Advanced Dynamic Programming for Spring 2010
OR 774/IT774
Prereq: OR 674 or permission of instructor
Instructor: Dr. Rajesh Ganesan

This course covers advanced topics on the theory and practice of dynamic programming, i.e. optimal sequential decision making over time in the presence of uncertainties. The course will stress on the mathematical foundations and will introduce the theory, computational aspect, and applications of approximate dynamic programming for stochastic DP problems. The course will use Matlab.

Text: Approximate DP, Warren Powell, Wiley Publishers
Notes prepared from
- Neuro-Dynamic Programming (Optimization and Neural Computation Series, 3) by Dimitri P. Bertsekas, John N. Tsitsiklis
- Markov Decision Processes: Discrete Stochastic Dy. Programming by M. L. Puterman
- Reinforcement Learning by Sutton and Barto

Topics

Review: Dynamic Programming Algorithm
               Deterministic Systems, Shortest path Algorithm
               Stochastic DP
Introduction to ADP Algorithms
Stochastic Approximation
Convergence and optimality
Reinforcement learning
ADP: value function approximation

    Student Evaluation Criteria
    Mid-term: 40%
    Case study discussion 20%
    Final Exam: 40%

Academic Policy:
All academic policies as given in the Honor System and code will be strictly followed. Visit URL
http://www.gmu.edu/catalog/apolicies/#Anchor12

Grades:
Letter grades will be decided as follows:
               97% and above –A$^+$, 94-96%- A, 90-93% -A$^-$, 86-89- B+, 83-85%-B, 80-82%-B-,
               76-79%- C$^+$, 73-75%- C, 70-72%-C$^-$, 66-69%-D$^+$, 63-65%-D, 60-62%-D$^-$, at or
               below 59%-F

Please visit http://classweb.gmu.edu/rganesan to check for announcements, Hw problems, and solutions