

AWS Machine Learning with Amazon SageMaker AI

Module 1 – Introduction to Machine Learning

- Understand Artificial Intelligence and Machine Learning
- Difference between AI, Machine Learning, and Deep Learning
- Understand models, training, and predictions
- Understand features and labels
- Supervised and unsupervised learning
- Classification, regression, and clustering
- Examples of Machine Learning use cases in organizations

Module 2 – Machine Learning on AWS

- Overview of the AWS Machine Learning ecosystem
- Introduction to Amazon SageMaker AI
- Role of Amazon S3 in a Machine Learning project
- Introduction to IAM roles and permissions
- Introduction to Amazon CloudWatch for monitoring
- Understand the architecture of a Machine Learning solution on AWS

Module 3 – Preparing the AWS Environment

- Access the AWS Management Console
- Select an AWS Region
- Create and configure an Amazon S3 bucket
- Upload a dataset to Amazon S3
- Access Amazon SageMaker AI
- Configure the required permissions
- Explore the SageMaker working environment

Module 4 – Understanding and Preparing Data

- Identify the data used for model training
- Understand rows, columns, and variables
- Identify numerical and categorical variables
- Identify the target variable
- Handle missing values
- Identify and remove duplicates
- Understand outliers
- Select relevant features

Module 5 – Transforming Data for Machine Learning

- Clean a dataset

- Transform data types
- Encode categorical variables
- Normalize data
- Introduction to Feature Engineering
- Create a dataset ready for model training
- Discover data preparation tools in SageMaker

Module 6 – Preparing Training and Test Data

- Understand the purpose of training data
- Understand the purpose of validation data
- Understand the purpose of test data
- Split a dataset into multiple subsets
- Understand model generalization

Module 7 – Training a Model with Amazon SageMaker AI

- Understand how Machine Learning algorithms work
- Select an algorithm appropriate for the problem
- Understand SageMaker Training Jobs
- Configure a training job in SageMaker
- Understand hyperparameters
- Select computing resources
- Launch a training job
- Monitor the training process
- Retrieve the trained model

Module 8 – Evaluating Model Performance

- Understand why models must be evaluated
- Measure Accuracy
- Understand Precision
- Understand Recall
- Discover the F1 Score
- Understand the confusion matrix
- Compare predictions with expected results
- Interpret model performance

Module 9 – Understanding Underfitting and Overfitting

- Understand Underfitting
- Understand Overfitting
- Identify a model that does not generalize well
- Understand how data affects model performance
- Understand the impact of hyperparameters
- Improve a model progressively

Module 10 – Optimizing a Machine Learning Model

- Compare multiple models

- Adjust hyperparameters
- Improve the features used by the model
- Introduction to Hyperparameter Tuning
- Retrain a model
- Compare performance results
- Select the best-performing model

Module 11 – Deploying a Model with Amazon SageMaker AI

- Understand the concept of inference
- Create a deployable model
- Understand SageMaker endpoints
- Configure an endpoint
- Deploy a model
- Send new data to the model
- Retrieve predictions
- Discover real-time inference
- Introduction to batch and serverless inference

Module 12 – Integrating a Model into an AWS Application

- Understand how an application consumes a Machine Learning model
- Call a SageMaker endpoint using Python
- Introduction to AWS Lambda
- Introduction to Amazon API Gateway
- Understand an API Gateway, Lambda, and SageMaker architecture
- Test predictions from an application

Module 13 – Introduction to MLOps on AWS

- Understand the Machine Learning model lifecycle
- Version Machine Learning models
- Introduction to SageMaker Model Registry
- Register different versions of a model
- Understand model approval workflows
- Manage model deployment to production
- Understand the model retraining process

Module 14 – Monitoring a Model in Production

- Introduction to Amazon CloudWatch
- Review endpoint metrics
- Review logs
- Understand Data Drift
- Understand Model Drift
- Introduction to SageMaker Model Monitor
- Identify model performance degradation
- Determine when a model should be retrained

Module 15 – Security and Cost Management

- Understand IAM roles used by SageMaker
- Apply the principle of least privilege
- Control access to data stored in Amazon S3
- Understand data protection
- Identify the main SageMaker cost drivers
- Understand Training Job costs
- Understand endpoint costs
- Remove unused AWS resources after labs

Module 16 – Introduction to Generative AI on AWS

- Understand the difference between traditional Machine Learning and generative AI
- Understand Foundation Models
- Understand Large Language Models
- Introduction to Amazon Bedrock
- Understand the concept of prompting
- Introduction to embeddings
- Introduction to Retrieval-Augmented Generation (RAG)
- Identify use cases for SageMaker AI and Amazon Bedrock

Module 17 – Hands-on Workshop: End-to-End Machine Learning Project

- Import a dataset into Amazon S3
- Explore and prepare the data
- Split the data into training and test datasets
- Configure Amazon SageMaker AI
- Train a Machine Learning model
- Evaluate model performance
- Improve the model
- Deploy the model to a SageMaker endpoint
- Send new data to the deployed model
- Generate and interpret predictions
- Review metrics and logs
- Remove AWS resources used during the workshop