
Course Code: W7139GS

Course Title: Machine Learning Specialist - Supervised Learning: Regression and Classification

Description:

This course introduces you to two of the main types of modelling families of supervised Machine Learning: Regression and Classification. You start by learning how to train regression models to predict continuous outcomes and how to use error metrics to compare across different models. You then learn how to train predictive models to classify categorical 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. 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.

IBM Customers and Sellers: If you are interested in this course, consider purchasing it as part of one of these Individual or Enterprise Subscriptions:

- IBM Learning for Data and AI Individual Subscription ([SUBR022G](#))
- IBM Learning for Data and AI Enterprise Subscription ([SUBR004G](#))
- IBM Learning Individual Subscription with Red Hat Learning Services ([SUBR023G](#))

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, and use decision tree and tree-ensemble models.
- Use a variety of error metrics to compare and select a linear regression model or classification model that best suits your data.
- Articulate why regularization might help prevent overfitting.
- Use regularization regressions: Ridge, LASSO, and Elastic net.
- Use oversampling 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:

22.4 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
4. Logistic Regression
5. K Nearest Neighbors
6. Support Vector Machines
7. Decision Trees
8. Ensemble Models
9. Modeling Unbalanced Classes

Audience:

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