Syllabus

Machine Learning

CSCI-B555 - Spring 2015

Class Meets

When: Monday and Wednesday 11:00am-12:30pm
Where: Informatics East, Room 130

Instructor

Martha White
Office: LH 401E
Email: martha@indiana.edu
Web: http://homes.soic.indiana.edu/martha

Office Hours (instructor)

Mondays 2pm-4pm, or by appointment

Course Objective

The course objective is to study the theory and practice of constructing algorithms that learn (functions) and make optimal decisions from data and experience. Machine learning is a field with goals overlapping with other disciplines, in particular, statistics, algorithms, engineering, or optimization theory. It also has wide applications to a number of scientific areas such as finance, life sciences, social sciences, or medicine.

Prerequisites

Graduate student standing or permission of the instructor.

Textbooks

Recommended
Pattern Recognition and Machine Learning - by C. M. Bishop, Springer 2006.

Additional
Machine Learning - by Tom M. Mitchell, McGraw-Hill, 1997
The Elements of Statistical Learning - by T. Hastie, R. Tibshirani, and J. Friedman, 2009

Supplementary material will be provided in class.

Grading

Thought questions: 15%
Final exam: 25%
Homework assignments: 40%
Class (mini) project: 20%

Topics: about 75% of the following topics depending on the year

- mathematical foundations of machine learning
  - random variables and probabilities
  - probability distributions
  - high-dimensional spaces
- overview of machine learning
  - supervised, semi-supervised, unsupervised learning
  - inductive and transductive frameworks
- basics of parameter estimation
  - maximum likelihood and maximum a posteriori
  - Bayesian formulation
- classification algorithms: linear and non-linear algorithms
  - perceptrons
  - logistic regression
  - naive Bayes
  - decision trees
  - neural networks
  - support vector machines
- regression algorithms
  - least squares linear regression
  - neural networks
- kernel methods (taught within classification and regression)
- representation learning and matrix factorization
  - (nonlinear) dimensionality reduction
  - sparse coding
- basics of graphical models
  - Bayesian networks, e.g., hidden Markov model
  - inference and estimation
- ensemble methods
  - bagging
  - boosting
  - random forests
- practical aspects in machine learning
  - data preprocessing
  - overfitting
  - accuracy estimation
  - parameter and model selection
- special topics (if time permits)
  - introduction to PAC learning
  - sample selection bias
  - learning from graph data
  - learning from sequential data

Late Policy and Academic Honesty

All assignments and exams are individual, except when collaboration is explicitly allowed. All the sources used for problem solution must be acknowledged, e.g. web sites, books, research papers, personal communication with people, etc. Academic honesty is taken seriously; for detailed information see Indiana University Code of Student Rights, Responsibilities, and Conduct.