

Deep Learning on AWS

In this one-day course, you will learn cloud-based deep learning solutions on the AWS platform. You will learn how to run your models on the cloud using Amazon EC2–based deep learning Amazon Machine Image (AMI) and Apache MXNet on AWS frameworks. In addition, you will learn how to use Amazon SageMaker and deploy your deep learning models using AWS services like AWS Lambda and Amazon Elastic Container Service (Amazon ECS)—all while designing intelligent systems on AWS.

Course Objectives

In this course, you will learn to:

- Define machine learning and deep learning
- Identify the concepts in a deep learning ecosystem
- Leverage Amazon SageMaker and MXNet programming frameworks for deep learning workloads
- Fit AWS solutions for deep learning deployments

Intended Audience

This course is intended for:

- Developers who are responsible for developing deep learning applications
- Developers who want to understand concepts behind deep learning and how to implement a deep learning solution on AWS