Introduction to Support Vector Machines and Neural Networks

A graduate level special-topics course

Shu-Cherng Fang
Spring 2024
fang@ncsu.edu

- Special thanks to Fengming Lin and Xinyu Xu for helping prepare the lecture notes.

Systems Analytics

- Modern information technology (IT) provides industries with the ability to collect vast amounts of data (Big Data) of underlying business systems.

- Data requires more than just flashy dashboards and reports. Information and knowledge (Data Mining and Machine Learning) are the key.

- OR / Optimization models & techniques are useful for machine learning and decision making with dynamic or transactional data (Artificial Intelligence).

- Systems analytics help us make better data-driven business decisions for industries (Analytic Decision Making).

Artificial Intelligence

1. The main goal of artificial intelligence (AI) is to equip machines (computers) with human (or super-human, hopefully) intelligence for dynamic decision making.

2. Computers examine data to extract embedded information (data mining) to form useful knowledge (machine leaning) for right decision making (analytics).

3. Algorithms guide machines to perform each desired task step by step.

Machine Learning – Common approaches

I. Supervised learning for classification and prediction

- Support Vector Classifier & Regression (SVM/SVR)
- Artificial Neural Networks (ANN)

II. Unsupervised learning for clustering and featuring

- Similarity Learning and Sparse Solutions

III. Reinforcement learning in dynamic environment

- Markov Decision Process and Dynamic Programming

Machine learning and optimization

1. Many machine learning problems are formulated as minimization of some loss function that measures discrepancy between the predictions of the model being trained and the actual problem instances, or as maximization of some reward function that affirms an expected decision.

2. One major difference between machine learning and optimization lies in the goal of generalization - optimization intends to minimize the loss/maximize the reward on a set of seen examples while machine learning is concerned with minimizing the loss/maximizing the reward on unseen samples.

Outline of the course

- Course objective
- Pre-requisites
- Contents
- Grading and exams
- Classroom rules
- Textbook and references

Course objective

- An introductory course of supervised learning with the aim to introduce the basic concepts, models, methods and applications of "Support Vector Machines (SVM)" and "Neural Networks (NN)" for machine learning.

- The course is designed for engineering and management students to
  - understand the involved OR/Optimization models for classification, prediction, and generalization of machine learning techniques;
  - experience the process of modeling, computing and analytic decision making in machine learning;
  - work on some data banks with real applications.

Important to know – (4-Not)

- This special topics course is designed mainly for graduate students with proper background, research interests and self-learning capabilities.

- This course does NOT teach you how to concretely model a system problem.
- This course does NOT tell you the details of all theory and algorithms of mathematical optimization.
- This course does NOT tell you how to manage existing data banks.
- This course does NOT tell you how to use Excel, SAS OPT, MATLAB, Python, LINGO, CPLEX, CVX, or any commercial solver and software platform.

Important to know – (3-Intends)

- This course INTENDS to prepare you for understanding the basic SVM, SVR, and ANN models and methodologies involved in supervised leaning.

- This course INTENDS to prepare you for appreciating more in reading existing literatures.

- This course INTENDS to prepare you for conducting further research in your own area.

Prerequisites

1. Linear Algebra and Linear Programming

2. Ability to self-conduct computational experiments using platform solvers and programming languages such as MATLAB, Python, ILOG CPLEX, etc.

Course content - I

Part I: Review (3 Lectures)

A. Linear Algebra

- vectors and matrices, linear transformation, ranks,
  row/column/null spaces,
  eigenvalue/eigenvector & eigen-decomposition,
  singular value & singular value decomposition (SVD),
  principal component analysis (PCA)

B. Optimization

- functions, gradients, unconstrained optimization,
  linear programming, quadratic programming

C. Software Platform and Data

- MATLAB, Python, Data banks

Course content - II

Part II: Support Vector Machines (5 Lectures)

A. Bi-classification

- linear SVM, soft-margin linear SVM, SVM with kernel, kernel-free SVM

B. Multi-classification

- OVA and OVO, Twin SVM

C. Prediction

- Support vector regression (SVR)

Course content - III

Part III: Artificial Neural Networks (5 Lectures)

A. Basic structure of NN

- Neuron and activation function, perceptron, basic feed-forward neural network (NN)

B. Backpropagation and learning

- Error/Reward function, batch vs. online learning, learning algorithm

C. Multi-layer neural network and deep learning

- Scale, feature and computation, ReLU and SGD

D. Radial basis function neural network (RBFN)

E. Convolutional neural network (CNN)

Grading and standard

1. Grading of the course is mainly homework-project based.
2. Homework assignment with each lecture.(60%)
3. Final project with presentation. (40%)

Reference Books

1. Gilbert Strang, Introduction to Linear Algebra, 5th Edition, 2016, Wellesley-Cambridge Press, ISBN-13: 978-0980232776
2. Gilbert Strang, Linear Algebra and Learning from Data, First Edition, 2019, Wellesley Cambridge Press, ISBN-13: 978-0692196380
3. N. Cristianini and J. Shawe-Taylor, An Introduction to Support Vector Machines and Other Kernel-based Learning Methods, 2000 Cambridge University Press, ISBN-0-521-78019-5
4. Charu C. Aggarwal, Neural Networks and Deep Learning, 2018, Springer, ISBN-13: 978-3319944623
5. P. Flach, Machine Learning, 2012, Cambridge University Press, ISBN 978-0-511973-00-0
6. T. Hastie, R. Tibshirani and J. Friedman, The Elements of Statistical Learning, 2nd Edition, 2009, Springer, ISBN 978-0-387-84857-0