

CS343

Artificial Intelligence

Prof: Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Morning, Colleagues



Good Morning, Colleagues

Are there any questions?

Logistics

- Final

Logistics

- Final
- Stephanie Tellex talk on Friday

Logistics

- Final
- Stephanie Tellex talk on Friday
- Multiagent assignment questions?

Some Context

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

Some Context

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

Some Context

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

Some Context

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

Some Context

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

Some Context

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

Some Context

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

Some Context

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

Some Context

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