

Machine Learning Engineering on AWS with AWS Jam

AWS Classroom Training

Course description

Machine Learning (ML) Engineering on Amazon Web Services (AWS) is a 4-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.

The final day is an AWS Jam, a gamified event, with teams competing to score points by completing a series of challenges according to established best practices based on concepts covered in the course. You get to experience a wide range of AWS services in a series of real-world scenarios that represent common operational and troubleshooting tasks. The end result is developing, enhancing, and validating your skillsets in the AWS Cloud through real-world problem solving, exploring new services, features, and understanding how they interoperate.

- Course level: Intermediate
- Duration: 4 days

Activities

This course includes instructor presentation, hands-on labs, demonstrations, group exercises, and team-based gamified challenges.

Course objectives

In this course, you will learn to:

- 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
- Apply the concepts to build ML pipelines for building, deploy, and monitoring ML models leveraging Amazon SageMaker AI features.

Intended audience

This course is intended for:

Machine Learning Engineering on AWS with AWS Jam

AWS Classroom Training

- 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:

- 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 with AWS Jam

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: Amazon SageMaker Autopilot
- Topic C: Selecting built-in training algorithms
- 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
- Topic B: Techniques to reduce training time
- Topic C: Hyperparameter tuning techniques

Machine Learning Engineering on AWS with AWS Jam

AWS Classroom Training

- 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

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

Day 4

Module 12: AWS Jam

- Participate in team-based challenges in a real AWS environment
- Compete with your colleagues in a gamified, hands-on learning experience
- Apply your learning from the course on various AWS services

Machine Learning Engineering on AWS with AWS Jam

AWS Classroom Training

Number	Challenge	Description	Services	Difficulty	Est Duration *
1	Bits, Bytes, and Beyond: Fine-Tuning Expedition	Learn how to fine-tune a LLM by running Python code in Amazon SageMaker Notebook and modifying a Lambda function.	Amazon SageMaker Amazon S3	Medium	45 mins
2	Build an Automated ML Pipeline with Amazon SageMaker	Create and execute a machine learning pipeline in Amazon SageMaker that automates model training, evaluation, and registration.	Amazon SageMaker Studio Amazon SageMaker Pipelines	Medium	45 mins
3	This Instructor is in danger: Amazon SageMaker Data Wrangler in Class	Use Amazon SageMaker Data Wrangler to address and resolve issues related to importing and processing Customer Satisfaction (CSAT) data.	Amazon SageMaker Amazon S3 Amazon Lambda Amazon VPC	Medium	45 mins
4	Apply A/B Testing machine learning models with Amazon SageMaker MLOps	How to implement A/B Testing in Amazon SageMaker	Amazon SageMaker Studio	Medium	30 mins
5	Sports Predictor: Will They Win The Game?	Build an end-to-end machine learning model using Amazon SageMaker. Learn the fundamental steps of creating and deploying ML models to predict match outcomes, and essential SageMaker features.	Amazon SageMaker AI Amazon SageMaker AI Canvas Amazon SageMaker Data Wrangler Amazon S3	Medium	60 mins
6	Building Credit Risk Assessment Model Responsibly	Develop and deploy a responsible AI model using Amazon	Amazon SageMaker Studio	Hard	60 mins

Machine Learning Engineering on AWS with AWS Jam

AWS Classroom Training

		SageMaker AI to predict credit risk.	Amazon SageMaker Clarify Amazon S3		
7	Talent Retention Analytics: Develop an XGBoost Model for HR Insights on SageMaker	Build an ML model with SageMaker and XGBoost to predict employee attrition, helping HR implement retention strategies and reduce turnover costs.	Amazon SageMaker AI	Medium	30 mins
8	Cheaters not allowed!	Learn how to handle datasets and create an ML model using no code. Deploy and share ML models using SageMaker Canvas.	Amazon SageMaker Canvas Amazon SageMaker Data Wrangler Amazon S3	Medium	30 mins

* Duration times are estimated.