CS478 - Machine Learning
Tentative Syllabus - Spring 2002

M = Mithcel, D = Duda, Hart & Stork

- Introduction
  M chap. 1, D chap. 1 (Week 1)
  - What is machine learning?
- Non-metric methods, inductive inference
  D chap. 8, M chap. 2-3 (Week 1-3)
  - Concept learning
  - Decision trees
  - Strings based methods (maybe)
- Bayesian Learning
  D chap. 2-3, M chap. 6 (Weeks 3-7)
  - Bayesian decision theory
  - ML, MAP and Bayesian Parameter Estimation
  - Sufficient statistics (maybe)
  - Hypothesis evaluation using Bayes Theorem
  - Bayes optimal classifier and Naive Bayes
  - Gibbs algorithm
  - Bayesian belief networks
  - The EM algorithm
  - Hidden Markov models
- Nonparametric Techniques
  D chap. 4, M chap. 8 (Week 7-8)
  - 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 (maybe..)
- Artificial neural networks
  D chap. 6, M chap. 4 (Weeks 9-10)
  - Feedforward operation
  - Backpropagation algorithm
  - Feature mapping
  - Improving performance
- Stochastic methods
  D chap. 7, M chap. 9 (Week 10-11)
  - Genetic algorithms
  - Genetic programming
  - Simulated annealing
- Hypothesis evaluation and model selection
  D chap. 9, M chap. 5 (Week 11-12)
  - Sample error vs. true error
  - Confidence intervals
  - Comparing hypotheses
  - K-fold cross validation
  - The no free lunch theorem
  - MDL principle
- Unsupervised Learning
  D chap. 10 (Weeks 12-14)
  - Mixture densities
  - The maximum likelihood estimates
  - The iterative EM clustering algorithm
  - The K-means clustering algorithm
  - Hierarchical clustering
  - Principal component analysis
  - Multidimensional scaling
  - Kohonen networks