SYLLABUS
ISL 439 E
KREDİSİ 3 AKTS 5

INTRODUCTION TO MACHINE LEARNING WITH BUSINESS APPLICATIONS

Contact Information

Instructor: Tolga Kaya, kayatolga@itu.edu.tr

Teaching Assistant: Gizem Kaya, kayagizem@itu.edu.tr Hakan Yorulmuş, yorulmus@itu.edu.tr

Course Description

Machine learning is based on techniques that improve their own performance by learning patterns in data. This course is designed to introduce undergraduate level Management Engineering students to methodologies and algorithms for machine learning. R programming labs and weekly assignments will include hands-on practice with various supervised and unsupervised machine learning algorithms. Several software libraries and data sets publicly available will be used to illustrate the business applications of these algorithms. The term project will also give participants an opportunity to undertake the solution of a business problem of their own interest. Students participating the course are expected to have an existing knowledge of statistics, linear algebra, and entry level programming.

Course Objectives

The aim of this course is,

1- to introduce students several fundamental concepts and methods for machine learning.
2- to familiarize the audience with learning algorithms and their applications in business environments.
3- to develop an understanding of the modern data analysis, strengths and weaknesses of many popular machine learning approaches.

Course Outcomes

Students who pass the course will be able to,

1- have a good understanding of the basic concepts, fundamental problems and limitations of machine learning.
2- understand the mathematical models and the assumptions behind various supervised and unsupervised learning algorithms.
3- conduct exploratory analysis, visualize and manipulate data effectively in R programming environment.
4- implement and apply various machine learning techniques in a range of real world business applications.
5- compare the performances of alternative learning algorithms using appropriate tools and criteria.

Assesment Criteria

Final Grade will be made out of,

Midterm 20% - Term Project 25% - Assignments 25% - Final 30%

Course Contents

Week 1

What is machine learning? - Applications of machine learning - Assessing model accuracy - Regression vs classification - Measuring the quality of fit - The bias-vaariance trade off - Supervised vs unsupervised learning.

Week 2

Simple linear regression - Estimating the coefficients - Accuracy of the estimates - Accuracy of the model - Multiple linear regression - Estimating the multiple regression coefficients.

Week 3

Other considerations in the regression model - Qualitative predictors - Extensions of the linear model - Potential problems - Non-parametric regression - K-Nearest Neighbors algorithm.

Week 4

An overview of classification - Logistic regression - The logistic model - Estimating the regression - Making predictions - Multiple logistic regression - Logistic regression for multiple (>2) response classes.

Week 5

Linear discriminant analysis - Using Bayes' Theorem for classification - Linear discriminant analysis for multiple (p>1) features - Quadratic discriminant analysis - K-Nearest Neighbors for classification - A comparison of classification methods.

Week 6

Understanding Naive Bayes - Basic concepts of Bayesian methods - Understanding the joint probability notion - computing conditional probability with Bayes Theorem - The Naive Bayes algorithm - The Laplace estimator - Using numeric features with Naive Bayes.

Week 7

Resampling methods - Cross validation - The validation set approach - Leave-One-Out Cross-Validation - k-Fold Cross-Validation - The bootstrap.

Week 8

Shrinkage methods - Ridge regression - The Lasso - Selecting the tuning parameter - Dimension reduction methods - Principal components regression - Partial least squares - Considerations in high dimensions.

Week 9

Tree based methods - Basics of decision trees - Regression trees - Classification trees - Trees vs linear models - Advantages and disadvantages of trees - Bagging - Random forests – Boosting.

Week 10

Support vector classifiers - Maximal margin classifier - What is a hyperplane? - Support vector machines - Classification with non-linear decision boundaries - SVMs with multiple classes - Relationship to logistic regression.

Week 11

Understanding neural networks - From biological to artifical neurons - Activation functions - Network topology - Number of layers - The direction of information travel - The number of nodes in each layer - Training neural networks with backpropagation – Deep learning.

Week 12

Introduction to unsupervised learning - Principal components analysis - What are principal components? - Interpretation of principal components.

Week 13

Clustering methods - K-means clustering - Hierarchical clustering - Practical issues in clustering.

Week 14

Understanding association rules - The Apriori algorithm for association rule learning - Measuring rule interest - Support and confidence - Building a set of rules with the Apriori principle - Market basket analysis.

Textbooks

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning with applications in R*. New York: Springer.

Lantz, B. (2015). *Machine learning with R*. 2nd Edition. Packt Publishing Ltd.

Alpaydin, E. (2014). *Introduction to machine learning*. MIT press.

Grinberg, N. F.,Reed, R. J. (2013) *Programming exercises for R*. Warwick University.

Supplemantary References

Nielsen, M. A. (2015). *Neural networks and deep learning*. Determination Press.

Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning*. Cambridge: MIT press.

Jockers, M. L. (2014). Text analysis with R for students of literature. New York: Springer.

Friedman, J., Hastie, T., & Tibshirani, R. (2009). *The elements of statistical learning: Data Mining, Inference, and Prediction,* 2nd Edition. New: Springer.

MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge university press York.