

Practical Data Science with SageMaker

Course Benefits & Agenda



Overview

1

Course Benefits

1

Agenda

2

Overview

This bootcamp is made up of nine modules (not including the introductory module). Each of the nine core modules includes either presentation content and/or a demo/notebook exercise. The course is not meant to teach the ML pipeline to students, because that is a prerequisite. It is meant to provide a full-day, guided, hands-on experience of the ML pipeline. To that end, the module PPT slides are there for you to use at the beginning of each module before getting into the notebook demos and exercises – which is where most of the time will be spent.

Course Benefits

This course teaches you how to:

- An introduction to datasets
- A review of the machine learning process
- Amazon SageMaker (workflow, features, and a demo)
- Setting up Qwiklabs and
- Launching a Jupyter notebook

Agenda

Day 1

Module	Topic
Module 1	Machine Learning and Artificial Intelligence Setting up Qwiklabs and launching a Jupyter Notebook Module
Module 2	The Business Challenge and Preparing a Dataset
Exercise 1	Relating features to target variables Exercise 2: Examining relationships between attributes
Module 3	Training and Evaluating a Model
Exercise 2	Finish the estimator definition
Exercise 3	Setting hyperparameters
Exercise 4	Deploying the mode
Module 4	Automatically Tune the Model
Exercises 6 – 9	Tuning jobs
Module 5	Deploying the Best Model to an Endpoint, A/B Testing, and Auto Scaling
Module 6	Production Readiness
Exercises 10 – 11	Setting up AWS Auto Scaling
Module 7	Relative Cost of Errors
Module 8	Amazon SageMaker Architecture
Module 9	Amazon SageMaker Features