

Syllabus for Dynamic Programming Spring 2017

Course Instructor: Diego Klabjan

Course Time and Location: MW 11:00-12:15 pm; Tech M228

Best Way to Contact in Order of Preference:

skype dklabjan (do not add me as a contact)

hangout dklabjan

sms 847 691 1148

email d-klabjan@northwestern.edu

Response time directly correlated with the order of preference.

Textbook: Warren B. Powell; Approximate Dynamic Programming: Solving the Curse of Dimensionality; John Wiley & Sons, 2007. ISBN 978-0-470-17155-4

Goal: This course will cover reinforcement learning aka dynamic programming, which is a modeling principle capturing dynamic environments and stochastic nature of events. The main goal is to learn dynamic programming and how to apply it to a variety of problems. The course will cover both theoretical and computational aspects.

Tentative list of topics:

1. Introduction to dynamic programming (Chapters 1 and 2)
2. Value and policy iterations (Chapter 3)
3. Stochastic gradient algorithm (Chapter 6)
4. Q-learning and temporal differences (Chapter 8)
5. Value function approximation and Monte-Carlo sampling (Chapter 4)
6. Linear and dynamic programming (time permitting, not in the text book)

Grading: There will be mandatory individual homework assignments. In addition, each student will have two options: either performing a quarter long project, or having two take home exams. They will be weighed as follows:

- Homework assignments: 30%
- Midterm and final exam: 35% each (if you choose the homework/exam route)
- Project: 70% (if you choose the homework/project route).