Machine Learning
Principles & Applications

Spring 2020
Draft Syllabus

Daniel Egger

Overview

This course focuses on understanding how machine learning (ML) works, and case studies of its successful application to a wide range of problem types, from better forecasting customer behavior, to understanding and responding appropriately to human speech. Students will learn the basic mathematical principles behind establishing reliable ML performance, and have an opportunity to experiment with various ML algorithms (in Python) and observe how they perform on real-world data.

Prerequisite: Fundamentals of Data Science or the equivalent. The course does not require any other prior programming experience.

Required text: Andriy Burkov, *The Hundred-Page Machine Learning Book*, and supplemental case studies.

Learning Objectives

Students:

-will be able to identify the four primary machine learning areas, and what type of projects are suitable for each. They will be able to classify real-world projects as instance-based or model-based, and batch or online learning.

-will be able to apply the K nearest neighbor algorithm to classification problems, and know its strengths and weaknesses relative to decision tree-based classification.

-will have hands-on exposure to the superiority of ensemble methods over most individual classification algorithms, comparing results from an individual model to an ensemble of decision trees (Random Forest).

-will know how to apply a perceptron model to separable data, and how to use support vector machines for both separable (hard margin) and mixed (soft margin) training sets.

-will develop a regression-based forecasting model after pre-processing data for dimensionality reduction through Principal Component Analysis (PCA)

-will be familiar with k means clustering and other unstructured learning algorithms for identifying both hierarchical and non-hierarchical structure in data.

-will be able to identify cases where reinforcement learning is the appropriate machine learning approach.

-will understand the role of backpropagation and hidden layers (deep learning) in solving machine learning problems.

Week One

Introduction: Types of Learning from Data

Week Two

Classification 1:
K Nearest Neighbor and Decision Trees

Week Three

Classification 2:
Ensemble Methods and Random Forests

Week Four

Classification 3:
Support Vector Machines

Week Five

Regression

Week Six

Dimensionality Reduction

Week Seven

Midterm Exam

Week Eight

Clustering and Other Unsupervised Learning Methods
Case Study

Week Nine

Reinforcement Learning and Self-teaching Systems
Case Study

Week Ten

Neural Networks I
Case Study

Week Eleven

Neural Networks II: Deep Learning
Case Study

Week Twelve

Review

Week Thirteen

Project Presentations

Final Exam