor ML solutions using Amazon SageMaker Studio

Intended audience

This course is intended for:

- Experienced data scientists who are proficient in ML and deep learning fundamentals

Prerequisites

We recommend that all attendees of this course have:

- Experience using ML frameworks
- Python programming experience
- At least 1 year of experience as a data scientist responsible for training, tuning, and deploying models
- *AWS Technical Essentials* [digital](#) or [classroom training](#)

Enroll today

Visit aws.training to find a class today.

Amazon SageMaker Studio for Data Scientists

AWS Classroom Training

Course outline

Day 1

Module 1: Amazon SageMaker Studio Setup

- JupyterLab Extensions in SageMaker Studio
- Demonstration: SageMaker user interface demo

Module 2: Data Processing

- Using SageMaker Data Wrangler for data processing
- Hands-On Lab: Analyze and prepare data using Amazon SageMaker Data Wrangler
- Using Amazon EMR
- Hands-On Lab: Analyze and prepare data at scale using Amazon EMR
- Using AWS Glue interactive sessions
- Using SageMaker Processing with custom scripts
- Hands-On Lab: Data processing using Amazon SageMaker Processing and SageMaker Python SDK
- SageMaker Feature Store
- Hands-On Lab: Feature engineering using SageMaker Feature Store

Module 3: Model Development

- SageMaker training jobs
- Built-in algorithms
- Bring your own script
- Bring your own container
- SageMaker Experiments
- Hands-On Lab: Using SageMaker Experiments to Track Iterations of Training and Tuning Models

Day 2

Module 3: Model Development (continued)

- SageMaker Debugger
- Hands-On Lab: Analyzing, Detecting, and Setting Alerts Using SageMaker Debugger
- Automatic model tuning
- SageMaker Autopilot: Automated ML
- Demonstration: SageMaker Autopilot
- Bias detection
- Hands-On Lab: Using SageMaker Clarify for Bias and Explainability
- SageMaker Jumpstart

Module 4: Deployment and Inference

- SageMaker Model Registry
- SageMaker Pipelines
- Hands-On Lab: Using SageMaker Pipelines and SageMaker Model Registry with SageMaker Studio
- SageMaker model inference options

Amazon SageMaker Studio for Data Scientists

AWS Classroom Training

- Scaling
- Testing strategies, performance, and optimization
- Hands-On Lab: Inferencing with SageMaker Studio

Module 5: Monitoring

- Amazon SageMaker Model Monitor
- Discussion: Case study
- Demonstration: Model Monitoring

Day 3

Module 6: Managing SageMaker Studio Resources and Updates

- Accrued cost and shutting down
- Updates

Capstone

- Environment setup
- Challenge 1: Analyze and prepare the dataset with SageMaker Data Wrangler
- Challenge 2: Create feature groups in SageMaker Feature Store
- Challenge 3: Perform and manage model training and tuning using SageMaker Experiments
- (Optional) Challenge 4: Use SageMaker Debugger for training performance and model optimization
- Challenge 5: Evaluate the model for bias using SageMaker Clarify
- Challenge 6: Perform batch predictions using model endpoint
- (Optional) Challenge 7: Automate full model development process using SageMaker Pipeline