CS343
Artificial Intelligence

Prof: Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Afternoon, Colleagues

Good Afternoon, Colleagues

Are there any questions?

Logistics

- Classification assignment: Avg: 22.8 (out of 25), stddev: 2.79

Logistics

- Classification assignment: Avg: 22.8 (out of 25), stddev: 2.79
  - Mini-contest: 1. Jonny (86.2%) 2. George (85.6%) 3. Craig (76.4%)

Logistics

- Classification assignment: Avg: 22.8 (out of 25), stddev: 2.79
  - Mini-contest: 1. Jonny (86.2%) 2. George (85.6%) 3. Craig (76.4%)
- Late readings accepted until the night before the exam (5/9)
- Final: Thursday May 10th, 9am-noon

Tournament Qualifiers

-

- ShotCallaKoalas, NikolaAndTheNomNoms, SHARON, ManicM
  Mayonnaise, AgentStanley, AgentSea, MediocreAgentsV2,
  GreyhoundAgents, Magnesium, Grepped, CaesarAgents,
  helloAgent, PsychoAgents, TenPoolWithSpeedAgents, O__O,
  NotMelsPacman, GentlemenAgents, RoleReversal

Tournament Qualifiers

-

- ShotCallaKoalas, NikolaAndTheNomNoms, SHARON, ManicM
  Mayonnaise, AgentStanley, AgentSea, MediocreAgentsV2,
  GreyhoundAgents, Magnesium, Grepped, CaesarAgents,
  helloAgent, PsychoAgents, TenPoolWithSpeedAgents, O__O,
  NotMelsPacman, GentlemenAgents, RoleReversal

Congratulations to all!

Pending questions

- How could we verify strong AI?
- What's the role of emotions?
- What's the role of the Turing test?
- Can AI achieve "thinking"?
- When will there be an ultraintelligent agent?
- Does it matter to you if our "descendents" aren't human?

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
- **Before Midterm:** 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
- **Before Midterm:** 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
- **Before Midterm:** 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 5:** Probability and utilities

Course Recap (cont.)

- **Probabilistic Reasoning:** Now state is unknown
- Bayesian networks – state estimation/inference
- **Prior, net structure, and CPT's known**
  - **Week 5:** Probability and utilities
  - **Week 9:** 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 5:** Probability and utilities
  - **Week 9:** Conditional independence and inference (exact and approximate)
  - **Week 10:** Exact 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 5:** Probability and utilities
  - **Week 9:** Conditional independence and inference (exact and approximate)
  - **Week 10:** Exact state estimation over time
  - **Week 11:** Approximate 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 5:** Probability and utilities
  - **Week 9:** Conditional independence and inference (exact and approximate)
  - **Week 10:** Exact state estimation over time
  - **Week 11:** Approximate state estimation over time
- **Week 12:** 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 5:** Probability and utilities
  - **Week 9:** Conditional independence and inference (exact and approximate)
  - **Week 10:** Exact state estimation over time
  - **Week 11:** Approximate state estimation over time
- **Week 12:** 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 5:** Probability and utilities
  - **Week 9:** Conditional independence and inference (exact and approximate)
  - **Week 10:** Exact state estimation over time
  - **Week 11:** Approximate state estimation over time
- **Week 12:** What if they're not known?
  - Also Bayesian networks for **classification**
  - A type of **machine learning**

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification
  - Perceptrons, SVMs, MIRA
  - Didn't cover unsupervised learning

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification
  - Perceptrons, SVMs, MIRA
  - Didn't cover unsupervised learning
- **Week 14:** Classical planning
  - Reasoning with first order representations

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification
  - Perceptrons, SVMs, MIRA
  - Didn't cover unsupervised learning
- **Week 14:** Classical planning
  - Reasoning with first order representations
  - So far we had dealt only with propositions

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification
  - Perceptrons, SVMs, MIRA
  - Didn't cover unsupervised learning
- **Week 14:** Classical planning
  - Reasoning with first order representations
  - So far we had dealt only with propositions
  - Back to known transitions, known state, etc.

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification
  - Perceptrons, SVMs, MIRA
  - Didn't cover unsupervised learning
- **Week 14:** Classical planning
  - Reasoning with first order representations
  - So far we had dealt only with propositions
  - Back to known transitions, known state, etc.
- **Week 15:** Philosophical foundations and ethics

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification
  - Perceptrons, SVMs, MIRA
  - Didn't cover unsupervised learning
- **Week 14:** Classical planning
  - Reasoning with first order representations
  - So far we had dealt only with propositions
  - Back to known transitions, known state, etc.
- **Week 15:** Philosophical foundations and ethics

It's all about building agents

Sense, decide, act

Course Recap (cont.)

- **Week 13:** Machine Learning
  - Just a taste – focus on concept learning = classification
  - Perceptrons, SVMs, MIRA
  - Didn't cover unsupervised learning
- **Week 14:** Classical planning
  - Reasoning with first order representations
  - So far we had dealt only with propositions
  - Back to known transitions, known state, etc.
- **Week 15:** Philosophical foundations and ethics

It's all about building agents

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

Topics not covered

- Constraint satisfaction (Chapter 6)
- 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)

Final Exam

- Thursday May 10th, 9am-noon

Final Exam

- Thursday May 10th, 9am-noon
- I may be late - Daniel will proctor beginning

Final Exam

- Thursday May 10th, 9am-noon
- I may be late - Daniel will proctor beginning
  - Office hours on Tuesday may change
  - Available by appointment if you want to meet

Final Exam

- Thursday May 10th, 9am-noon
- I may be late - Daniel will proctor beginning
  - Office hours on Tuesday may change
  - Available by appointment if you want to meet
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning

Final Exam

- Thursday May 10th, 9am-noon
- I may be late - Daniel will proctor beginning
  - Office hours on Tuesday may change
  - Available by appointment if you want to meet
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm

Final Exam

- Thursday May 10th, 9am-noon
- I may be late - Daniel will proctor beginning
  - Office hours on Tuesday may change
  - Available by appointment if you want to meet
- 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

- Thursday May 10th, 9am-noon
- I may be late - Daniel will proctor beginning
  - Office hours on Tuesday may change
  - Available by appointment if you want to meet
- 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
- Sample - Dan Klein's Spring 2009 exam with solutions

My Perspective

- I've enjoyed teaching this class!

My Perspective

- I've enjoyed teaching this class!
- Wasn't sure at first if I could cover all this material

My Perspective

- I've enjoyed teaching this class!
- Wasn't sure at first if I could cover all this material
- I've been impressed by the levels of questions and understanding

My Perspective

- I've enjoyed teaching this class!
- Wasn't sure at first if I could cover all this material
- 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!
- Wasn't sure at first if I could cover all this material
- I've been impressed by the levels of questions and understanding
  - You kept me on my toes
  - I learned tons!
- Thanks to Daniel for handling all the programming assignments!

My Perspective

- I've enjoyed teaching this class!
- Wasn't sure at first if I could cover all this material
- I've been impressed by the levels of questions and understanding
  - You kept me on my toes
  - I learned tons!
- Thanks to Daniel for handling all the programming 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!
- Wasn't sure at first if I could cover all this material
- I've been impressed by the levels of questions and understanding
  - You kept me on my toes
  - I learned tons!
- Thanks to Daniel for handling all the programming assignments!
- I'm proud of all of you for sticking with it through what I think was a demanding course

THANKS!!!

Surveys

- Daniel's and my surveys
- Positive **and** negative feedback useful
- Invitation to send more feedback by email
  - If/when I teach the course again, how should it change?
- Most important: course rating, instructor rating, written comments