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

- Questions about the syllabus?

Logistics

- Questions about the syllabus?
- Class registration

Logistics

- Questions about the syllabus?
- Class registration
- Problems with the assignment?

Logistics

- Questions about the syllabus?
- Class registration
- Problems with the assignment?
- Piazza — useful discussion yesterday

Logistics

- Questions about the syllabus?
- Class registration
- Problems with the assignment?
- Piazza — useful discussion yesterday
  - CC Kim (houck@cs), and me on everything

Logistics

- Questions about the syllabus?
- Class registration
- Problems with the assignment?
- Piazza — useful discussion yesterday
  - CC Kim (houck@cs), and me on everything
- Assignments up through week 3

Logistics

- Questions about the syllabus?
- Class registration
- Problems with the assignment?
- Piazza — useful discussion yesterday
  - CC Kim (houck@cs), and me on everything
- Assignments up through week 3

Example Intelligent (autonomous) Agents

- Autonomous robot

Example Intelligent (autonomous) Agents

- Autonomous robot
- Information gathering agent
  - Find me the cheapest?

Example Intelligent (autonomous) Agents

- Autonomous robot
- Information gathering agent
  - Find me the cheapest?
- E-commerce agents
  - Decides what to buy/sell and does it

Example Intelligent (autonomous) Agents

- Autonomous robot
- Information gathering agent
  - Find me the cheapest?
- E-commerce agents
  - Decides what to buy/sell and does it
- Air-traffic controller

Example Intelligent (autonomous) Agents

- Autonomous robot
- Information gathering agent
  - Find me the cheapest?
- E-commerce agents
  - Decides what to buy/sell and does it
- Air-traffic controller
- Meeting scheduler

Example Intelligent (autonomous) Agents

- Autonomous robot
- Information gathering agent
  - Find me the cheapest?
- E-commerce agents
  - Decides what to buy/sell and does it
- Air-traffic controller
- Meeting scheduler
- Computer-game-playing agent

Not Intelligent Agents

- Thermostat
- Telephone
- Answering machine
- Pencil
- Java object

Environments

Environment $\Longrightarrow$ sensations, actions

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential
- static vs. dynamic

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous
- known vs. unknown

Student Examples

- game bot
- robot waiter
- bowling robot, ping pong player
- kiva robots, Mars rover, robot suturing agent
- Wall-E
- Words with friends word checker
- thermostat
- trading agent
- Siri
- Briggo
- piano playing agent
- unhappiness agent

BE a learning agent

BE a learning agent

- You, as a group, act as a learning agent

BE a learning agent

- You, as a group, act as a learning agent
- **Actions**: Wave, Stand, Clap

BE a learning agent

- You, as a group, act as a learning agent
- **Actions**: Wave, Stand, Clap
- **Observations**: colors, reward

BE a learning agent

- You, as a group, act as a learning agent
- **Actions**: Wave, Stand, Clap
- **Observations**: colors, reward
- **Goal**: Find an optimal *policy*

BE a learning agent

- You, as a group, act as a learning agent
- **Actions**: Wave, Stand, Clap
- **Observations**: colors, reward
- **Goal**: Find an optimal *policy*
  - Way of selecting actions that gets you the most reward

How did you do it?

How did you do it?

- What is your policy?
- What does the world look like?

How did you do it?

- What is your policy?
- What does the world look like?

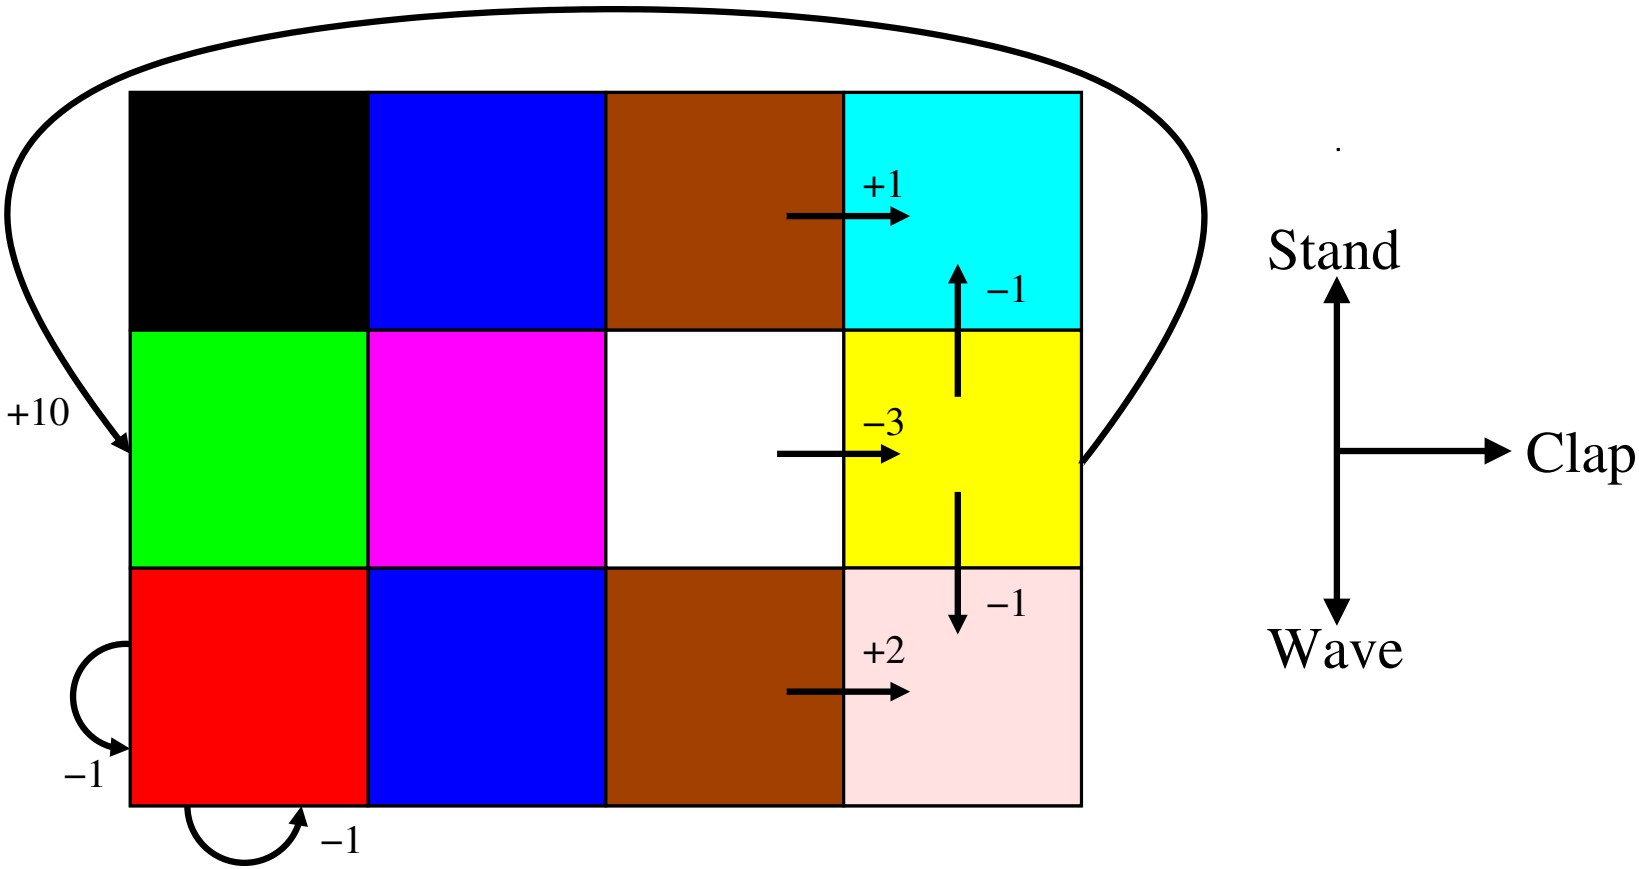

Formalizing what Just Happened

Knowns:

Formalizing what Just Happened

Knowns:

- $\mathcal{O} = \{\text{Blue}, \text{Red}, \text{Green}, \text{Black}, \ldots\}$
- Rewards in $\mathbb{R}$
- $\mathcal{A} = \{Wave, Clap, Stand\}$

$$o_0, a_0, r_0, o_1, a_1, r_1, o_2, \ldots$$

Formalizing what Just Happened

Knowns:

- $\mathcal{O} = \{\text{Blue}, \text{Red}, \text{Green}, \text{Black}, \ldots\}$
- Rewards in $\mathbb{R}$
- $\mathcal{A} = \{Wave, Clap, Stand\}$

$$o_0, a_0, r_0, o_1, a_1, r_1, o_2, \ldots$$

Unknowns:

Formalizing what Just Happened

Knowns:

- $\mathcal{O} = \{\text{Blue}, \text{Red}, \text{Green}, \text{Black}, \ldots\}$
- Rewards in $\mathbb{R}$
- $\mathcal{A} = \{Wave, Clap, Stand\}$

$$o_0, a_0, r_0, o_1, a_1, r_1, o_2, \ldots$$

Unknowns:

- $\mathcal{S} =$ 4x3 grid
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
- $\mathcal{P} = \mathcal{S} \mapsto \mathcal{O}$
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$

Formalizing what Just Happened

Knowns:

- $\mathcal{O} = \{\text{Blue}, \text{Red}, \text{Green}, \text{Black}, \ldots\}$
- Rewards in $\mathbb{R}$
- $\mathcal{A} = \{Wave, Clap, Stand\}$

$$o_0, a_0, r_0, o_1, a_1, r_1, o_2, \ldots$$

Unknowns:

- $\mathcal{S} =$ 4x3 grid
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
- $\mathcal{P} = \mathcal{S} \mapsto \mathcal{O}$
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$

$o_i = \mathcal{P}(s_i)$

Formalizing what Just Happened

Knowns:

- $\mathcal{O} = \{\text{Blue, Red, Green, Black, } \ldots\}$
- Rewards in $\mathbb{R}$
- $\mathcal{A} = \{Wave, Clap, Stand\}$

$$o_0, a_0, r_0, o_1, a_1, r_1, o_2, \ldots$$

Unknowns:

- $\mathcal{S} =$ 4x3 grid
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
- $\mathcal{P} = \mathcal{S} \mapsto \mathcal{O}$
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$

$o_i = \mathcal{P}(s_i)$ $\qquad$ $r_i = \mathcal{R}(s_i, a_i)$

Formalizing what Just Happened

Knowns:

- $\mathcal{O} = \{\text{Blue}, \text{Red}, \text{Green}, \text{Black}, \ldots\}$
- Rewards in $\mathbb{R}$
- $\mathcal{A} = \{Wave, Clap, Stand\}$

$$o_0, a_0, r_0, o_1, a_1, r_1, o_2, \ldots$$

Unknowns:

- $\mathcal{S} =$ 4x3 grid
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
- $\mathcal{P} = \mathcal{S} \mapsto \mathcal{O}$
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$

$$o_i = \mathcal{P}(s_i) \qquad r_i = \mathcal{R}(s_i, a_i) \qquad s_{i+1} = \mathcal{T}(s_i, a_i)$$

Describe the environment

Environment $\Longrightarrow$ sensations, actions

Describe the environment

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)

Describe the environment

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent

Describe the environment

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)

Describe the environment

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential

Describe the environment

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential
- static vs. dynamic

Describe the environment

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous

Describe the environment

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- single-agent vs. multiagent
- deterministic vs. non-deterministic (stochastic)
- episodic vs. sequential
- static vs. dynamic
- discrete vs. continuous
- known vs. unknown

Next week: Search

- Textbook readings
- Responses both Monday and Wednesday
- Python tutorial due