CS344M
Autonomous Multiagent Systems

Prof: Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Afternoon, Colleagues

Are there any questions?

Good Afternoon, Colleagues

Are there any questions?

- Pending questions:
  - Can an agent stop being an agent due to environment
  - Does it need to be persistent?
  - Is social ability essential?
  - What does "autonomy" mean?
  - How can you categorize discrete/continuous?

Logistics

- First assignment: how did it go?

Logistics

- First assignment: how did it go?
- Next soccer assignment: score a goal

Logistics

- First assignment: how did it go?
- Next soccer assignment: score a goal
  - Help each other with C issues — parsing strings

Logistics

- First assignment: how did it go?
- Next soccer assignment: score a goal
  - Help each other with C issues — parsing strings
  - Evaluating mostly on the logic — does the agent "do the right thing?"

Logistics

- First assignment: how did it go?
- Next soccer assignment: score a goal
  - Help each other with C issues — parsing strings
  - Evaluating mostly on the logic — does the agent "do the right thing?"
- 2D or 3D?

Self-Introductions

- Speak loudly

Self-Introductions

- Speak loudly
- Name, year, major

Self-Introductions

- Speak loudly
- Name, year, major
- At least one other thing about yourself

Discussion

An autonomous agent is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to affect what it senses in the future.

- Is this a good definition?
- The authors claim is is a "formal" definition of agents. Is it?

Discussion

An autonomous agent is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to affect what it senses in the future.

- Is this a good definition?
- The authors claim is is a "formal" definition of agents. Is it?
- Can you do better?

Discussion

An autonomous agent is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to affect what it senses in the future.

- Is this a good definition?
- The authors claim is is a "formal" definition of agents. Is it?
- Can you do better?
- Do they need to be social? persistent?
- Can they cease to be agents in a different environment?
- Autonomy

Varieties of Autonomy

- Do we have complete freedom over our beliefs, goals, and actions?

Varieties of Autonomy

- Do we have complete freedom over our beliefs, goals, and actions?
- Software service has no autonomy — does what it's told.

Varieties of Autonomy

- Do we have complete freedom over our beliefs, goals, and actions?
- Software service has no autonomy — does what it's told.
- What's Wooldridge's take on where autonomous agents lie on the spectrum?

Varieties of Autonomy

- Do we have complete freedom over our beliefs, goals, and actions?
- Software service has no autonomy — does what it's told.
- What's Wooldridge's take on where autonomous agents lie on the spectrum?
  - Decide how to act so as to accomplish **delegated** goals

Varieties of Autonomy

- Do we have complete freedom over our beliefs, goals, and actions?
- Software service has no autonomy — does what it's told.
- What's Wooldridge's take on where autonomous agents lie on the spectrum?
  - Decide how to act so as to accomplish **delegated** goals
- Also mentions **adjustable** autonomy

My Requirements of Agents

- They must **sense** their environment.
- They must **decide** what action to take ("think").
- They must **act** in their environment.

My Requirements of Agents

- They must **sense** their environment.
- They must **decide** what action to take ("think").
- They must **act** in their environment.

Complete Agents

My Requirements of Agents

- They must **sense** their environment.
- They must **decide** what action to take ("think").
- They must **act** in their environment.

Complete Agents

Multiagent systems: Interact with other agents

My Requirements of Agents

- They must **sense** their environment.
- They must **decide** what action to take ("think").
- They must **act** in their environment.

Complete Agents

Multiagent systems: Interact with other agents

Learning agents: Improve performance from experience

My Requirements of Agents

- They must **sense** their environment.
- They must **decide** what action to take ("think").
- They must **act** in their environment.

Complete Agents

Multiagent systems: Interact with other agents

Learning agents: Improve performance from experience

Autonomous Bidding, Cognitive Systems,
Traffic management, **Robot Soccer**

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)
- deterministic vs. non-deterministic

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- deterministic vs. non-deterministic
- episodic vs. non-episodic

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- deterministic vs. non-deterministic
- episodic vs. non-episodic
- static vs. dynamic

Environments

Environment $\Longrightarrow$ sensations, actions

- fully observable vs. partially observable (accessible)
- deterministic vs. non-deterministic
- episodic vs. non-episodic
- static vs. dynamic
- discrete vs. continuous

Environments

Environment $\Longrightarrow$ sensations, actions

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

The Decision

The Decision

- reactive vs. deliberative

The Decision

- reactive vs. deliberative
- multiagent reasoning?

The Decision

- reactive vs. deliberative
- multiagent reasoning?
- learning?

Formalizing My Example

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

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

Standard/Reactive/State-based Agents

- Standard agent:

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$
- Reactive agent:

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$
- Reactive agent: $action : \mathcal{P} \mapsto \mathcal{A}$

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$
- Reactive agent: $action : \mathcal{P} \mapsto \mathcal{A}$
  - Decision based entirely on the present

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$
- Reactive agent: $action : \mathcal{P} \mapsto \mathcal{A}$
  - Decision based entirely on the present
- State-based agent:

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$
- Reactive agent: $action : \mathcal{P} \mapsto \mathcal{A}$
  - Decision based entirely on the present
- State-based agent: $action : \mathcal{I} \mapsto \mathcal{A}$, $next : \mathcal{I} \times \mathcal{P} \mapsto \mathcal{I}$

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$
- Reactive agent: $action : \mathcal{P} \mapsto \mathcal{A}$
  - Decision based entirely on the present
- State-based agent: $action : \mathcal{I} \mapsto \mathcal{A}$, $next : \mathcal{I} \times \mathcal{P} \mapsto \mathcal{I}$

  It is worth observing that state-based agents as defined here are in fact no more powerful than the standard agents we introduced earlier. In fact, they are *identical* in their expressive power.

Standard/Reactive/State-based Agents

- Standard agent: $action : \mathcal{P}^* \mapsto \mathcal{A}$
- Reactive agent: $action : \mathcal{P} \mapsto \mathcal{A}$
  - Decision based entirely on the present
- State-based agent: $action : \mathcal{I} \mapsto \mathcal{A}$, $next : \mathcal{I} \times \mathcal{P} \mapsto \mathcal{I}$

It is worth observing that state-based agents as defined here are in fact no more powerful than the standard agents we introduced earlier. In fact, they are *identical* in their expressive power.

Reactive agents for next Thursday's assignment task?

Discussion

What new autonomous agents do you expect to see in the next 10 years?