Machine Learning in Python with ML Ops.

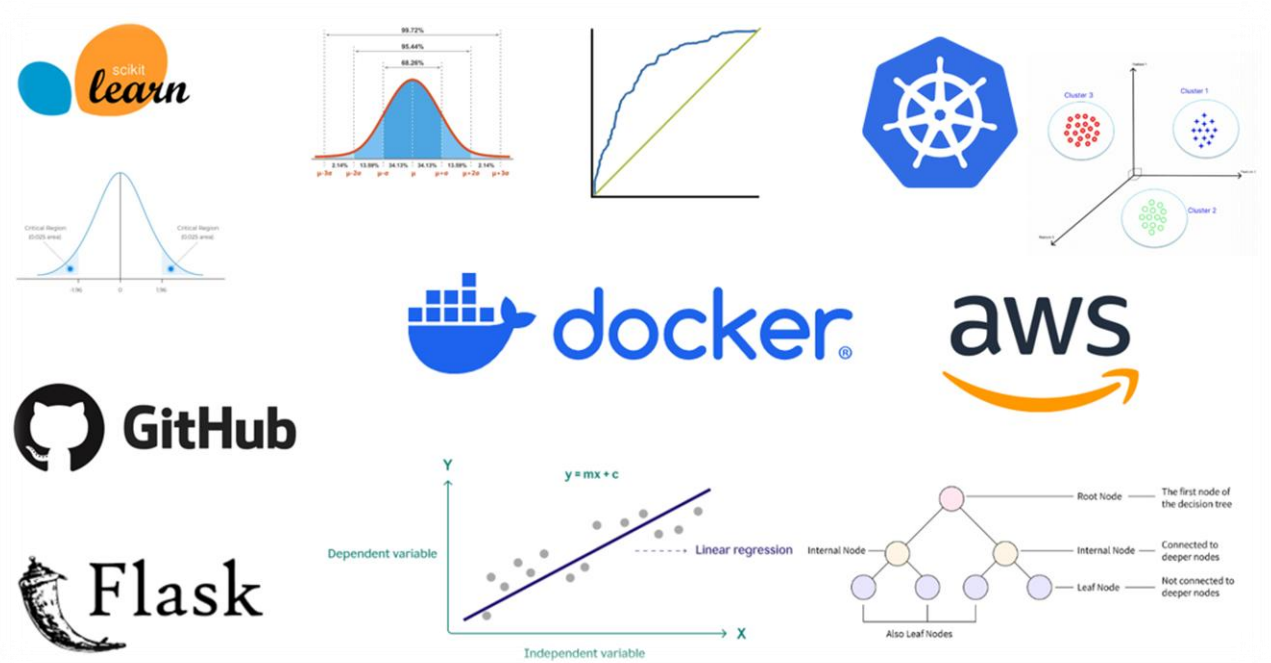

Trainer & Career Counsellor

Saiful Hoda, Exeter MBA, B. Tech. in Computer Science

Ex. Citibank, Refinitiv, Virtu Financials and Nomura Holdings (Hong Kong)

Overview

It is expected that you possess a good understanding of Core concepts in Python and have completed the course on Data Analytics in Python before you begin with Data Sciences in Python.

This is one of the most comprehensive courses out there and along with conceptual understanding, also provides hand on learning experience through practical real-world projects.

Along with Data preprocessing, Model building, evaluation and fine tuning; the course also teaches you deployment and maintenance of Machine learning models. The python libraries and tools taught throughout the course are high in demand.

Upon gaining mastery of the topics in the course, you become a complete Data Scientist in Python. We further help you in shaping your resume and interview preparation, until you ace your dream job.

All the best.

Table of Contents

1. Understanding Machine Learning

2. Simple Linear Regression

2.1. Background

2.2. Gradient Descent Linear Regression

2.2.1. Linear Regression Cost Function

2.2.2. Optimisation Methods

2.2.3. Learning Rate

2.2.4. Gradient Descent in Python

2.2.5. Practise Questions

3. Multiple Linear Regression

3.1. Introduction

3.2. Multicollinearity

3.3. Dealing with Categorical values

3.4. Model Assessment and Comparison

3.5. Feature Selection

3.6. Multiple Linear Regression in Python

3.7. Understanding the Data

3.8. Preparing the Data

3.9. Building the Model

3.10. Residual Analysis and Prediction

3.11. Variable Selection using RFE

3.12. Summarise out learning

3.13. Questions for the interview

3.14. Where are Logistic models used in the industry

3.15. Prediction versus Projection

3.16. Case Study

4. Logistic Regression’

4.1. Univariate Logistic Regression

4.2. Binary Classification

4.3. Sigmoid Curve

4.4. Finding the Best fit Sigmoid curve

4.5. Odds and Log Odds

4.6. Summarise

5. Multi-variate Logistic Regression in Python using a Case Study

5.1. Data Cleaning and Transformation

5.2. Building a Logistic model

5.3. Feature elimination using RFE

5.4. Confusion Matrix and Accuracy

5.5. Manual Feature elimination

5.6. Summarize our Learning

6. Model Evaluation

Introduction

Metric Beyond Accuracy: Sensitivity and Specificity

Sensitivity and Specificity in Python

Understanding the ROC curve

ROC Curve in Python

Finding the Optimal Threshold

Precision and Recall

Making Predictions

Summarise our Learning

7. Multi classification using Logistic Regression

7.1. Introduction

7.2. Business Examples

7.3. Business Versus Multiclass Classification

7.4. One versus One

7.5. One versus Rest

7.6. Application on an Industry Problem

7.7. Understanding of Business Problem Statement

7.8. Application of One versus Rest Classifier

7.9. Analysing the Model Parameters

7.10. Understanding the internal mathematics

7.11. Application of One versus One classifier

7.12. Summary

8. Industry applications of Logistic Regression

8.1. Commonly faced challenges in Implementation

8.2. A closer look at Model Evaluation

8.3. Model validation and importance of stability

8.4. Tracking of model performance over time

8.5. Summary of Learnings

9. Tree Models

9.1. Introduction to decision trees

9.2. Interpreting a Decision Tree

9.3. Building decision trees

9.4. Decision Tree classifier in Python

9.5. Tree models over linear models

9.6. Algorithms for Decision tree construction

9.7. Spitting and Homegeneity

9.8. Impurity measures

9.9. The Gini Index

9.10. Introduction to Hyperparameter Tuning iun Decision Trees

9.11. Disadvantages of Decision Trees

9.12. Tree Truncation

9.13. Building decision trees in Python

9.14. Choosing Tree Hyperparameters in Python

9.15. Decision Tree Regression in Python

9.16. Question and Practise

10. Ensemble and Random Forests

10.1. Ensemble

10.2. Example of some popular Ensembles

10.3. Introduction to Random Forests

10.4. Comprehension – Out Of Bag Error

10.5. Feature Importance in Random Forests

10.6. Random Forests in Python

10.7. Random Forests Regression in Python

10.8. Case Study - Telecom Churn Prediction

11. Clustering: Unsupervised Learning

11.1. Introduction

11.2. Understanding to Clustering

11.3. Customer Segmentation as an example of clustering

12. K-Means Clustering

12.1. Introduction

12.2. Euclidean Distance

12.3. Centroid

12.4. Steps of the Algorithm

12.5. K-Means algorithm

12.6. K Means ++ Algorithm

12.7. Visualizing the K Means algorithm

12.8. Practical consideration in K Means algorithm

12.9. Cluster Tendency

12.10. Practise and Questions

13. Let us build a K-means algorithm in Python

13.1. Introduction

13.2. Data understanding and Cleaning

13.3. Data Preparation

13.4. Creating clusters

13.5. Optimal number of clusters

13.6. Cluster Analysis

13.7. Other Behavioural Segmentation Types

13.8. Summary and Practise

14. Hierarchical Clustering

14.1. Introduction

14.2. Hierarchical Clustering Algorithm

14.3. Interpreting the Dendrogram

14.4. Types of Linkages

14.5. Hierarchical Clustering in Python

14.6. Industry Insights

15. Other Forms of Clustering

15.1. K mode Clustering

15.2. K mode in Python

15.3. K Prototype in Python

15.4. DB scan clustering

15.5. Practise and Exercise

16. Model Selection

16.1. Introduction to Model Selection

16.2. Model and Learning Algorithm

16.3. Simplicity, Complexity and Overfitting

16.4. Bias-Variance Trade-off

16.5. Regularization

16.6. Summary and Practise

17. Model Evaluation

17.1. Regularization versus Hyper parameters

17.2. Model Evaluation and Cross Validation

17.3. Model Evaluation example in Python

17.4. Cross Validation

17.5. Cross Validation in Python

17.6. Cross Validation – Hyperparameter Tuning

18. Boosting

18.1. Introduction

18.2. Ensemble models

18.3. Introduction to Boosting

18.4. Weak learners

18.5. Introduction to Ada Boost algorithm

18.6. Adaboost Distribution and Parameter Calculation

19. Gradient Boosting

19.1. Introduction

19.2. Understanding to Gradient Boosting

19.3. Gradient in Gradient Boosting

19.4. Understanding XG Boost algorithm

20. Unsupervised Learning (PCA)

20.1. Introduction to PCA

20.2. Vectorial Representation of PCA

20.3. Vector Operations

20.4. Matrix multiplication

20.5. Basis

20.6. Change of Basis- Introduction, Demonstration and Calculation

20.7. Introduction to Variance

20.8. Directions of Maximum Variance

20.9. The working of PCA

20.10. Covariance Matrix

20.11. Transforming the Covariance matrix

20.12. Eigen decomposition

20.13. PCA in Python

20.14. Scree Plots

20.15. Dimensionality Reduction

20.16. Improve Performance Model

21. Deploying a Machine Learning Model

21.1. Introduction

21.1.1. Machine learning Pipeline

21.1.2. Software Architecture of Deployment

21.2. Introduction to Git and GitHub

21.2.1. Version Control Systems

21.2.2. Introduction to Git and GitHub

21.2.3. GitHub installation and Configuration

21.2.4. Branch, Commit, Merge and Pull

21.3. Deploying ML models in Production as APIs

21.3.1. API

21.3.2. REST

21.3.3. Deployment of a Data set

21.3.4. Deployment of a Model

21.4. Deploying ML model using Flask and Docker

21.4.1. What is Docker

21.4.2. Terminologies of Docker

21.4.3. Docker Installation

21.4.4. Dockerise a Flask application

21.4.5. Docker CLI commands

21.4.6. Deploy an ML model using Dockers

21.5. Deploying an ML model on AWS (IaaS)

21.5.1. Introduction to AWS

21.5.2. Applications and Advantages of AWS Service

21.5.3. Deploying the Object Identification Model

22. Project Work

- Build and Deploy a Machine learning model
- Indutry relevant use cases

23. Resume Preparation and Job Search