CS 7641 Syllabus

Machine Learning is that area of Artificial Intelligence that is concerned with computational artifacts that modify and improve their performance through experience. The area is concerned with issues both theoretical and practical. This particular class is a part of a series of classes in the Intelligence thread, and as such takes care to present algorithms and approaches in such a way that grounds them in larger systems. We will cover a variety of topics, including: statistical supervised and unsupervised learning methods, randomized search algorithms, Bayesian learning methods, and reinforcement learning. The course also covers theoretical concepts such as inductive bias, the PAC and Mistake-bound learning frameworks, minimum description length principle, and Ockham's Razor. In order to ground these methods the course includes some programming and involvement in a number of projects.

Objectives

There are four primary objectives for the course:

1. To provide a broad survey of approaches and techniques in ML
2. To develop a deeper understanding of several major topics in ML
3. To develop the design and programming skills that will help you to build intelligent, adaptive artifacts
4. To develop the basic skills necessary to pursue research in ML

The last objective is the core one: you should develop enough background that you can pursue any desire you have to learn more about specific techniques in ML, either to pursue ML as a research career, or to apply ML techniques in other research areas in interesting (as opposed to uninteresting) ways.

Topics

1. Supervised Learning
    1. Regression and Classification
    2. Decision Trees
    3. Neural Networks
    4. Instance Based Learning
    5. Ensemble Learning
    6. Kernel Methods and Support Vector Machines
    7. Computational Learning Theory
    8. VC Dimensions
    9. Bayesian Learning and Inference

2. Unsupervised Learning
    1. Randomized Optimization
    2. Clustering and Expectation Maximization
    3. Feature Selection
    4. Feature Transformation
    5. Information Theory

3. Reinforcement Learning
    1. Markov Decision Processes
    2. Reinforcement Learning
    3. Game Theory