

Course Code: W7103GW

Course Title: Supervised Learning: Classification

Description:

This course introduces you to one of the main types of modeling families of supervised Machine Learning: Classification. You will learn how to train predictive models to classify categorical outcomes and how to use error metrics to compare across different models. The hands-on section of this course focuses on using best practices for classification, including train and test splits, and handling data sets with unbalanced classes.

Objectives:

By the end of this course you should be able to:

- Differentiate uses and applications of classification and classification ensembles.
- Describe and use logistic regression models.
- Describe and use decision tree and tree-ensemble models.
- Describe and use other ensemble methods for classification.
- Use a variety of error metrics to compare and select the classification model that best suits your data.
- Use oversampling and undersampling as techniques to handle unbalanced classes in a data set.

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. Logistic Regression
2. K Nearest Neighbors
3. Support Vector Machines
4. Decision Trees
5. Ensemble Models
6. Modeling Unbalanced Classes

Audience:

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