CS343H
Honors Artificial Intelligence

Prof: Peter Stone

Department of Computer Science
The University of Texas at Austin

Logistics

- Late assignments accepted until Tuesday (12/12)
- Final: Monday Dec. 18th, 2pm-5pm
  - Open notes - handwritten (2 pages)
  - No books, no printouts, no electronics

Final Exam

- Monday Dec. 18th, 2pm-5pm

Final Exam

- Monday Dec. 18th, 2pm-5pm
- TAs and I will proctor (I may need to leave for part in the middle)

Final Exam

- Monday Dec. 18th, 2pm-5pm
- TAs and I will proctor (I may need to leave for part in the middle)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning

Final Exam

- Monday Dec. 18th, 2pm-5pm
- TAs and I will proctor (I may need to leave for part in the middle)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm

Final Exam

- Monday Dec. 18th, 2pm-5pm
- TAs and I will proctor (I may need to leave for part in the middle)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm
- 3 hours rather than 1 hour and 15 minutes

Final Exam

- Monday Dec. 18th, 2pm-5pm
- TAs and I will proctor (I may need to leave for part in the middle)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm
- 3 hours rather than 1 hour and 15 minutes
- Samples - Berkeley exams

Question

- Would you have rather been born 100 years earlier or 100 years later?

Question

- Does it matter to you if our "descendents" aren't human?

Question

- If an AI technology runs amok, who is responsible?

Question

- If an AI technology runs amok, who is responsible?
- Are there some types of research we shouldn't do?

Question

- Can computers perfectly simulate a human's decision-making (weak AI)?

Question

- Can computers perfectly simulate a human's decision-making (weak AI)?
- Will computers ever be better than people at everything?

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**
- **Next:** MDPs —

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**
- **Next:** MDPs — towards reinforcement learning

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**
- **Next:** MDPs — towards reinforcement learning
  - Still know transition and reward function

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**
- **Next:** MDPs — towards reinforcement learning
  - Still know transition and reward function
  - Looking for a **policy** — optimal action from every state

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**
- **Next:** MDPs — towards reinforcement learning
  - Still know transition and reward function
  - Looking for a **policy** — optimal action from every state
- **Action learning:** Reinforcement learning

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**
- **Next:** MDPs — towards reinforcement learning
  - Still know transition and reward function
  - Looking for a **policy** — optimal action from every state
- **Action learning:** Reinforcement learning
  - Policy without knowing transition or reward functions

Course Recap

- **First weeks:** search (BFS, A*, minimax, alpha-beta)
  - Find an optimal plan (or solution)
  - Best thing to do from the current state
  - Know transition and cost (reward) functions
  - Either execute complete solution (deterministic) or search again at every step
  - **Know current state**
- **Next:** MDPs — towards reinforcement learning
  - Still know transition and reward function
  - Looking for a **policy** — optimal action from every state
- **Action learning:** Reinforcement learning
  - Policy without knowing transition or reward functions
  - **Still know state**

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 4:** Utilities

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 4:** Utilities
  - **Week 7:** Conditional independence and inference (exact and approximate)

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 4:** Utilities
  - **Week 7:** Conditional independence and inference (exact and approximate)
  - **Week 9:** State estimation over time

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 4:** Utilities
  - **Week 7:** Conditional independence and inference (exact and approximate)
  - **Week 9:** State estimation over time
  - **Week 9:** Utility-based decisions

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 4:** Utilities
  - **Week 7:** Conditional independence and inference (exact and approximate)
  - **Week 9:** State estimation over time
  - **Week 9:** Utility-based decisions
- **Week 10:** What if they're not known?

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 4:** Utilities
  - **Week 7:** Conditional independence and inference (exact and approximate)
  - **Week 9:** State estimation over time
  - **Week 9:** Utility-based decisions
- **Week 10:** What if they're not known?
  - Also Bayesian networks for **classification**

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 4:** Utilities
  - **Week 7:** Conditional independence and inference (exact and approximate)
  - **Week 9:** State estimation over time
  - **Week 9:** Utility-based decisions
- **Week 10:** What if they're not known?
  - Also Bayesian networks for **classification**
  - A type of **machine learning**

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)
  - **Week 12:** SVMs, Kernels, and Clustering

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)
  - **Week 12:** SVMs, Kernels, and Clustering
- **Week 13:** Classical planning
  - Reasoning with first order representations

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)
  - **Week 12:** SVMs, Kernels, and Clustering
- **Week 13:** Classical planning
  - Reasoning with first order representations
  - So far we had only dealt with propositions

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)
  - **Week 12:** SVMs, Kernels, and Clustering
- **Week 13:** Classical planning
  - Reasoning with first order representations
  - So far we had only dealt with propositions
  - Back to known transitions, known state, etc.

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)
  - **Week 12:** SVMs, Kernels, and Clustering
- **Week 13:** Classical planning
  - Reasoning with first order representations
  - So far we had only dealt with propositions
  - Back to known transitions, known state, etc.
- **Week 14:** Philosophical foundations and ethics

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)
  - **Week 12:** SVMs, Kernels, and Clustering
- **Week 13:** Classical planning
  - Reasoning with first order representations
  - So far we had only dealt with propositions
  - Back to known transitions, known state, etc.
- **Week 14:** Philosophical foundations and ethics

It's all about building agents

Sense, decide, act

Course Recap (cont.)

- **After that:** More machine learning
  - **Week 11:** Perceptrons and Neural Nets (Deep Learning)
  - **Week 12:** SVMs, Kernels, and Clustering
- **Week 13:** Classical planning
  - Reasoning with first order representations
  - So far we had only dealt with propositions
  - Back to known transitions, known state, etc.
- **Week 14:** Philosophical foundations and ethics

It's all about building agents

Sense, decide, act **Maximize expected utility**

Topics not covered

- Knowledge representation and reasoning
  . (Chapters 7-9, 11, 12)
- Game theory and auctions (Sections 17.5, 17.6)
- Aspects of learning (Chapters 18, 19)
- Natural language (Chapters 22, 23)
- Vision (Chapter 24)
- Robotics (Chapter 25)

Surveys

- TA's and my surveys
- Negative **and** positive feedback useful

Surveys

- TA's and my surveys
- Negative **and** positive feedback useful
- Invitation to send more feedback by email

Surveys

- TA's and my surveys
- Negative **and** positive feedback useful
- Invitation to send more feedback by email
  - When I teach the course next, what should I do the same? What should change?

Surveys

- TA's and my surveys
- Negative **and** positive feedback useful
- Invitation to send more feedback by email
  - When I teach the course next, what should I do the same? What should change?
- Most important: course rating, instructor rating, written comments

My Perspective

- I've enjoyed teaching this class!

My Perspective

- I've enjoyed teaching this class!
- I've been impressed by the levels of questions and understanding

My Perspective

- I've enjoyed teaching this class!
- I've been impressed by the levels of questions and understanding
  - You kept me on my toes
  - I learned tons!

My Perspective

- I've enjoyed teaching this class!
- I've been impressed by the levels of questions and understanding
  - You kept me on my toes
  - I learned tons!
- Thanks to Josiah, Yinan, and Rohan for handling all the assignments!

My Perspective

- I've enjoyed teaching this class!
- I've been impressed by the levels of questions and understanding
  - You kept me on my toes
  - I learned tons!
- Thanks to Josiah, Yinan, and Rohan for handling all the assignments!
- I'm proud of all of you for sticking with it through what I think was a demanding course

My Perspective

- I've enjoyed teaching this class!
- I've been impressed by the levels of questions and understanding
  - You kept me on my toes
  - I learned tons!
- Thanks to Josiah, Yinan, and Rohan for handling all the assignments!
- I'm proud of all of you for sticking with it through what I think was a demanding course

THANKS!!!