

OR 674/SYST 674 Dynamic Programming (Graduate) for Fall 2010
Pre Req: OR 442 or OR542 or Permission of Instructor

Instructor: Rajesh Ganesan, Ph.D.
rganesan@gmu.edu
703-993-1693
ENGR 2217

Course Description:

This is a course 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 intuition, the mathematical foundations being for the most part elementary. It will introduce the theory, applications (finance, engineering, and biology), and computational aspects of dynamic programming for deterministic and stochastic problems. The course will use some Matlab and spread sheets to solve DP problems, however prior knowledge of Matlab or spread sheet use is not needed.

Course Objectives:

At the conclusion of this course the student will have learned the art of formulating recursive equations, and how and why dynamic programming is the only method that can solve many of the large scale optimization problems involving sequential decision making. The student will also have advanced knowledge of solving optimization problems that involve sequential decision making in both deterministic and stochastic environment.

The field of dynamic programming provides alternate approaches to solving many optimization and control problems. Large scale optimization and control has become very important in recent years, which involve complexity and are often model free. Also adaptive systems with sequential decision making powers are sought after for effective real time optimization and control. The above is achievable only via dynamic programming, which is often solved using artificial intelligence (stochastic approximation methods). Many large scale industrial and defense applications such as airline pricing, optimal power flow, helicopter control, and missile control use approximate dynamic programming approaches, which has created a demand for this knowledge.

This is a new course in the SEOR program that has been designed to provide a wealth of knowledge that is directly applicable to the needs of applications that are complex, adaptable, and large scale. The course compliments the fundamentals learnt in OR441/OR541 and OR442/OR542 and introduces a new and powerful alternate to solve many optimization problems that involve sequential decision making.

Text: Eric Denardo, Dynamic Programming: Models and Applications, Dover, May 2003.

Problems from Operations Research by Winston. W. L. (Chapters 18 and 19). This book is the same book that you have used for OR 541 and OR 542.

References Books

- Introduction to Mathematical Programming : Applications and Algorithms Wayne L. Winston, Munirpallam Venkataramanan

- Approximate Dynamic Programming: Warren Powell.
- Dynamic Programming by Richard Bellman
- Dynamic Programming and Optimal Control (Volumes 1 and 2) by Dimitri P. Bertsekas
- Introduction to Stochastic Dynamic Programming by Sheldon M. Ross
- Neuro-Dynamic Programming (Optimization and Neural Computation Series, 3) by Dimitri P. Bertsekas, John N. Tsitsiklis
- Markov Decision Processes: Discrete Stochastic Dynamic Programming by Martin L. Puterman

Deterministic Dynamic Programming:

Week 1	Course introduction, Finite Decision Trees
Week 2	Dynamic Programming Networks and the Principle of Optimality
Week 3	Formulating dynamic programming recursions, Shortest Path Algorithms, Critical Path Method, Resource Allocation (including Investments)
Week 4	Knapsack Problems, Production Control, Capacity Expansion, and Equipment Replacement
Week 5	Infinite Horizon Optimization including Equipment Replacement over an Unbounded Horizon
Week 6	Infinite Decision Trees and Dynamic Programming Networks
Week 7	Midterm

Stochastic Dynamic Programming

Week 8	Stochastic Shortest Path Problems with examples in Inventory Control
Week 9	Markov Decision Processes, value and policy iteration for average cost criteria
Week 10	Markov Decision Processes, value and policy iteration for discounted cost criteria
Week 11	Markov Decision Processes, value and policy iteration for discounted cost criteria
Week 12	MDP with examples in Equipment Replacement and inventory problems
Week 13	Semi-Markov Decision Process
Week 14	Reinforcement learning for MDP (Approximate dynamic programming)
Week 15	Final exam (exam week)

Student Evaluation Criteria

Mid-term:	40%
Project:	20% (Due on Dec 15 th)
Final Exam:	40% (Due on Dec 15 th)

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

Class Website: <http://classweb.gmu.edu/rganesan/>

For hw and midterm grades log in to blackboard and for final grades check patriotweb