

Deep Learning on AWS

Learn via: **Classroom/AFA**

Duration: **1 Day**

Overview

Deep Learning on AWS is a one-day course that introduces you to cloud-based Deep Learning solutions on Amazon Web Services (AWS). The training will detail how deep learning is useful and explain its different concepts. This course also teaches you how to run your models on the cloud using Amazon Elastic Compute Cloud (Amazon EC2)-based Deep Learning Amazon Machine Image (AMI) and MXNet framework. In addition, you will gain a better understanding of deploying your deep learning models using AWS services like AWS Lambda and Amazon EC2 Container Service (Amazon ECS) while designing intelligent systems on AWS, based on Deep Learning.

Prerequisites

We recommend that attendees of this course have the following prerequisites:

- Basic understanding of machine learning processes
- Basic understanding of AWS core services like Amazon EC2 and knowledge of AWS SDK
- Basic knowledge of a scripting language e.g. Python

Who Should Attend

- Developers 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

What You Will Learn

This course teaches you how to:

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

Outline

This course covers the following concepts:

- Introduction to Machine Learning
- Introduction to Deep Learning
 - Lab 1: Setting up a Deep Learning AMI instance and running a multi-layer perceptron model
- Introduction to MXNet on AWS
 - Lab 2: Running a convolutional neural network model and predicting images on CIFAR-10 dataset
- Deploying Deep Learning Workloads on AWS
 - Lab 3: Deploying a Deep Learning model for predicting images using AWS Lambda