

AWS Certified Machine Learning Engineer - Associate (MLA-C01)

Harness the power of AI and put it to work for your company's needs. Prepare data, train AI and machine learning models, and deploy them securely with this AWS Machine Learning course. You'll explore AI services like AWS SageMaker and AWS Glue and learn how to apply them in your everyday work. Don't just study the entire AWS data pipeline — get hands-on with AI tools that simplify, accelerate, and enhance the work you're already doing in the AWS ecosystem. At the same time, prepare for the AWS Certified Machine Learning Engineer – Associate, a career-advancing credential from the world's cloud and AI powerhouse.

[CBT Nuggets course material](#) →

WEEK 1

Getting Started with the AWS ML Ecosystem

- Certification Landscape and Overview
- Course Approach and Methods
- Overview of the AWS ML Ecosystem
- Creating Our Account
- Enabling MFA
- Setting up CloudTrail
- Set up a Budget Alert
- Validation

Securing Cloud Resources with IAM

- Skill Introduction
- Shared Responsibility Model
- IAM Basics
- IAM Roles and Policies
- Role-Based and Attribute-Based Access Control
- A Cautionary Tale
- Setting up an IAM Identity Store
- Validation

Network Protections and Isolation

- Skill Introduction
- Virtual Private Clouds
- Connecting to the AWS Network
- Exposing Stuff to the Cloud
- Creating a VPC
- Setting Up Subnets

Internet Access and Routings

Security Groups and Network Access Control Lists

Creating an Endpoint

Validation

Data Storage Options in AWS

Skill Introduction

Cloud Storage as an Enabler

Simple Storage Service

Amazon Glacier

Network Storage

Block Storage

Data Transfer Methods

WEEK 2

Preparing Our Datasets

Importing Data into a Directory Bucket

Validation

Database Options on AWS

Skill Introduction

Amazon Relational Database Service

Amazon Aurora

Amazon DynamoDB

Amazon ElastiCache

Other Specialized AWS Databases

RDS Setup

RDS Failover

Teardown

Validation

Data Wrangling Basics with AWS Glue

Skill Introduction

Common Data Formats for ML

Meet AWS Glue

Setting up Security Roles

Our First Glue Project

Transforms with AWS Glue

Using Notebooks

Teardown

Validation

Advanced Data Handling with AWS Glue

Skill Introduction

Aquiring Supplemental Weather Data

Merging Datasets

WEEK 3

Cleaning Up Our UFO Sightings Data

SQL Transforms

Fixing the Description Values

Building a Data Catalog

Validation

Data Exploration with AWS Glue DataBrew

Skill Introduction

Meet AWS DataBrew

Data Profiling

Data Aggregation

Binning

One-Hot Encoding

Teardown

Validation

Setting Up a SageMaker Domain

Skill Introduction

The Many Faces of SageMaker

Setting Up Our VPC

Creating Our SageMaker Domain

Updating Our SageMaker Role

Adding User Profiles

SageMaker Studio

SageMaker Canvas

Venturing to the Other SageMaker

Teardown

Validation

Labeling and Data Augmentation with SageMaker

Skill Introduction

Data Prep for Image Classification

Creating Our SageMaker Domain

Using Ground Truth to Label Our Dataset

SageMaker Data Wrangler

Reviewing Our Augmented Data and Teardown

Validation

Handling Inbound Streaming Data

Skill Introduction

Inbound Streaming Data Options on AWS

Our Kinesis Project

Creating a Kinesis Data Stream

Dad Joke Generator

Amazon Data Firehose

Using EventBridge for Delivery to SQS

Teardown

Validation

Creating a Feature Store

Skill Introduction

What's the Story with Feature Store?

Spinning Up our SageMaker Domain...again

Using Athena to Query Data

Creating Our Feature Store

Data Wrangling UFO Sightings

Feature Store Population

Teardown

Validation

Framing Business Problems

Skill Introduction

Problem Framing

Common Supervised Approaches

Common Unsupervised and Hybrid Approaches

WEEK 4

Model Validation

Our Project

Creating our SageMaker AI Domain

Jupyter Notebook Fun with Two Models

Are Our Models Any Good?

Validation

Using SageMaker Built-In Algorithms

Skill Introduction

Requesting a Quota Increase

Managed ML Services

Fun With Managed ML Services

Amazon Bedrock

Augmenting Our Image Data (Part Deux)

Training Our Model

Deploying an Endpoint and Testing

Teardown

Validation

SageMaker JumpStart and Hyperparameter Tuning

Skill Introduction

The ROC and AUC Metric

SageMaker JumpStart

Making Use of a JumpStart Model

Hyperparameter Tuning Approaches

Finding the Best Fit

Reviewing Our Results

Validation

SageMaker Script Mode and Importing Custom Models

Skill Introduction

Containers Everywhere

Preparing Our Training Data

WEEK 5

Training with a Built-In Algorithm

Training with a Custom Script

Comparison of Both Methods

Importing a Local Model

Validation

Analyzing Bias, Explainability & Drift on AWS

Skill Introduction

Bias

Explainability

Drift

Visual Analysis of Our Data

Explainability for Our Model

Drift Monitoring

Validation

Automating ML Workflows with Pipelines

Skill Introduction

Deployment Types

Creating Our Pipeline

Deploying Our Model Version 1

Retraining Our Model for Version 2

Deploying Version 2 Alongside Version 1

Ramping Up Version 2

Validation

Cost Optimization Methods for SageMaker

Skill Introduction

Cost Optimization Methods

Training with Spot Instances

Spot Savings and Neo Compilation

MultiModel Endpoints

Serverless Inference

Validation

Data Stewardship and Compliance

Skill Introduction

Governance and Compliance

Data Stewardship and Management

AWS Tools for Compliance and Governance

Amazon Macie

A Tour of AWS Artifact

AWS Config

Automatic Remediation with AWS Config

Validation