

# **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

---

- Late readings accepted until the night before the exam (5/11)
- Final: Wednesday May 12th, 2-5pm

# Tournament Qualifiers

---

- dtkraftAgents (David)
- NormChomsky (Lon)
- DbAgents (Diego)
- AdamNick (Adam and Nick)
- SpecialAgents (Jackie and Dustin)
- TheAvengingBlowfish (Bethany)

# Round Robin

---

1. AdamNick: 7 matches won (qualified)
2. SpecialAgents: 6 matches won (qualified)
3. DbAgents: 6 matches won (qualified)
4. dtkraftAgent: 5 matches won (qualified)
5. TheAvengingBlowfish: 2 matches won
6. NormChompsky: 1 matches won

# Semifinals

---

- AdamNick vs. dtkraftAgent

# Semifinals

---

- AdamNick vs. dtkraftAgent (7-2)



# Semifinals

---

- AdamNick vs. dtkraftAgent (7-2)
- SpecialAgents vs. DbAgents

# Semifinals

---

- AdamNick vs. dtkraftAgent (7-2)
- SpecialAgents vs. DbAgents (0-9)

# Semifinals

---

- AdamNick vs. dtkraftAgent (7-2)
- SpecialAgents vs. DbAgents (0-9)

**3rd place game:** dtkraftAgent vs. SpecialAgents

# Semifinals

---

- AdamNick vs. dtkraftAgent (7-2)
- SpecialAgents vs. DbAgents (0-9)

**3rd place game:** dtkraftAgent vs. SpecialAgents (7-2)

# Semifinals

---

- AdamNick vs. dtkraftAgent (7-2)
- SpecialAgents vs. DbAgents (0-9)

**3rd place game:** dtkraftAgent vs. SpecialAgents (7-2)

**Final:** AdamNick vs. DbAgents

# Semifinals

---

- AdamNick vs. dtkraftAgent (7-2)
- SpecialAgents vs. DbAgents (0-9)

**3rd place game:** dtkraftAgent vs. SpecialAgents (7-2)

**Final:** AdamNick vs. DbAgents (8-1)

# Final Rankings

---

1. AdamNick (5% bonus)
2. DbAgents (4% bonus)
3. dtkraftAgent (3% bonus)
4. SpecialAgents (1% bonus)
5. TheAvengingBlowfish (1% bonus)
6. NormChompsky (1% bonus)

**Congratulations to all!**

# Pending Questions

---

- Is *qualia* important? (subjective experience)



# Pending Questions

---

- Is *qualia* important? (subjective experience)
- What's so important about agents being in the physical world?

# Pending Questions

---

- Is *qualia* important? (subjective experience)
- What's so important about agents being in the physical world?
- Is the brain a machine?
  - Is the brain just a bunch of neurons? Can they be replaced?
  - What is the mind? (Is there a soul?)

# Pending Questions

---

- Is *qualia* important? (subjective experience)
- What's so important about agents being in the physical world?
- Is the brain a machine?
  - Is the brain just a bunch of neurons? Can they be replaced?
  - What is the mind? (Is there a soul?)
- What do I think about ethical side?
  - Who's liable for a crash?
  - Would declaring AI unethical stop progress?

# Pending Questions

---

- Is *qualia* important? (subjective experience)
- What's so important about agents being in the physical world?
- Is the brain a machine?
  - Is the brain just a bunch of neurons? Can they be replaced?
  - What is the mind? (Is there a soul?)
- What do I think about ethical side?
  - Who's liable for a crash?
  - Would declaring AI unethical stop progress?
- What am I looking forward to?

# 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

---

- Wednesday May 12th, 2-5pm

# Final Exam

---

- Wednesday May 12th, 2-5pm
- I will be out of town - Daniel will proctor

# Final Exam

---

- Wednesday May 12th, 2-5pm
- I will be out of town - Daniel will proctor
  - Leaving on Sunday, but in email contact
  - Available by appointment tomorrow (Friday)

# Final Exam

---

- Wednesday May 12th, 2-5pm
- I will be out of town - Daniel will proctor
  - Leaving on Sunday, but in email contact
  - Available by appointment tomorrow (Friday)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning

# Final Exam

---

- Wednesday May 12th, 2-5pm
- I will be out of town - Daniel will proctor
  - Leaving on Sunday, but in email contact
  - Available by appointment tomorrow (Friday)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm

# Final Exam

---

- Wednesday May 12th, 2-5pm
- I will be out of town - Daniel will proctor
  - Leaving on Sunday, but in email contact
  - Available by appointment tomorrow (Friday)
- 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

---

- Wednesday May 12th, 2-5pm
- I will be out of town - Daniel will proctor
  - Leaving on Sunday, but in email contact
  - Available by appointment tomorrow (Friday)
- 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



# Surveys

---

- Daniel's and my surveys
- Positive **and** negative feedback useful
- Invitation to send more feedback by email

# 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?

# 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