

Machine Learning with Amazon SageMaker

Course Content

Introduction

- Overview of key Machine Learning and Deep Learning concepts
- Getting about in AWS
- Overview of SageMaker features
- Taking your first look at SageMaker studio

Preparing your dataset

- Identifying your data and articulating your problem
- Format data for consistency
- Cleaning and validating your data
- Uploading to SageMaker

Data Analysis

- Clustering
- Trend analysis
- Finding other relationships between different types of data

Data Visualisation

- Frequency tables
- Cross-tabulation tables
- Bar charts
- Line graphs
- Pie charts
- Heat Maps
- Scatter graphs.

Training your Model

- Creating a Training job
- Assigning Compute resources
- Selecting the right algorithm
- Overview of using custom code (Python, TensorFlow)

Deploying your Model

- SageMaker Hosting Services
- Configuring and creating an HTTPS endpoint

SageMaker Batch Transforms

- Making inferences from your dataset
- Indexing and real-time indices
- Using Batch Transform to preprocess data to train a new model

Validating your Model

- SageMaker Debugger
- Offline testing
- Online testing
- Validating using a holdout set

Model Tuning

- Defining metrics
- Hyperparameter tuning
- Automatic model tuning

Deploying and sharing using SageMaker Feature Store

- Creating, Storing and Sharing Features
- Online / Offline
- Feature Groups
- Discovery
- Batch Inference
- Feature Data Ingestion