OIDD 934 – Dynamic Programming
Q3, Spring 2018, March 13 –April 24, 2018

Instructor: Dr. Maria T. Rieders
Adjunct Full Professor
Operations, Decisions and Information Department
The Wharton School
University of Pennsylvania
Office: 517 Jon M. Huntsman Hall
Phone: (215) 898-0535
e-mail: rieders@wharton.upenn.edu

Lectures: Tuesday, Thursday 10:30-11:50, JMHH G88
Office Hours: Tuesday 1:30-2:30; Friday 10:30-noon
Website: https://canvas.upenn.edu/

Evaluation: Weekly/biweekly homework assignments (50%);
Student project and participation (50%);

Prerequisites:

- Stochastic processes, in particular Markov chains in discrete and continuous time (as in OIDD 930)
- Basics of real analysis; linear algebra
- Some exposure to linear programming, including the construction of basic solutions and duals
- Graduate standing; target audience: PhD students

Course Materials:
We will cover material that draws mostly from the following texts:

- Martin Puterman, *Markov Decision Processes – Discrete Stochastic Dynamic Programming,* Wiley, 1994;
- Dimitri Bertsekas, *Dynamic Programming and Optimal Control*, volumes 1 and 2, 3rd edition, Athena Scientific, 2007;
- Warren Powell, *Approximate Dynamic Programming – Solving the Curses of Dimensionality,* Wiley, 2007

The flavors of these texts differ. Puterman carefully constructs the mathematical foundation for Markov decision processes. His focus is on theory such as conditions for the existence of solutions and convergence properties of computational procedures. Bertsekas uses an engineering type approach and offers a wide variety of dynamic programming formulations in applied settings. Powell gives a concise, but very readable introduction to the basic theory of dynamic programming and then focuses on computational methods suitable to dealing with very large state spaces. Approximations and aggregation of states play a large role. You are not required to purchase a particular text but will benefit from choosing one of these texts as your primary source. I will supplement your choice by making these texts available at the reserve desk at Lippincott library and by posting readings on Canvas.

Communication:
This class utilizes a Canvas website where relevant information for this course will be posted.

Course Description:
This course presents a mathematical modeling framework for sequential decision processes. A decision maker is to choose a sequence of actions towards an objective goal such as to maximize expected reward (or minimize expected cost). At each stage, his choice of action as well as events outside of his control determines the future development of the system. The difficulty that the decision maker faces consists of considering both the present reward and anticipated future rewards. Dynamic programming captures this tradeoff.
The course goal is to provide a brief but fairly rigorous introduction to the formulation and solution of dynamic programs. Its focus is primarily methodological. We will mostly address discrete-time problems with discrete state spaces, over finite or infinite time horizon; as time allows we will consider some continuous time models as well. We will cover computational methods and some approximation techniques that allow for calculation of large systems. Of particular interest are properties that result in structured policies as this will increase our insight in a problem's solution. Applications are presented throughout the course, including inventory policies, production control, financial decisions, and scheduling.

Course Topics:
A tentative course outline follows. Actual contents may vary depending on interests by students and/or the instructor.

Session	Date	Topic
1	Tue 3/13	Introduction – Elements of DP
2	Thu 3/15	Finite Horizon MDP
3	Tue 3/20	Finite Horizon MDP
4	Thu 3/22	Structured Policies for Finite Horizon Problems
5	Tue 3/27	Structured Policies for Finite Horizon Problems
6	Thu 3/29	Infinite Horizon Problems, Discounted MDPs
7	Tue 4/3	Infinite Horizon Problems, Discounted MDPs
8	Thu 4/5	Average Reward MDPs
9	Tue 4/10	Computational Methods
10	Thu 4/12	Computational Methods
11	Tue 4/17	Continuous Time MDPs
12	Thu 4/19	Wrap-up; Student Presentations
13	Tue 4/24	Student Presentations

Student Participation:
This is a Ph.D. level course. Students are expected to thoroughly review lecture material before coming to class, and actively participate through discussions, questions, and answers.

Homework:
Assignments will be handed out on a weekly/biweekly basis. Some assignments may require programming. Unless otherwise stated, homework is to be done individually. You may discuss the problems with each other; yet the work that you submit must be your own. If you do obtain a solution from some other source (e.g. from a peer, another text, or an online site), it is expected that you cite the source.

Project/Presentation:
Each student is to design a project and submit a report by the end of the term. The topic is to be approved by the instructor and may consist of a survey of one or more research papers, an attempt to formulate an interesting question as a dynamic programming problem, or the computational implementation and experimentation of a computational algorithm. Each student will give a short presentation during the last class week of classes.