



INTRODUCTION TO MACHINE LEARNING

PROF. BALARAMAN RAVINDRAN

Department of Computer Science and Engineering
IIT Madras

INTENDED AUDIENCE: This is an elective course. Intended for senior UG/PG students. BE/ME/MS/PhD

PRE-REQUISITES: We will assume that the students know programming for some of the assignments. If the students have done introductory courses on probability theory and linear algebra it would be helpful. We will review some of the basic topics in the first two weeks as well.

INDUSTRY SUPPORT: Any company in the data analytics/data science/big data domain would value this course.

COURSE OUTLINE :

With the increased availability of data from varied sources there has been increasing attention paid to the various data driven disciplines such as analytics and machine learning. In this course we intend to introduce some of the basic concepts of machine learning from a mathematically well motivated perspective. We will cover the different learning paradigms and some of the more popular algorithms and architectures used in each of these paradigms.

ABOUT INSTRUCTOR :

Prof. Balaraman Ravindran is currently an Professor in Computer Science at IIT Madras and Mindtree Faculty Fellow . He has nearly two decades of research experience in machine learning and specifically reinforcement learning. Currently his research interests are centered on learning from and through interactions and span the areas of data mining, social network analysis, and reinforcement learning.

COURSE PLAN :

Week 0: Probability Theory, Linear Algebra, Convex Optimization - (Recap)

Week 1: Introduction: Statistical Decision Theory - Regression, Classification, Bias Variance

Week 2: Linear Regression, Multivariate Regression, Subset Selection, Shrinkage Methods, Principal Component Regression, Partial Least squares

Week 3: Linear Classification, Logistic Regression, Linear Discriminant Analysis

Week 4: Perceptron, Support Vector Machines

Week 5: Neural Networks - Introduction, Early Models, Perceptron Learning, Backpropagation, Initialization, Training & Validation, Parameter Estimation - MLE, MAP, Bayesian Estimation

Week 6: Decision Trees, Regression Trees, Stopping Criterion & Pruning loss functions, Categorical Attributes, Multiway Splits, Missing Values, Decision Trees - Instability Evaluation Measures

Week 7: Bootstrapping & Cross Validation, Class Evaluation Measures, ROC curve, MDL, Ensemble Methods - Bagging, Committee Machines and Stacking, Boosting

Week 8: Gradient Boosting, Random Forests, Multi-class Classification, Naive Bayes, Bayesian Networks

Week 9: Undirected Graphical Models, HMM, Variable Elimination, Belief Propagation

Week 10: Partitional Clustering, Hierarchical Clustering, Birch Algorithm, CURE Algorithm, Density-based Clustering

Week 11: Gaussian Mixture Models, Expectation Maximization

Week 12: Learning Theory, Introduction to Reinforcement Learning, Optional videos (RL framework, TD learning, Solution Methods, Applications)