MASSACHUSETTS INSTITUTE OF TECHNOLOGY
Department of Electrical Engineering and Computer Science

6.231: Dynamic Programming—Fall 2006

Course Information and Outline

Course Information

Instructor Prof. Munther A. Dahleh (dahleh@mit.edu), Room 32-D734, x3-3892.

TAs
- Sleiman Itani (`ssolomon@MIT.EDU`)and Ilan Lobel (`lobel@mit.edu`) will be your TAs, and will conduct the recitation section.

Secretary Ms. Fifa Monserrate (`fifa@mit.edu`), Room 32-D733, x3-2184.

Lectures M,W 11-12:30, Room 36-112.

Recitation F 11-12, Seminar Room in the 7th Floor of 32D.

Homework Generally handed out every Wednesday, and due in class a week later (except as noted on schedule), at which time solutions will be handed out.

Tests There will be one exam conducted during the week of Nov 27th.

Grading The course grade will depend on: (a) your involvement in the subject (30%), as evidenced mainly by your homework, but also by your interaction with the TAs and instructor; (b) your performance on the the exam (40%), and a final project (30%).

Homework Policy

The homeworks are *not* intended as tests, but as vehicles for learning, complemented by the homework solutions that we hand out, and by any discussions that you have about the problems. *Moderate* collaboration on homework with your *classmates* is permitted. Discussions with the TAs and instructor are encouraged. There is no harm in seeking minor assistance from others who are knowledgeable but not involved in the class, although we would much prefer that your discussions be with those in the class.

We expect each of you to put in enough time *alone* to understand the specific difficulties and issues raised by each homework problem. We also expect that you will *independently* write up the actual solutions that you turn in, and not give us direct copies of a classmate's solutions! You should note on your solutions the names of those you have collaborated with or obtained help from.

We will feel free to use problems from previous terms. For obvious reasons, *any* use of written solutions from previous terms is *strongly prohibited* — you will gain very little in

the short run, and you will lose significantly on the opportunity to genuinely master the material.

Computer Use

It is important that you give us an e-mail address for yourself, and important that you check it frequently (daily?), as there will quite likely be administrative and other messages sent out from time to time by the 6.241 teaching staff.

We expect all of you to have (or to immediately arrange to get) an account on Athena. There is a 6.231 locker on Athena. To access it, type `add 6.231`, then `cd /mit/6.231/Public`. Course material will be posted on the course web site: web.mit.edu/6.231/www/.

Homework this term may make some use of Matlab, a package for numerical linear algebra that is available on Athena. You should start familiarizing yourself with Matlab right away, if you have not worked with it before. You can get an introductory overview by typing `add matlab` and hitting the return key at the Athena prompt, then `matlab &` and return. Once in Matlab, you can type `demo` and return, then use the menu. Type `quit` and return to get out of Matlab. You can also find an interactive tutorial and useful links to other tutorials at http://web.mit.edu/matlab/www/home.html.

Notes and Texts

Required Dimitri Bertsekas. *Dynamic Programming and Optimal Control: Vol 1.*
Recommended Dimitri Bertsekas. *Dynamic Programming and Optimal Control: Vol 2.*

Team Projects:

Teams comprise 2-3 people. The projects should take no more than 2 weeks to complete. The outcome will be reported in a 30 min presentations made by each group (see schedule).

Groups as well as their 1-page proposals should all be decided by Nov 10. I encourage all of you to talk to the TAs and myself before then so that we can agree on the proposal.

Below is a suggested list of topics.

1. Approximate Dynamic Programming–LP methods
2. Simulation methods and Q-Learning
3. Application to systems with large number of states
4. Incomplete information with continuous states
5. Continuous-time problems (e.g., controlled Markov Processes)
6. Queuing applications
7. Policy Search Algorithms
8. Stochastic Games

9. TD-Learning Algorithms
10. Actor-Critic Algorithms
11. Decentralized Policies and Control
12. Bandit problems
13. Applications in Combinatorial Optimization

Tentative Schedule of Topics

Date,Day		Topic	Ch.
Sept. 6, W	L1:	Introduction to Dynamic Programming, Examples, Problem Formulation	1.1, 1.2
		PS1 handed out.	
8, F	R1:	probability review, Markov Chains	
11, M	L2:	The Dynamic Programming Algorithm	1.3, 1.4
13, W	L3:	Deterministic Systems and the Shortest Path Problem	2.1,2.2
		PS1 turned in, PS2 handed out.	
15, F	R2:	derivation of the DP algorithm, Label correcting algorithm	
18, M	L4:	Linear Quadratic Problems, Inventory problem	4.1, 4.2
20, W	L5:	Stopping Problems and scheduling problems	4.4,4.5
		PS2 turned in; PS3 handed out.	
22, F	R3:	Jump parameter linear quadratic	
25, M		Student holiday	
27, W	L6:	Rollout policies and Model Predictive control. Introduction to Infinite horizon problems	5.1
		PS3 turned in, PS4 handed out.	
29, F	R4:	minimax problems, e.g., LQ	
Oct. 2, M	L7:	Infinite Horizon Problems: Discounted cost; finite and infinite states	VII, 1.1, 1.2
4, W	L8:	Finite states systems. Value iteration and Policy iteration. LP Formulation	VII, 1.3, 1.5
		PS4 turned in, PS5 handed out.	
6, F	R5:	Examples of discounted Infinite horizon formulations (Add Date)	
9, M		Columbus Day.	
11, W	L9:	Stochastic shortest path problem	VI 7.1, VII 2.1, 2.2
		PS5 turned in, PS6 handed out.	
13, F	R6:	Scheduling and Multiarmed Bandit problem	
16, M	L10:	Stochastic Games	Supplement
18, W	L11:	unbounded cost problems	VII 3.1
		PS6 turned in, PS7 handed out	
20, F	R7	Examples of SSP	
23, M		Class Holiday	
25, W	L12:	unbounded cost problems: infinite-Horizon LQ	VII 3.2
		PS7 turned in, PS8 handed out	
27, F	R8:	Counter examples	

Date,Day		Topic	Ch.
30, M	L13:	Finish Unbounded cost problems. Average cost problem, value and policy iterations, LP formulation	7.4
Nov. 1, W	L14:	Incomplete Information: New formulation	5.1
		PS8 turned in, PS9 handed out.	
3, F	R9:	Average cost examples	
6, M	L15:	Examples: LQ with output feedback	5.2
8 ,W	L16:	Finite state problems and Sufficient statistics	5.4
		PS9 turned in. All proposals are due on Fri	
10, F		Veterans day	
13, M	L17:	Finite dimensional Examples. Discuss Test	5.4
15, W	L18:	Continuous-time problems: Decisions in Markov processes	VII 5.1, 5.2
		Test: Available on Wed and Due on Fri	
17, F	R10:	Example with POMDP	
20, M	L19:	Semi-Markov Problems	VII 5.3
22, W	L20:	No Class	
		PS10 handed in	
		(Drop Date)	
24, F		Thanksgiving.	
27, M	L21:	Approximation and Simulation methods	VII 2.3
29, W	L22:	Q-Learning	VII 2.3
Dec.1, F	R11:	summary of stochastic convergence results	
4, M	L23:	Projects	
6, W	L24	Projects	
8, F	R12:	Projects	
11, M		Projects	
13, W		No Lecture	