CSE 150. Intro to AI

Probabilistic Reasoning and Decision-Making




Welcome to CSE 150!

“I've always considered the most boring 20 minutes of the semester the time I spend reading the syllabus on the first day of class.

Students come in, potentially excited about getting started, only to end up listening to me read aloud.

I imagine them paraphrasing in their heads one of my favorite Woody Allen lines: Thanks, but I've been doing my own reading since about the first grade.”

http://chronicle.com/weekly/v53/i02/02c00201.htm

Out of date

CSE150 - Introduction to Artificial Intelligence: Search and Reasoning

Units: 4

Search algorithms including BFS, DFS, iterative deepening and A*, radomized search algorithms including Walksat, syntas and semantics of first-order logic (FOL), knowledge representation in FOL including reasoning, basic reasoning with probabilities, basic Bayesian learning.

Course Objectives:

Introductory course in artificial intelligence programming, concentrating on the use of the LISP and Prolog languages. Familiarity with computer programming, especially the ideas of procedural abstraction, recursion, list and tree data structures, and elementary tree traversal algorithms is presupposed.

Format:

3 hours of lecture per week, 1 hour of discussion section, and 8 hours of outside preparation.

Web page

CSE 150 - Summer 2016
Introduction to Artificial Intelligence:
Probabilistic Reasoning and Decision Making

Prof. Lawrence Saul



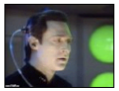


Administrivia	Syllabus	Piazza	GradeSource	CAPEs

Subject

This course will introduce students to the probabilistic and statistical models at the heart of modern artificial intelligence. Specific topics to be covered include: probabilistic methods for reasoning and decision-making under uncertainty; inference and learning in Bayesian networks; prediction and planning in Markov decision processes; applications to intelligent systems, speech and natural language processing, information retrieval, and robotics.

Prerequisites

This course is aimed very broadly at undergraduates in mathematics, science, and engineering. Prerequisites are elementary probability, linear algebra, and calculus, as well as basic programming ability in some high-level language such as C, Java, Matlab, R, or Python. (Programming assignments are completed in the language of the student's choice.) Students of all backgrounds are welcome.

Texts

The course will not closely follow a particular text. The following texts, though not required, may be useful as general references:

- K. Korb and A. Nicholson, Bayesian Artificial Intelligence.
- S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach.
- R. Sutton and A. Barto, Reinforcement Learning: An Introduction.

Instructors

- **Lecturer:** Lawrence Saul (saul@cs.ucsd.edu)
- **Teaching assistants:**
  Zhen Zhai (zzhai@ucsd.edu)
  Sahil Aggrawal (saa034@eng.ucsd.edu)

Meetings

- Lectures: Tue/Thu 11:00a-1:50p, WLH 2005.
- Office hour: Tue/Thu 2:00-3:00p, EBU3B-3214.
- Discussion sections:
  - Mon TBA (Sahil)
  - Wed TBA (Zhen)
- Tutoring hours in CSE basement:
  - Thu TBA (Zhen)
  - Fri TBA (Sahil)
- Final exam: Sat July 30, 11:30a-2:30p

Grading

- homework (54%)
- final exam (46%)

Syllabus

Tue June 28	Administrivia and course overview. Modeling of uncertainty, review of probability.	HW 1
Thu June 30	Examples of probabilistic reasoning. Belief networks: from probabilities to graphs.	HW 2
Tue July 05	Conditional independence, d-separation. Inference in polytrees and loopy networks.	HW 3
Thu July 07	Learning, maximum likelihood estimation. Naive Bayes and Markov models.	HW 4
Tue July 12	Latent variable models, EM algorithm. Examples: noisy-OR, word clustering.	HW 5
Thu July 14	Hidden Markov models, automatic speech recognition. Inference and learning in HMMs.	HW 6
Tue July 19	Introduction to reinforcement learning. Markov decision processes.	HW 7
Thu July 21	Policy evaluation, greedy policies, policy improvement.	HW 8
Tue July 26	Policy and value iteration. Temporal difference learning, Q-learning.	HW 9
Thu July 28	Course wrap-up, odds and ends, review.	
Sat July 30	**Final exam**	

http://cseweb.ucsd.edu/classes/wi15/cse150-a/

Who uses probabilistic methods in AI?

- **Google – pattern matching**
- **Amazon - customer profiling**
- **IBM – Watson & Jeopardy Q/A**
- **Microsoft - multimedia OS**
- **Game makers - enemy AI**
- **FBI - forensic speaker ID**

Instructors & Face-to-Face Help



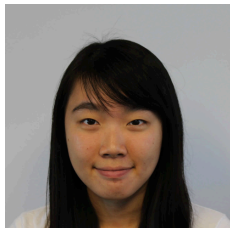

Professor: Lawrence Saul

Teaching assistants: Sahil Agarwal, Zhen Zhai

Discussion sessions: Wed/Fri 11am, CSE 2154

Tutor hours: Mon 11am, Thu 2:30pm, CSE 4262

Office hours: Mon/Fri 9am, CSE 3214

Prerequisites

- **Programming**
  - **Homeworks 2-9 will involve coding.**
  - **Also: basic data analysis and visualization.**
  - **Solutions accepted in any language!**
  - **Python, MATLAB, Java, C/C++, Perl, etc.**
  - **No hand-holding with compiling, debugging.**

Non-CS majors are welcome.

Prerequisites

- **Elementary probability**
  - Random variables
  - Expected values
- **Linear algebra**
  - Matrix multiplication, inverses
  - Solving systems of linear equations
- **Calculus**
  - Computing derivatives
  - Computing maxima and minima

Readings vs lectures

- **Readings**
  - No truly required texts.
  - Some handouts in class.

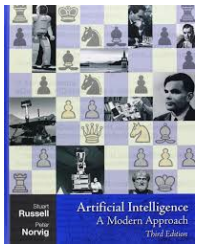

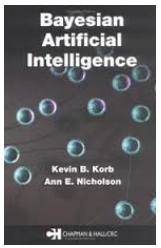

- **Lectures**
  - Designed to be self-contained.
  - Important for homework assignments.
  - Emphasis on mathematical development.
  - Blackboard, not powerpoint!

Homework (54%)

- **Rules of the game:**
  - No extensions: the course moves too fast.
  - Other late assignments: half credit.
  - No credit once solutions are out.
  - Submit in class on Tue/Thu (or Fri by 10am).
- **Best practices**
  - Keep up, or lectures will be hard to digest.
  - Please do not typeset your solutions (except for source code).

Do's and don'ts

- **What is allowed:**
  - **You may work in groups to start the problems (but not to finish them).**
  - **Write up all your own work.**
  - **You may consult published texts.**
- **What is not allowed:**
  - **Using old course materials.**
  - **Copying from current or former students.**
  - **Uploading current materials to archives.**

Academic dishonesty will be severely punished.

Online resources

- **Piazza**
  - Posting of homeworks, notes, handouts
  - Online Q/A and discussion boards

https://piazza.com/ucsd/summer2016/cse150/home

- **Gradesource**
  - Scores on homeworks and final
  - Please check routinely for accuracy

http://www.gradesource.com