CS478 - Machine Learning, Pattern Recognition and Data Mining. Tentative Syllabus - Spring 2001

M = Mithcel, D = Duda, Hart & Stork, C = Cover & Thomas

- Introduction
  M chap. 1, D chap. 1 (Week 1)
  - What is machine learning?
- Concept learning
  M chap. 2 (Week 1)
- Non-metric methods
  D chap. 8, M chap. 3 (Week 2)
  - Decision trees
  - Strings based methods
  - Grammatical methods
  - Rule-based methods
- Bayesian Learning
  D chap. 2-3, M chap. 6 (Weeks 3-6)
  - Bayesian decision theory
    * Continuous features
    * Discrete features
    * Bayesian belief networks
  - ML and Bayesian Parameter Estimation
    * MAP and ML learners
    * Naive Bayes
    * Bayes optimal classifier
    * Gibbs algorithm
    * Sufficient statistics
    * The EM algorithm
    * Hidden Markov models
- Nonparametric Techniques
  D chap. 4, M chap. 8 (Week 7)
  - Density Estimation
  - Parzen Windows
  - The nearest neighbor algorithm
- Linear discriminant functions
  D chap. 5, M chap. 4 (Weeks 8-9)
  - LD functions and decision surfaces
  - The perceptron function
  - Relaxation and MSE procedures
  - Support vector machines
- Multilayer neural networks
  D chap. 6, M chap. 4 (Weeks 9-10)
  - Feedforward operation
  - Backpropagation algorithm
  - Error surfaces
- Stochastic methods
  D chap. 7, M chap. 9 (Week 10-11)
  - Simulated annealing
  - Genetic algorithms
  - Genetic programming
- Algorithm-independent machine learning
  D chap. 9, M chap. 5 (Week 11)
  - Model selection
  - MDL principle
  - Overfitting
  - Evaluation of learning systems
  - Combining classifiers
- Unsupervised Learning
  D chap. 10 (Weeks 12-13)
  - K-means clustering
  - Unsupervised Bayesian learning
  - Kohonen networks
  - Hierarchical clustering
  - Problems of dimensionality
  - Component analysis
  - Multidimensional scaling
- Computational Learning Theory
  M chap. 7 (Week 14)
  - PAC learning
  - VC dimension
  - Mistake bounds
- Basic concepts in Information Theory
  C (Week 14)