

Deep Learning on AWS

AWS Classroom Training

Course description

In this course, you'll learn about AWS's deep learning solutions, including scenarios where deep learning makes sense and how deep learning works. You'll learn how to run deep learning models on the cloud using Amazon SageMaker and the MXNet framework. You'll also learn to deploy your deep learning models using services like AWS Lambda while designing intelligent systems on AWS.

- Course level: Intermediate
- Duration: 1 day

Activities

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

Course objectives

In this course, you will:

- Learn how to define machine learning (ML) and deep learning
- Learn how to identify the concepts in a deep learning ecosystem
- Use Amazon SageMaker and the MXNet programming framework 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 the concepts behind deep learning and how to implement a deep learning solution on AWS Cloud

Prerequisites

We recommend that attendees of this course have:

- A basic understanding of ML processes
- Knowledge of AWS core services like Amazon EC2 and AWS SDK
- Knowledge of a scripting language like Python

Deep Learning on AWS

AWS Classroom Training

Course outline

Module 1: Machine learning overview

- A brief history of AI, ML, and DL
- The business importance of ML
- Common challenges in ML
- Different types of ML problems and tasks
- AI on AWS

Module 2: Introduction to deep learning

- Introduction to DL
- The DL concepts
- A summary of how to train DL models on AWS
- Introduction to Amazon SageMaker
- Hands-on lab: Spinning up an Amazon SageMaker notebook instance and running a multi-layer perceptron neural network model

Module 3: Introduction to Apache MXNet

- The motivation for and benefits of using MXNet and Gluon
- Important terms and APIs used in MXNet
- Convolutional neural networks (CNN) architecture
- Hands-on lab: Training a CNN on a CIFAR-10 dataset

Module 4: ML and DL architectures on AWS

- AWS services for deploying DL models (AWS Lambda, AWS IoT Greengrass, Amazon ECS, AWS Elastic Beanstalk)
- Introduction to AWS AI services that are based on DL (Amazon Polly, Amazon Lex, Amazon Rekognition)
- Hands-on lab: Deploying a trained model for prediction on AWS Lambda