8. AI 832: Reinforcement Learning

a. Textbook

- Reinforcement Learning: An Introduction by Sutton and Barto. Also available at http://incompleteideas.net/book/the-book-2nd.html
- Handouts and research articles may also be used by the instructor.

b. Objective

This class will provide a solid introduction to the field of reinforcement learning and students will learn about the core challenges and approaches, including generalization and exploration. Through a combination of lectures, and written and coding assignments, students will become well versed in key ideas and techniques for RL.

c. Pre-Requisite

Linear Algebra, Probability, and Machine Learning

d. Course Outcome

- Define the key features of reinforcement learning that distinguishes it from AI and non-interactive machine learning.
- Given an application problem (e.g. from computer vision, robotics, etc), decide if it should be formulated as a RL problem; if yes be able to define it formally (in terms of the state space, action space, dynamics and reward model), state what algorithm (from class) is best suited for addressing it and justification.
- Implement in code common RL algorithms.
- Describe (list and define) multiple criteria for analyzing RL algorithms and evaluate algorithms on these metrics: e.g. regret, sample complexity, computational complexity, empirical performance, convergence, etc.
- Describe the exploration vs exploitation challenge and compare and contrast at least two approaches for addressing this challenge (in terms of performance, scalability, complexity of implementation, and theoretical guarantees)

e. Course Outline:

- Introduction to Reinforcement Learning
- Markov Processes
- Policy Search and Iteration
- Value Iteration
- Policy Evaluation
- Model Free Learning: Q-Learning and SARSA

- RL with function approximation
- Imitation Learning in Large spaces
- Exploration/Exploitation
- Batch Reinforcement Learning