Computer Science 6804: Learning and Sequential Decision-Making

Spring 2013

1 Course Description

1.1 Instructor

Instructor: Sanmay Das
Office: KWII 2226 (for now)
Phone: 540-231-0908
Office hours: Wednesdays from 2:30-4:00 PM (McBryde 122), and by appointment (KWII 2226).

1.2 Overview

This course will study, from a computer science perspective, the problems faced by agents who are situated in dynamic, changing environments. The central question is how agents should reason in order to make the best decisions in these environments. In order to do so, they must be able to learn from experience and perform complex decision-making tasks under uncertainty. Topics to be covered include, but are not limited to, Markov decision processes, partial observability, reinforcement learning, bandit problems, sequential search, and reasoning and learning in games.

1.3 Format

The first few weeks will consist of lectures that go over necessary background and should help to get everyone on the same page. After that, the course will largely follow a seminar format. Each student will present one paper from the reading list. All students will be expected to complete the assigned readings before class and come to class ready to engage in discussing the readings. Most of the reading list will be drawn from recent literature (such as the ICML, NIPS, AAAI, IJCAI, AAMAS and EC conferences).

1.4 Prerequisites

Undergraduate class in AI (CS 4804 or equivalent) or machine learning or data analytics. Graduate work in algorithms/machine learning/data mining strongly recommended. Familiarity with probability and ability to program (ideally in a language that is mathematically flexible and well-suited for rapid prototyping, e.g. Matlab, R, or Python) are absolutely essential.
Note: If any student needs special accommodations because of a disability, please contact the instructor during the first week of classes.

2 Policies

2.1 Website and Communication

We will make extensive use of electronic communication and the course website. You are responsible for checking the course website regularly for announcements and course materials, as well as your e-mail for communications related to the class. The course website is hosted on Piazza and located at:
`https://piazza.com/vt/spring2013/cs6804/home`

2.2 Course meetings

The course will meet from 4:00PM-5:15PM on Mondays and Wednesdays in McBryde 210. You are responsible for all material covered and announcements made in class. Participation will count towards your grade.

2.3 Paper Presentations, Summaries, and Discussions

Each student is required to choose one paper from the reading list and present that paper in class. At least **one week** in advance of your presentation, you must post a summary of the paper to Piazza as a new discussion. In the course of the next week, before the scheduled presentation, every student in the class must make one comment or ask one question on that thread. This could be a question you have upon reading the paper, something you liked about the paper, an answer to another student, a clarification of something in the original summary, or anything substantive along these lines. You are also expected to participate actively in discussions during class.

2.4 Project

The major deliverable for this class is a research paper that summarizes the results of a semester-long project to be done in groups of 2 or 3 people. You will be required to turn in a proposal for this project before spring break. In the last few class meetings, your group must give a 15 minute presentation on the project.

2.5 Assessment and Course Grade

Your overall course score will be determined using the following weights:

1. Project, including proposal and presentation: 50 %
2. Paper summary and presentation: 25 %
3. Participation in discussions on Piazza and in class: 25 %

There are no exams. Students will be graded on a curve. There are no fixed grade thresholds.

2.6 Honor Code, Collaboration, and Academic Integrity

The Virginia Tech honor code applies to all work you submit in this class. Please cite any resources you use in any work you do in this class. You are welcome to discuss and collaborate with other students, but please be sure to acknowledge the collaboration appropriately.

3 Textbook and Syllabus

3.1 Textbook

There is no textbook for this class. There is a reading list of papers from the literature, and the instructor will provide lecture notes and other reading resources as necessary.

3.2 Syllabus

The field of learning and sequential decision-making is vast. We will cover some necessary prerequisites:

1. Markov Decision Processes (both fully and partially observable).
2. Bandit problems and the exploration-exploitation tradeoff.
3. Reinforcement learning.
4. Optimal stopping.
5. Game theory.
6. Auctions and Mechanism Design.

After that, the topics will depend on the readings chosen by students to present in class. In addition to theoretical papers on the topics above, some examples of the applications these readings could cover include (but are not limited to):

1. AI for difficult games, like backgammon and poker.
2. Sponsored search auctions: design and bidding strategies.
3. Trading agent design.
4. Market modeling and design.
5. Online advertising.
6. Energy and smart grids.