Who's Who

Instructor: Malik Magdon-Ismail, magdon@cs.rpi.edu, Lally Rm 312, x4857.
TA: Brian Osman, osmanb@acm.org, Lally Rm 316, Office Hours: M, 4-5pm.
Class: MR 2-3:20 pm. Office Hours: M 3:30-5pm or by appointment.

Course Description

A basic introduction to the theory, algorithms and applications of automated learning. I will assume that you are familiar with basic mathematics (including elementary probability and matrix theory) and computer programming. The (adaptive) outline is:

Supervised Learning

- *The setup of the learning problem*: Main issues, simple learning models and algorithms, neural networks.
- *Generalization*: How the learning process will perform on unseen data. Model complexity-over fitting tradeoff, bias-variance, regularization, the effect of noise.
- *Learning algorithms*: Computation issues, what to minimize and how to minimize it, error functions and optimization techniques.
- *Learning models*: Alternative structures for hypothesis sets - nearest neighbor, radial basis functions, support vector machines, hidden Markov models, mixtures of experts.
- *Learning aides*: Preprocessing, post-processing, committees, hints.
- *The Bayesian approach to learning and Gaussian Processes.*
- *Learning aides*: Mahalanobis distance, principal component analysis, incorporation of prior knowledge (e.g. invariances) into learning systems.

Reinforcement Learning

- *Basic setup*: Exploration versus exploitation tradeoff.
- *Value functions and Dynamic programming.*
- *Temporal difference learning.*

Unsupervised Learning (If time permits)

- *Density Estimation.*
- *Clustering.*

Policies

Course Grade

Homeworks will be handed out roughly twice every three weeks, and will closely follow lectures. Your grade will be based on your homeworks.

Collaboration Policy

- Discussions are allowed in the homeworks. Nothing written can be taken away from any collaboration. You should write and understand all solutions yourself.
- Computer logistics and debugging help can be obtained from anyone, but you should write all the code and report only results from your own programs.
- Books and notes can be consulted but not copied from.

Late Policy

All assignments are due 1 week after they are handed out. Late assignments are penalized 20% per day (except in institute established illness or emergency). Late assignments should be turned into the TA in charge.

Academic Dishonesty

An automatic grade of F will result in any cases of academic dishonesty.

Course References

1. C. Bishop, *Neural Networks for Pattern Recognition*, Clarendon Press, Oxford.**(Suggested Text)**
2. T. M. Mitchell, *Machine Learning*, McGraw Hill.
3. R. Sutton, A. Barto, *Reinforcement Learning*, MIT Press.
4. V. Vapnik, *Statistical Learning Theory*, Wiley.
5. J. Hertz et al., *Introduction to the Theory of Neural Computation*, Addison-Wesley.
6. R. Duda & P. Hart, *Pattern Classification and Scene Analysis*, John Wiley & Sons.
7. S. Haykin, *Neural Networks – A Comprehensive Foundation*, Macmillan Publishing.
8. N. Nilsson, *The Mathematical Foundations of Learning Machines*, Morgan Kaufmann.
9.B. Ripley, *Pattern Recognition and Neural Networks*, Cambridge Press.

General References

1. J. Pitman, *Probability*, Springer-Verlag.
2. M. Degroot, *Probability and Statistics*, Addison Wesley.
3. S. Lang, *Undergraduate Algebra*, Springer-Verlag.
4. R. Horn and C. Johnson, *Matrix Analysis*, Cambridge Press.