Course Code: W7102GW

Course Title: Supervised Learning: Regression

Description:

This course introduces you to one of the main types of modelling families of supervised Machine Learning: Regression. You will learn how to train regression models to predict continuous outcomes and how to use error metrics to compare across different models. This course also walks you through best practices, including train and test splits, and regularization techniques.

Objectives:

By the end of this course you should be able to:- Differentiate uses and applications of classification and regression in the context of supervised machine learning.

- Describe and use linear regression models.

- Use a variety of error metrics to compare and select a linear regression model that best suits your data.

- Articulate why regularization may help prevent overfitting.

- Use regularization regressions: Ridge, LASSO, and Elastic net.

Prerequisites:

To make the most out of this course, you should have familiarity with programming on a Python development environment, as well as fundamental understanding of Data Cleaning, Exploratory Data Analysis, Calculus, Linear Algebra, Probability, and Statistics.

Duration:

11.2 Hrs

Topics:

1. Introduction to Supervised Machine Learning and Linear Regression

2. Data Splits and Cross Validation

3. Regression with Regularization Techniques: Ridge, LASSO, and Elastic Net

Audience:

This course targets aspiring data scientists interested in acquiring hands-on experience with Supervised Machine Learning Regression techniques in a business setting.