CSC2542
Topics in KR&R:
Algorithms for Sequential Decision Making

Sheila McIlraith and Rick Valenzano
Fall, 2016

UNIVERSITY OF TORONTO

General Information

URL: http://www.cdf.toronto.edu/~csc2542h/fall/
Lectures: Thursday 2:00 – 4:00 PM, OI4410 (252 Bloor St West)
Tutorials : Thursday 4:00 – 5:00 PM, OI4410
Instructor: Sheila McIlraith and Richard Valenzano
Email: csc2542profs@cs.toronto.edu
goes to sheila@cs.toronto.edu rvalenzano@cs.toronto.edu
Office: Pratt 398 (6 King's College Road, 3rd floor)
Office Hours: By appointment and during tutorial house

Bulletin Board: We have a **Piazza** instance. Please register. This is largely for the class to use to communicate w/ each other about assignments and the project.

Announcements: On the course Web page. It's your responsibility to watch this.

UNIVERSITY OF TORONTO

Course Description

- CSC2542 is a seminar-style topics course that explores recent advances in knowledge representation and automated reasoning. In the fall of 2016, the topic being covered is "**Algorithms for Sequential Decision Making**." Sequential decision making is the task of deciding what to do in the context of an extended interaction with the environment. It is a core competency of most intelligent agents.
- Central to sequential decision making is the notion of making a good decision now with respect to both immediate and longer term consequences. A driver deciding on a particular route to work, a robot deciding what to say to a person it is conversing with, a bank electing to embark on a major advertising campaign, or a computer program deciding on a move in a game of Go are all examples of sequential decision making.

UNIVERSITY OF TORONTO

Course Description (cont.)

- This course will examine select principles and algorithmic techniques that are exploited for sequential decision making. The course will not be exhaustive but rather will focus on two themes: advances in techniques for **search** including suboptimal search algorithms, real-time search, and Monte Carlo search; and techniques for **sequential decision making under uncertainty**, including Markov Decision Processes (MDPs) and Reinforcement Learning (RL).
- The course will draw predominantly on research readings. The format of the course will be a mix of class lectures, seminars, videos, assigned readings, and student paper presentations. A background assignment and a course project will be major components of a student's course mark.

** *This should be a fun and interesting course!*

UNIVERSITY OF TORONTO

Course Description (cont.)

Format: class lectures, research paper readings and presentations.
Prerequisites:
introductory AI course (CSC384 or comparable)
knowledge of logic.
Breadth Area: Area 1

Each week we will have a lecture and/or selected readings.
Watch the course Web page. The first month will be lectures.

Reference Material

There is no one good textbook for this course. We'll point you to readings or other resources that we think will be interesting/helpful.

- **Artificial Intelligence: A Modern Approach** (3rd edition)
  Stuart Russell and Peter Norvig
  (at least) on hold in the library

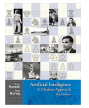

- **Reinforcement Learning: An Introduction**
  Richard S. Sutton and Andrew G. Barto
  2nd Edition (not the 1998 edition)
  (available online)
  https://webdocs.cs.ualberta.ca/~sutton/book/ebook/the-book.html

Reference Material (cont.)

There is no one good textbook for this course. We'll point you to readings or other resources that we think will be interesting/helpful.

- **Heuristic Search: Theory and Applications**
  Stefan Edelkamp
  (electronic copy available in library)
  http://search.library.utoronto.ca/detils?8124175

- **Automated Planning: Theory and Practice**
  Authors: Ghallab, Nau, Traverso
  (online copy available free through science direct)
  http://www.sciencedirect.com/science/book/9781558608566

Course Work

- Each student will **present one paper** or **be involved in the active critiquing of two papers**
- Each week, students will be required to write a **short (1-2 page) summary of assigned readings** for that week, except in the case where they are presenting one of the papers.
- There will be a **warm-up assignment** to get your hands dirty.
- Students will complete a **course project** to be due at the end of the exam period. The course project is to be completed individually. We will have a set of possible projects and will discuss potential topics with students individually. If you have an idea for something you' d like to work on, let' s discuss it! (A project of your own conception is ideal!)

Grading

The grading breakdown is as follows:

- Written paper review & class participation: 10%
- Class paper presentations: 15%
- Warm-up assignment: 20%
- Course project: 55%

There is no exam.

UNIVERSITY OF TORONTO

Caveat: Details may change

The class is larger than we anticipated. Depending on how many students stay enrolled in the course we may modify expectations slightly to accommodate finite resources. The grade distribution will stay the same. Changes that may occur: more than 1 person may be assigned to a paper presentation, there will be an option for projects to be performed in groups.

UNIVERSITY OF TORONTO

Research Paper Summaries (10%)

- Once we start reading research papers, each week students will be required to **hand in a 1-2 page written summaries** of the assigned readings. Summaries are not required by students on weeks they are presenting a paper.
- Your goal in the summaries is to summarize the contents of the paper, and comment on the contribution to the field and its significance. You may write a separate summary of each reading on a given week, or one summary that discusses all of the assigned readings together.

In addition to these summaries, a portion of your mark will be assessed with respect to your participation in class.

UNIVERSITY OF TORONTO

Subject to minor change to accommodate class size

Presentations (15%)

- Students taking the course for credit must give one **class presentation or facilitate discussion** of an assigned reading being presented by a classmate.
- Presentation and discussion of each assigned reading will together take approx. 45 minutes. One student will present the results of the paper. This should take approximately 20-30 minutes in length and should help stimulate discussion. The presenter should provide an overview of the paper, identify the important contributions of the paper and situate the paper within a broader research context. The “facilitator”/”discussant” has several roles. (We will elaborate on this in class. They should situate the paper in the context of previous and subsequent work, and should what the contribution of the paper was and why it was significant at the time. Where relevant, they should discuss the practical significance of the paper in terms of the realization of applications. Finally they should prompt discussion by asking questions of the presenter and audience. This can be during or after the presentation.

UNIVERSITY OF TORONTO

Subject to minor change to accommodate class size

Presentations (15%) (cont.)

- Presenting/facilitating students must **make an appointment to meet with Rick and/or Sheila (several days) prior to their presentation** to go over the material they plan to present. Students should have a substantial draft of the presentation ready to show at that time.
- Students presentations will be posted on the course Web page. Presenting students also have the option of linking any relevant supplementary material.

Student paper presentations will likely start later in October.

Subject to minor change to accommodate class size

Warm-up Assignment (20%)

- There will be a warm-up assignment to get your hands dirty. It will be on the topic of search.

The assignment will be handed out in late September/early October and will be due in mid-late October.

UNIVERSITY OF TORONTO

Class Project (55%)

The course project must be on the general topic of automated planning and reasoning about action. A set of potential topics will be provided, but we encourage students to choose their own topic and to use this as a vehicle to jumpstart a new research project or to investigate a new aspect of ongoing research. **You may work in groups of 1, 2 or 3.**

2-page **project proposal** due ***in late October***. Start thinking about your project early. **Come and talk to us soon and before submitting your proposal!** The proposal must comprise:

- a careful description of the problem your project will address;
- a set of approx 2-4 research papers from which the projects will be drawn;
- a description of the approach you will take to addressing the project;
- a description of how you will evaluate the success of the project;
- a rough schedule for when you'll accomplish the work

...

UNIVERSITY OF TORONTO

Subject to minor change to accommodate class size

Class Project (55%) (cont.)

Evaluation of the project (55 marks) will be as follows:

- (5 marks) Your project proposal.
- (10 marks) Your project presentation. Your presentation will be given in a class towards the end of term. As such, your presentation may have to be given before your project is completed.
- (40 marks) For the overall quality of your project, based in part on its level of difficulty, on the insights you exposed, and any novel ideas of your own that you are able to explore, and your final analysis of your project. A major proportion of this mark will depend on the students' presentation of their final results. This should usually be in the form of a formal written paper, perhaps with a well-structured web site to show results, if relevant.

...

UNIVERSITY OF TORONTO

Subject to minor change to accommodate class size

Class Project (55%) (cont 2)

- (40% cont.) For a "typical assignment": a major component of the report will be a review and analysis of the related literature, along with your assessment of the effectiveness and relative merits of each approach. This will focus mainly on the 2-4 papers you chose, but will also likely require several further sources in order to provide sufficient groundwork. The written report and/or website should will also include a detailed description of any algorithms you implemented. This should include problems you faced, the mathematical details of what was implemented, and an assessment of any empirical results.

Due Date: TBD but probably around the last day of exams.

Extra Incentive: ICAPS/IJCAI Workshop Deadline – early in 2017?

Subject to minor change to accommodate class size

The Weeks Ahead

Main Course Topics

- Search
- Sequential Decision Making Under Uncertainty

Rough Schedule

Week 1-6: Background lectures on search and reinforcement learning

Weeks 7-12: Paper readings and presentations

Getting to know each other

- What is your **home department**?
  - (CS, MIE, I-school, UTIAS, EE, …)
- What is your **primary area of research** right now
  - (undecided, AI, DB, SE, networks, theory, ...)
- What **preparation** do you have for the course?
  - Previous AI/ML courses/experience?
- What interests you about the course?
- If you're interested in a **particular aspect** of the coures, what is it?
  - (e.g., search, planning, reinforcement learning, MDPs,)
- If you're interest in applying sequential decision making techniques to a particular application, what is it?
  - (e.g., robots, software agents, advertising, ..._

Logistics Poll

- Class time OK?
- Tutorial time OK?
- Office hours?

Fun with search and RL!

Kiva Robots - https://www.youtube.com/watch?v=3eQAFVetNGI

AlphaGo - https://youtu.be/g-dKXOlsf98?t=145

Arcade Learning Environment (ALE)
https://www.youtube.com/watch?v=V1eYniJ0Rnk - Breakout
https://www.youtube.com/watch?v=5WXVJ1A0k6Q - Seaquest
https://www.youtube.com/watch?v=P-603qPMkSg - Width-based planners

Reinforcement Learning for Soccer Goalie
https://www.youtube.com/watch?v=CIF2SBVY-J0

Robot Path Planning and Arm Manipulation
https://www.youtube.com/watch?v=Sv-J37zcLU4