CS343H Honors Artificial Intelligence

Prof: Peter Stone

Department of Computer Science The University of Texas at Austin

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



• Monday Dec. 18th, 2pm-5pm



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



- 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



- 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



- 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



- 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



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





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



Peter Stone



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



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



 Can computers perfectly simulate a human's decisionmaking (weak AI)?



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



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



- 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



- 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



- 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



- 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 —



- 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



- 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



- 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



- 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



- 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



- 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



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



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



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



- **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)



- **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



- **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



- **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?



- 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**



- 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

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



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



- 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



- 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



- 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.



- 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



- 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



- 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
- Aspects of learning
- Natural language
- Vision
- Robotics

- (Sections 17.5, 17.6)
 - (Chapters 18, 19)
 - (Chapters 22, 23)
 - (Chapter 24)
 - (Chapter 25)



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



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



- 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?



- 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



• I've enjoyed teaching this class!



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



- 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!



- 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'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



- 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



