

CSE 250A. Intro to AI

Probabilistic Reasoning and Decision-Making



Welcome to CSE 250A!

“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/article/The-Promising-Syllabus/46748/>

<http://cseweb.ucsd.edu/classes/fa15/cse250A-a/>

CSE 250A. Principles of Artificial Intelligence: Probabilistic Reasoning and Decision-Making

Administrivia	Syllabus	Piazza	GradeSource
-------------------------------	--------------------------	------------------------	-----------------------------

Subject

Probabilistic methods for reasoning and decision-making under uncertainty. Topics include: inference and learning in directed probabilistic graphical models; prediction and planning in Markov decision processes; applications to computer vision, robotics, speech recognition, natural language processing, and information retrieval.

Prerequisites

The course is aimed broadly at advanced undergraduates and beginning graduate students in mathematics, science, and engineering. Prerequisites are elementary probability, multivariable calculus, linear algebra, and basic programming ability in some high-level language such as C, Java, or Matlab. Programming assignments are completed in the language of the student's choice.

Relation to other courses

CSE 250a covers largely the same topics as [CSE 150](#) (as I teach it), but at a faster pace and more advanced mathematical level. The homework assignments and exams in CSE 250A are also longer and more challenging. In general you should not take CSE 250a if you already have taken CSE 150 from me in a previous quarter.

Administrivia

- **Instructor:** [Lawrence Saul](#)
- **Teaching assistants (TAs):**
 - Kristjan Jonsson (kjonsson@eng.ucsd.edu)
 - Amanda Song (mas065@ucsd.edu)
 - Chaitanya (Chay) Ryali (rckrishn@eng.ucsd.edu)
- **Lectures:** Tue/Thu 12:30-1:50 pm, PCYNH 109
- **TA sections:** TBA
- **Instructor office hour:** Thu 3-4 pm @ CSE 3214 (starting Oct 01)
- **TA office hours:** TBA
- **Grading:** homework (25%), two quizzes (40%), final exam (35%).

Textbooks

The course does not closely follow a particular text; the lectures are meant to be self-contained. Nevertheless, the following texts (though not required) may be useful as general references:

- [Artificial Intelligence: A Modern Approach](#) (Russell & Norvig, 2010)
- [Machine Learning: A Probabilistic Perspective](#) (Murphy, 2012)
- [Reinforcement Learning: An Introduction](#) (Sutton & Barto, 1998)
- [Pattern Recognition and Machine Learning](#) (Bishop, 2006)

Syllabus

Thu Sep 24	Administrivia and course overview.	
Tue Sep 29	Modeling uncertainty, review of probability, explaining away.	HW 1 out.
Thu Oct 01	Belief networks: from probabilities to graphs.	
Tue Oct 06	Conditional independence, d-separation, polytrees.	HW 1 due. HW 2 out.
Thu Oct 08	Algorithms for exact and approximate inference.	
Tue Oct 13	Maximum likelihood estimation; Markov models of language; naive Bayes models of text.	HW 2 due. HW 3 out.
Thu Oct 15	Linear and logistic regression. Numerical optimization.	
Tue Oct 20	Latent variable modeling. Expectation-Maximization (EM) algorithm. Auxiliary functions.	HW 3 due.
Thu Oct 22	EM algorithm: derivation, proof of convergence.	
Tue Oct 27	Quiz #1	HW 4 out.
Thu Oct 29	Examples of EM; applications to language modeling.	
Tue Nov 03	Hidden Markov models, automatic speech recognition, Viterbi algorithm.	HW 4 due. HW 5 out.
Thu Nov 05	Forward-backward algorithm, Gaussian mixture models, Kalman filters.	
Tue Nov 10	Reinforcement learning (RL), Markov decision processes.	HW 5 due. HW 6 out.
Thu Nov 12	Policy evaluation, policy improvement.	
Tue Nov 17	Policy iteration, value iteration.	HW 6 due.
Thu Nov 19	Stochastic approximation theory, temporal difference prediction.	
Tue Nov 24	Quiz #2	HW 7 out.
Thu Nov 26	Thanksgiving Holiday	
Tue Dec 01	Q-learning, extensions of RL.	
Thu Dec 03	Bonus topic (if time).	HW 7 due.
Fri Dec 11	Final exam.	

Teaching Assistants

- **Kristjan Jonsson (A-D)**
- **Chay Riyali (E-P)**
- **Amanda Song (Q-Z)**



(Bonus: you can get help in at least four different languages...)

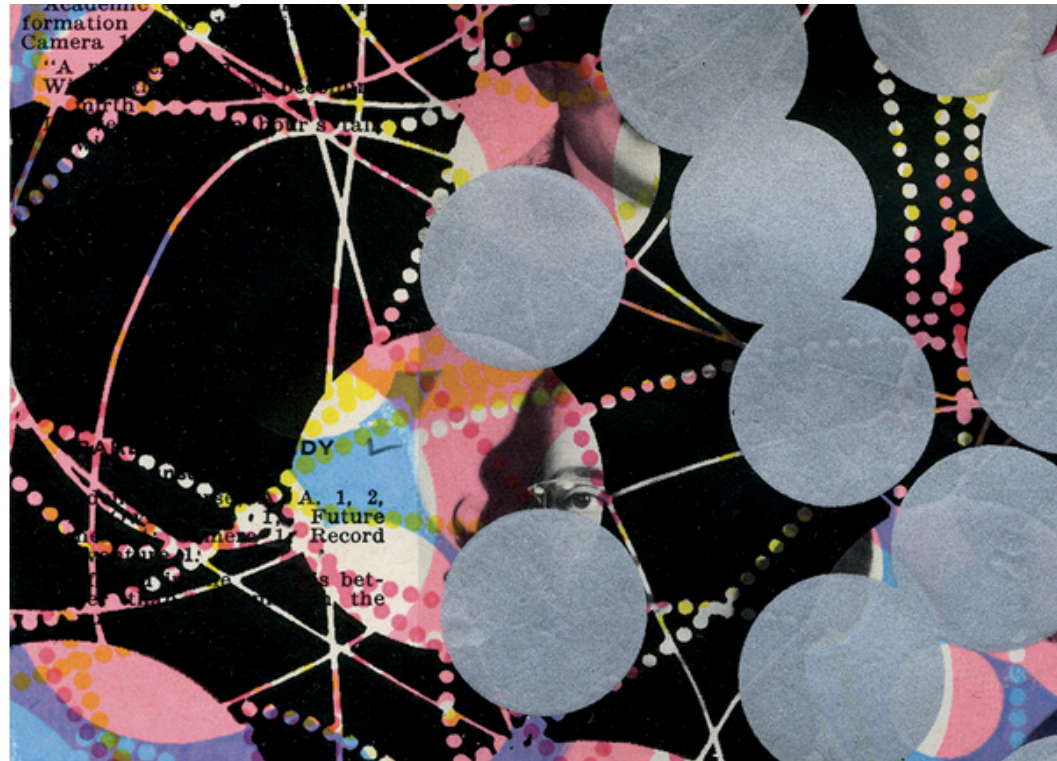
Enrollment priority

- **CSE graduate students**
 - PhD students
 - MS students
- **PhD students in other programs**
 - Bioinformatics
 - Cognitive science
 - ECE
 - Linguistics
- **MS students in other programs**

Who uses probabilistic methods in AI and ML?

- **Search & Ads** – Google, Microsoft, Yahoo
- **Sales & Recommendations** – Amazon, Etsy
- **Social media** – Facebook, Twitter, LinkedIn
- **Gaming & HCI** – XBox, Wii
- **Forensics & signal analysis** – FBI, NSA
- **Data science & analytics**

“Every company is a data company.”



DATA

Data Scientist: The Sexiest Job of the 21st Century

Prerequisites

- **Programming**

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

Non-CS backgrounds are welcome.

Prerequisites

- **Elementary probability**
 - Random variables
 - Expected values
- **Multivariable calculus**
 - Chain rule
 - Gradients and partial derivatives
 - Computing maxima and minima
 - Constrained optimization
 - Lagrange multipliers

Prerequisites

- **Linear algebra**
 - Vectors and matrices
 - Matrix multiplication, inverse, determinants
 - Systems of linear equations
- **Mathematical maturity**
 - Patience and persistence
 - Willingness to fill in gaps
 - Not for “hackers”

Readings versus lectures

- **Readings**

- No required texts.
- Some handouts (on web site).

- **Lectures**

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

Homework

- **Weekly problem sets**
 - Distributed (usually) on Tuesdays.
 - Due (usually) on Tuesdays in class.
- **Rules of the game**
 - One two-day extension (but only on HWs 2-6).
 - Other late assignments: no credit.
 - Okay to work in groups.
 - Write up your own solutions.
 - Typesetting is okay but not required.

Grading

- **Breakdown**
 - 25% problem sets
 - 40% two in-class quizzes (closed book)
 - 35% final exam (closed book)
- **Academic dishonesty**
 - Neither ethical nor in your self-interest.
 - Always credit your sources.
 - Plagiarism is severely punished.

Online resources

- **Web page**

- Syllabus, office hours, etc.
- Not much else

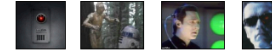
- **Piazza**

- Homework assignments
- Q/A forums
- Sign up ASAP

- **GradeSource**

- You will be emailed an ID after HW #1.
- It is your responsibility to check grades.

CSE 250A. Principles of Artificial Intelligence: Probabilistic Reasoning and Decision-Making



[Administrivia](#) [Syllabus](#) [Piazza](#) [GradeSource](#)

Subject

Probabilistic methods for reasoning and decision-making under uncertainty. Topics include: inference and learning in directed probabilistic graphical models; prediction and planning in Markov decision processes; applications to computer vision, robotics, speech recognition, natural language processing, and information retrieval.

Prerequisites

The course is aimed broadly at advanced undergraduates and beginning graduate students in mathematics, science, and engineering. Prerequisites are elementary probability, multivariable calculus, linear algebra, and basic programming ability in some high-level language such as C, Java, or Matlab. Programming assignments are completed in the language of the student's choice.

Relation to other courses

CSE 150 covers largely the same material as CSE 250A, but at a slower pace and less advanced mathematical level. The homework assignments (and exams) in CSE 250A are longer and more challenging.

Administrivia

- Professor: Lawrence Saul
- Teaching assistants: Sheeraz Ahmad and Long Jin
- Lectures: Tue/Thu 12:30 am - 1:50 pm, HSS 1330
- Sections: TBA
- Instructor office hour: Fri 9-10 am.
- TA office hours: TBD
- Grading: homework (25%), two quizzes (40%), final exam (35%).

Textbooks

The course does not closely follow a particular text; the lectures are meant to be self-contained. Nevertheless, the following texts (though not required) may be useful as general references:

- *Artificial Intelligence: A Modern Approach* (Russell & Norvig, 2010)
- *Machine Learning: A Probabilistic Perspective* (Murphy, 2012)
- *Reinforcement Learning: An Introduction* (Sutton & Barto, 1998)
- *Pattern Recognition and Machine Learning* (Bishop, 2006)

Syllabus

Thu Oct 02	Administrivia and course overview.	
Tue Oct 07	Modeling uncertainty, review of probability, explaining away.	HW 1 out.
Thu Oct 09	Belief networks: from probabilities to graphs.	

Not in CSE 250A

- **Things we won't cover**
 - Mathematical logic
 - Traditional search (A^*)
 - Theorem proving
 - Genetic algorithms
 - Philosophy of AI