Spring 2012

OPIM 7400: Stochastic Dynamic Programming with Applications

Dan Zhang
Leeds School of Business
University of Colorado at Boulder
Phone: (303) 492-2340; Fax: (303) 735-2915
Email: dan.zhang@colorado.edu
Office: Koelbel 452

CLASS MEETINGS

Monday 5:15PM-8:15PM, KOBL 350

OFFICE HOURS

By appointments.

COURSE DESCRIPTION

Stochastic dynamic programming (also known as Markov decision process) is widely used to study problems that involve sequential decision making under uncertainty. This course covers the basic models and solution techniques for problems with finite or infinite horizon. Approximate solution techniques (approximate dynamic programming) for problems involving large state/decision spaces and/or complex dynamics over time will also be discussed. The course is geared toward students who are interested in applying these methods in research. Application domains include, among others, revenue management and pricing, manufacturing, supply chains, service systems, economics, etc.

Prerequisites: An introductory course in Optimization and Probability, or instructor consent.

COURSE MATERIAL

Course material will be distributed by email, on D2L, or in class.

REQUIRED BOOK:

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York.

REFERENCE BOOKS (NOT REQUIRED)

Bertsekas, D. P. 2001. *Dynamic Programming and Optimal Control* , vol.I & II. Athena Scientific.

Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press.

Powell, W. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience.

Ross, S. 1983. *Introduction to stochastic dynamic programming*. Academic Press.

COURSE REQUIREMENTS

You are expected to do the readings before each session. Class time is dedicated to applications of the theoretical material. Each student is expected to lead the discussion of two papers in class. A sign-up sheet will be provided in the first class.

Regular homework assignments are designed to illustrate the principles learned in class. Homework should be done individually and handed in on time. No extensions will be accepted.

GRADING

Class participation/presentations	25%
Homework (2 sets)	20%
Midterm	20%
Final	35%

TOPICS (subject to change depending on student interest and the pace of the class)

1. Introduction (3 hours)

   Readings: Puterman (1994), Chapters 1 (skim), 2.1, 2.2, 3.1, 3.2, 3.3.

2. Finite-Horizon Markov Decision Processes. (6 hours)
   Principle of optimality; backward induction; policy evaluation

   Readings: Puterman (1994), Sections 4.1-4.6
   Lautenbacher, Conrad J., Shaler Stidham. 1999. The underlying Markov decision process in the single-leg airline yield management problem. *Transportation Science* 33 136-146.

3. Infinite-Horizon Discounted Problems (6 hours)
   Optimality criteria; Markov policies; value iteration; policy iteration; linear programming

   Readings: Puterman (1994), Sections 5.1, 5.3, 5.5, 5.6, 6.1, 6.2 (not 6.2.5), 6.3.2, 6.4.1, 6.4.2, 6.9
   Bhandari, A., N. Secomandi. 2011. Revenue management with bargaining. *Operations Research* 59(2) 498-506.
   Patrick, J., M. Puterman, M. Queyranne. 2008. Dynamic Multipriority Patient Scheduling for a Diagnostic Resource. *Operations Research*, Vol. 56, No. 6, pp. 1507–1525.

4. Infinite-Horizon Average Reward Problems (3 hours)

Classification of Markov decision processes; unichain models; value iteration and policy iteration; linear programming

Readings: Puterman (1994), sections 8.1, 8.2, 8.3 (skim), 8.4, 8.5, 8.6, 8.8
Chen, X., D. Simchi-Levi. 2004. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. *Mathematics of Operations Research* 698-723.

5. Structural properties (3 hours)
   Preservation of monotonicity and convexity properties; structural properties of the value function and the optimal solution; monotone policies; closure properties; time preservation; K-convexity in inventory problems

   Readings: Porteus, Sections 7.1, 8.1-8.3
   Smith, J. E., K. F. McCardle. 2002. Structural properties of stochastic dynamic programs. *Operations Research* 50(5) 796-809.

6. Continuous-Time Models (6 hours)
   Continuous-time Markov chain; discounted and average reward models

   Readings: Puterman (1994), Chapter 11
   Porteus (2002), Chapter 14
   Paschalidis, I.C., J.N. Tsitsiklis. 2000. Congestion-dependent pricing of network services. *IEEE/ACM Transactions on Networking* 8(2) 171-184.

7. Introduction to Approximate Dynamic Programming (3 hours)
   Curse of dimensionality; simulation-based methods; mathematical programming based methods

   Readings: Powell (2007), Chapters 1 (skim), 4 (skim)
   Bertsekas (2011). Chapter 6, Approximate Dynamic Programming, *Dynamic Programming and Optimal Control*, 3rd Edition, Volume II. Available online at http://web.mit.edu/dimitrib/www/dpchapter.pdf.

8. Topics in Approximate Dynamic Programming (6 hours)
   Value function approximation; information relaxation; Lagrangian relaxation

   Readings: Adelman, D. 2007. Dynamic bid-prices in revenue management. *Operations Research* 55(4) 647-661.
   Adelman, D., A. J. Mersereau. 2008. Relaxations of weakly coupled stochastic dynamic programs. *Operations Research* 56(3) 712-727.
   Brown, D.B., J.E. Smith, P. Sun. 2010. Information relaxations and duality in stochastic dynamic programs. *Operations research* 58(4) 785-801.

Longstaff, F.A., E.S. Schwartz. 2001. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies* 14(1) 113.

TENTATIVE SCHEDULE

Date	Topic	Comment
1/23/12	Introduction	Sign up for paper presentation
1/30/12	Finite-horizon problems	
2/6/12	Finite-horizon problems	
2/13/12	Infinite-horizon discounted problems	
2/20/12	Infinite-horizon discounted problems	
2/27/12	Infinite-horizon average reward problems	First homework assigned
3/5/12	Structural properties	
3/12/12	Continuous-time models	First homework due
3/19/12	Midterm	
3/26/12	Spring break	No class
4/2/12	Continuous-time models	
4/9/12	Topics in approximate dynamic programming	Second homework assigned
4/16/12	Topics in approximate dynamic programming	
4/23/12	Topics in approximate dynamic programming	Second homework due
4/30/12	Final	

Academic Honesty: All students of the University of Colorado at Boulder are responsible for knowing and adhering to the academic integrity policy of this institution. Violations of this policy may include: cheating, plagiarism, aid of academic dishonesty, fabrication, lying, bribery, and threatening behavior. All incidents of academic misconduct shall be reported to the Honor Code Council (honor@colorado.edu; 303-725-2273). Students who are found to be in violation of the academic integrity policy will be subject to both academic sanctions from the faculty member and non-academic sanctions (including but not limited to university probation, suspension, or expulsion). Other information on the Honor Code can be found at http://www.colorado.edu/policies/honor.html and at http://www.colorado.edu/academics/honorcode/

Disabilities: If you qualify for accommodations because of a disability, please submit to me a letter from Disability Services in a timely manner so that your needs may be addressed. Disability Services determines accommodations based on documented disabilities. Contact: 303-492-8671, Willard 322, and http://www.colorado.edu/disabilityservices/

Religious Observances: Campus policy regarding religious observances requires that faculty make every effort to reasonably and fairly deal with all students who, because of religious obligations, have conflicts with scheduled exams, assignments or required attendance. In this class, {{insert your procedures here}} See full details at http://www.colorado.edu/policies/fac_relig.html

Discrimination and Harassment: The University of Colorado at Boulder policy on Discrimination and Harassment, the University of Colorado policy on Sexual Harassment and the University of Colorado policy on Amorous Relationships apply to all students, staff and faculty. Any student, staff or faculty member who believes s/he has been the subject of discrimination or harassment based upon race, color, national origin, sex, age, disability, religion, sexual orientation, or veteran status should contact the Office of Discrimination and Harassment (ODH) at 303-492-2127 or the Office of Judicial Affairs at 303-492-5550. Information about the ODH, the above referenced policies and the campus resources available to assist individuals regarding discrimination or harassment can be obtained at http://www.colorado.edu/odh

Sexual Harassment: The University of Colorado Policy on Sexual Harassment applies to all students, staff and faculty. Sexual harassment is unwelcome sexual attention. It can involve intimidation, threats, coercion, or promises or create an environment that is hostile or offensive. Harassment may occur between members of the same or opposite gender and between any combination of members in the campus community: students, faculty, staff, and administrators. Harassment can occur anywhere on campus, including the classroom, the workplace, or a residence hall. Any student, staff or faculty member who believes s/he has been sexually harassed should contact the Office of Sexual Harassment (OSH) at 303-492-2127 or the Office of Judicial Affairs at 303-492-5550. Information about the OSH and the campus resources available to assist individuals who believe they have been sexually harassed can be obtained at: http://www.colorado.edu/sexualharassment/