# Practical Data Science with Amazon SageMaker

# Course Content

## Course Description:
In this course, you will cover how to solve a real-world use case with Machine Learning (ML) and produce actionable results using Amazon SageMaker. This course walks through the stages of a typical data science process for Machine Learning from analyzing and visualizing a dataset to preparing the data, and feature engineering.  Individuals will also learn practical aspects of model building, training, tuning, and deployment with Amazon SageMaker. Real life use case includes customer retention analysis to inform customer loyalty programs.

## Course Objectives:
In this course, you will:
- Prepare a dataset for training
- Train and evaluate a Machine Learning model
- Automatically tune a Machine Learning model
- Prepare a Machine Learning model for production
- Think critically about Machine Learning model results

## Prerequisites:
- Familiarity with Python programming language
- Basic understanding of Machine Learning

## Target Audience:
- Developers
- Data Scientists

## Topics:

**Lesson 1: Introduction to machine learning**
- Types of ML
- Job Roles in ML
- Steps in the ML pipeline

**Lesson 2: Introduction to data prep and SageMaker**
- Training and test dataset defined
- Introduction to SageMaker

**Demonstration:** SageMaker console
**Demonstration:** Launching a Jupyter notebook

# Practical Data Science with Amazon SageMaker

Course ID #: 7000-423-ZZ-Z

Hours: 7

**Lesson 3: Problem formulation and dataset preparation**
- Business challenge: Customer churn
- Review customer churn dataset

**Lesson 4: Data analysis and visualization**

**Demonstration:** Loading and visualizing your dataset
**Exercise 1:** Relating features to target variables
**Exercise 2:** Relationships between attributes
**Demonstration:** Cleaning the data

**Lesson 5: Training and evaluating a model**
- Types of algorithms
- XGBoost and SageMaker

**Lab:** Training the data
**Exercise 3:** Finishing the estimator definition
**Exercise 4:** Setting hyper parameters
**Exercise 5:** Deploying the model
Demonstration: hyper parameter tuning with SageMaker
**Demonstration:** Evaluating model performance

**Lesson 6: Automatically tune a model**
- Automatic hyper parameter tuning with SageMaker

**Exercise 6-9:** Tuning jobs

**Lesson 7: Deployment / production readiness**
- Deploying a model to an endpoint
- A/B deployment for testing
- Auto Scaling

**Demonstration:** Configure and test auto scaling
**Demonstration:** Check hyper parameter tuning job
**Demonstration:** AWS Auto Scaling
**Exercise 10-11:** Set up AWS Auto Scaling

**Lesson 8: Relative cost of errors**
- Cost of various error types

**Demonstration:** Binary classification cutoff

**Lesson 9: Amazon SageMaker architecture and features**
- Accessing Amazon SageMaker notebooks in a VPC
- Amazon SageMaker batch transforms
- Amazon SageMaker Ground Truth
- Amazon SageMaker Neo

**Register for this class by visiting us at:**
**www.tcworkshop.com or calling us at 800-639-3535**