

CS501 Machine Learning.

Prerequisites : MT204, CS303

Course Contents:

Introduction

(10%)

Basic concepts. Version space. Decision Trees.

Supervised learning

(35%)

Supervised learning setup. LMS. Logistic regression. Neural networks, Perceptron, Exponential family. Generative learning algorithms. Gaussian discriminant analysis. Naive Bayes. Support vector machines. Model selection and feature selection. Ensemble methods: Bagging, boosting. Evaluating and debugging learning algorithms.

Learning theory

(5%)

Bias/variance tradeoff. Union and Chernoff/Hoeffding bounds. VC dimension. Worst case (online) learning. Practical advice on how to use learning algorithms.

Unsupervised learning

(30%)

Clustering. K-means. EM. Mixture of Gaussians. Factor analysis. PCA (Principal components analysis). ICA (Independent components analysis).

Reinforcement learning and control

(20%)

MDPs. Bellman equations. Value iteration and policy iteration. Linear quadratic regulation (LQR). LQG. Q-learning. Value function approximation. Policy search. Reinforce. POMDPs.

Main Reading

1. Tom Michele, Machine Learning, McGraw-Hill.
2. Ethem Alpaydin, Introduction to Machine Learning, MIT Press.
3. Richard O. Duda, Peter E. Hart, David G. Stork Pattern Classification, Wiley.