

Machine Learning Engineering on AWS

AWS Classroom Training

Course description

Machine Learning (ML) Engineering on Amazon Web Services (AWS) is a 3-day intermediate course designed for ML professionals seeking to learn machine learning engineering on AWS. Participants learn to build, deploy, orchestrate, and operationalize ML solutions at scale through a balanced combination of theory, practical labs, and activities. Participants will gain practical experience using AWS services such as Amazon SageMaker AI and analytics tools such as Amazon EMR to develop robust, scalable, and production-ready machine learning applications.

- Course level: Intermediate
- Duration: 3 days

Activities

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

Course objectives

In this course, you will:

- Explain ML fundamentals and its applications in the AWS Cloud.
- Process, transform, and engineer data for ML tasks by using AWS services.
- Select appropriate ML algorithms and modeling approaches based on problem requirements and model interpretability.
- Design and implement scalable ML pipelines by using AWS services for model training, deployment, and orchestration.
- Create automated continuous integration and delivery (CI/CD) pipelines for ML workflows.
- Discuss appropriate security measures for ML resources on AWS.
- Implement monitoring strategies for deployed ML models, including techniques for detecting data drift.

Intended audience

This course is designed for professionals who are interested in building, deploying, and operationalizing machine learning models on AWS. This could include current and in-training machine learning engineers who might have little prior experience with AWS. Other roles that can benefit from this training are DevOps engineer, developer, and SysOps engineer.

Prerequisites

We recommend that attendees of this course have the following:

- Familiarity with basic machine learning concepts
- Working knowledge of Python programming language and common data science libraries such as NumPy, Pandas, and Scikit-learn
- Basic understanding of cloud computing concepts and familiarity with AWS
- Experience with version control systems such as Git (beneficial but not required)

Machine Learning Engineering on AWS

AWS Classroom Training

Course outline

Day 1

Module 0: Course Introduction

Module 1: Introduction to Machine Learning (ML) on AWS

- Topic A: Introduction to ML
- Topic B: Amazon SageMaker AI
- Topic C: Responsible ML

Module 2: Analyzing Machine Learning (ML) Challenges

- Topic A: Evaluating ML business challenges
- Topic B: ML training approaches
- Topic C: ML training algorithms

Module 3: Data Processing for Machine Learning (ML)

- Topic A: Data preparation and types
- Topic B: Exploratory data analysis
- Topic C: AWS storage options and choosing storage

Module 4: Data Transformation and Feature Engineering

- Topic A: Handling incorrect, duplicated, and missing data
- Topic B: Feature engineering concepts
- Topic C: Feature selection techniques
- Topic D: AWS data transformation services
- Lab 1: Analyze and Prepare Data with Amazon SageMaker Data Wrangler and Amazon EMR
- Lab 2: Data Processing Using SageMaker Processing and the SageMaker Python SDK

Day 2

Module 5: Choosing a Modeling Approach

- Topic A: Amazon SageMaker AI built-in algorithms
- Topic B: Selecting built-in training algorithms
- Topic C: Amazon SageMaker Autopilot
- Topic D: Model selection considerations
- Topic E: ML cost considerations

Module 6: Training Machine Learning (ML) Models

- Topic A: Model training concepts
- Topic B: Training models in Amazon SageMaker AI
- Lab 3: Training a model with Amazon SageMaker AI

Module 7: Evaluating and Tuning Machine Learning (ML) models

- Topic A: Evaluating model performance

Machine Learning Engineering on AWS

AWS Classroom Training

- Topic B: Techniques to reduce training time
- Topic C: Hyperparameter tuning techniques
- Lab 4: Model Tuning and Hyperparameter Optimization with Amazon SageMaker AI

Module 8: Model Deployment Strategies

- Topic A: Deployment considerations and target options
- Topic B: Deployment strategies
- Topic C: Choosing a model inference strategy
- Topic D: Container and instance types for inference
- Lab 5: Shifting Traffic A/B

Day 3

Module 9: Securing AWS Machine Learning (ML) Resources

- Topic A: Access control
- Topic B: Network access controls for ML resources
- Topic C: Security considerations for CI/CD pipelines

Module 10: Machine Learning Operations (MLOps) and Automated Deployment

- Topic A: Introduction to MLOps
- Topic B: Automating testing in CI/CD pipelines
- Topic C: Continuous delivery services
- Lab 6: Using Amazon SageMaker Pipelines and the Amazon SageMaker Model Registry with Amazon SageMaker Studio

Module 11: Monitoring Model Performance and Data Quality

- Topic A: Detecting drift in ML models
- Topic B: SageMaker Model Monitor
- Topic C: Monitoring for data quality and model quality
- Topic D: Automated remediation and troubleshooting
- Lab 7: Monitoring a Model for Data Drift

Module 12: Course Wrap-up