

Machine Learning with Sagemaker (AWS)

Course Summary

Description

This course is intended for data scientists and software engineers. It maintains an optimal balance of theory and practice. For each machine learning concept, we first discuss the foundations, its applicability and limitations. Then we explain the implementation and use, and specific use cases. This is achieved through a combination of about 50% lecture, 50% lab work.

Amazon SageMaker is a fully managed machine learning service. The course combines overview and understanding of Machine Learning concepts with specific implementation in SageMaker. In addition, it brings in other tools outside of SageMaker when required.

Machine Learning (ML) is the killer app for Big Data. Amazon Machine Learning brings the power of ML to a regular programmer and provides ML as a service. However, to use ML effectively, one needs to understand the models used and how to utilize them on Amazon.

Objectives

After taking this course, students will:

- Attain thorough understanding of popular machine learning algorithms, their applicability and limitations
- Practice the application of these methods in the Amazon machine learning environment
- Achieve clarity in the real-world use of machine learning by illustrating each method with practical use cases

Topics

- Introductions and overviews
- Supervised Learning
- Unsupervised learning
- Data visualization
- SageMaker

Audience

This course is designed for Data Scientists and Software Engineers.

Prerequisites

- Familiarity with programming in at least one language
- Be able to navigate Linux command line
- Basic knowledge of command line Linux editors (VI / nano)•basic familiarity with AWS (optionally may be Provided in the first day on the course)
-

Duration

Three Days

Machine Learning with Sagemaker (AWS)

Course Outline

- I. *Introductions and overviews***
 - A. Data ETL
 1. Go into one example in detail, implemented on AWS Redshift
 2. Provide pointer to other examples for self-study
 - B. Machine learning
 1. Goals, results, supervised/unsupervised
 2. Which part of ML is implemented in the Amazon Machine Learning
 3. SageMaker (AWS) Overview

- II. *Supervised Learning***
 - A. Linear regression
 - B. Logistic regression and multinomial logistic regression
 - C. SVM, decision trees, random forests, neural networks
 - D. Labs for every section above

- III. *Unsupervised learning***
 - A. K-Means
 - B. Other types of unsupervised learning
 1. Hierarchical clustering
 2. Mixture models
 3. DBSCAN

- IV. *Data visualization***
 - A. Visualization examples for the models above
 - B. Links to other visualizations for self-study

- V. *SageMaker***
 - A. Intro
 - B. SageMaker Details
 1. Using Built-in Algorithms
 2. Using Your Own Algorithms
 3. Using TensorFlow
 4. Using Apache MXNet
 5. Using Apache Spark
 6. Amazon SageMaker Libraries
 7. Authentication and Access Control
 8. Monitoring