CS378
Autonomous Multiagent Systems
Spring 2005

Prof: Peter Stone
TA: Nate Kohl

Department or Computer Sciences
The University of Texas at Austin

Week 4b: Thursday, February 9th

Good Afternoon, Colleagues

Are there any questions?

Good Afternoon, Colleagues

Are there any questions?

- Scientific community? Good for domains other than Thm proving?
- Legacy systems — saved by agents?
- Maintaining a hierarchy like distributed systems / fault tolerance?
- Bayesian uncertainty

Logistics

- Programming assignment 4 - any questions?

Logistics

- Programming assignment 4 - any questions?
- Orchestra (as a MAS)

Bayes Rule

- You run a cancer clinic.
- 10% of people tested have cancer: $P(C) = .1$
- What are the chances a new patient has cancer?

Bayes Rule

- You run a cancer clinic.
- 10% of people tested have cancer: $P(C) = .1$
- What are the chances a new patient has cancer?
- You find out he smokes. Now what are the chances? $P(C|S)$

Bayes Rule

- You run a cancer clinic.
- 10% of people tested have cancer: $P(C) = .1$
- What are the chances a new patient has cancer?
- You find out he smokes. Now what are the chances? $P(C|S)$
- Assume you have detailed information on your patients (people tested), but not others.
- What more information do you need?

Bayes Rule

- You run a cancer clinic.
- 10% of people tested have cancer: $P(C) = .1$
- What are the chances a new patient has cancer?
- You find out he smokes. Now what are the chances? $P(C|S)$
- Assume you have detailed information on your patients (people tested), but not others.
- What more information do you need?

- 80% of those tested who have cancer smoke: $P(S|C)$ = .8

Bayes Rule

- You run a cancer clinic.
- 10% of people tested have cancer: $P(C) = .1$
- What are the chances a new patient has cancer?
- You find out he smokes. Now what are the chances? $P(C|S)$
- Assume you have detailed information on your patients (people tested), but not others.
- What more information do you need?

- 80% of those tested who have cancer smoke: $P(S|C)$ = .8
- 50% of those tested smoke: $P(S)$ = .5

Bayes Rule

- You run a cancer clinic.
- 10% of people tested have cancer: $P(C) = .1$
- What are the chances a new patient has cancer?
- You find out he smokes. Now what are the chances? $P(C|S)$
- Assume you have detailed information on your patients (people tested), but not others.
- What more information do you need?

- 80% of those tested who have cancer smoke: $P(S|C)$ = .8
- 50% of those tested smoke: $P(S)$ = .5

$$\text{Bayes: } P(C|S) = \frac{P(C)*P(S|C)}{P(S)}$$

Some Definitions

- **Distributed Computing :**

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.

- **Distributed AI :**

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.
- **Distributed AI :** Control as well as data is distributed. Focus on problem solving, communication, and coordination.

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.
- **Distributed AI :** Control as well as data is distributed. Focus on problem solving, communication, and coordination.
- **Distributed Problem Solving :**

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.
- **Distributed AI :** Control as well as data is distributed. Focus on problem solving, communication, and coordination.
- **Distributed Problem Solving :** Task decomposition and/or solution synthesis.

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.
- **Distributed AI :** Control as well as data is distributed. Focus on problem solving, communication, and coordination.
- **Distributed Problem Solving :** Task decomposition and/or solution synthesis.
- **Multiagent Systems :**

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.
- **Distributed AI :** Control as well as data is distributed. Focus on problem solving, communication, and coordination.
- **Distributed Problem Solving :** Task decomposition and/or solution synthesis.
- **Multiagent Systems :** Behavior coordination or behavior management.

Some Definitions

- **Distributed Computing :** Processors share data, but not control. Focus on low-level parallelization, synchronization.
- **Distributed AI :** Control as well as data is distributed. Focus on problem solving, communication, and coordination.
- **Distributed Problem Solving :** Task decomposition and/or solution synthesis.
- **Multiagent Systems :** Behavior coordination or behavior management.
  - No necessary guarantees about other agents.
  - Individual behaviors typically simple relative to interaction issues.

Multiagent Systems

- Study, behavior, construction of **possibly preexisting** autonomous agents that interact with each other.
  - incomplete information for agents
  - no global control
  - decentralized data
  - asynchronous computation

Why Multiagent Systems?

(7)

Why Multiagent Systems?

(7)

- Some domains require it. (Hospital scheduling)
- Interoperation of legacy systems (works?)
- Parallelism.
- Robustness.
- Scalability
- Simpler programming.
- "Intelligence is deeply and inevitably coupled with interaction." – *Gerhard Weiss*

Organizations

- Hierarchy:

Organizations

- **Hierarchy:** authority from above

Organizations

- **Hierarchy:** authority from above
- **Community of Experts:**

Organizations

- **Hierarchy:** authority from above
- **Community of Experts:** specialists, mutual adjustment

Organizations

- **Hierarchy:** authority from above
- **Community of Experts:** specialists, mutual adjustment
- **Market:**

Organizations

- **Hierarchy:** authority from above
- **Community of Experts:** specialists, mutual adjustment
- **Market:** bid for tasks and resources; contracts

Organizations

- **Hierarchy:** authority from above
- **Community of Experts:** specialists, mutual adjustment
- **Market:** bid for tasks and resources; contracts
- **Scientific community:**

Organizations

- **Hierarchy:** authority from above
- **Community of Experts:** specialists, mutual adjustment
- **Market:** bid for tasks and resources; contracts
- **Scientific community:** full solutions (perhaps with varying information) combined

Organizations

- **Hierarchy:** authority from above
- **Community of Experts:** specialists, mutual adjustment
- **Market:** bid for tasks and resources; contracts
- **Scientific community:** full solutions (perhaps with varying information) combined

– Scientific community? Good for domains other than Thm proving?
– Maintaining a hierarchy like distributed systems / fault tolerance?

Issues and Challenges

- How to break down and resynthesize the problem among agents

Issues and Challenges

- How to break down and resynthesize the problem among agents
- Communication/interaction protocols

Issues and Challenges

- How to break down and resynthesize the problem among agents
- Communication/interaction protocols
- Maintain coherence, stability: guarantees?
  - Coherence is a global property

Issues and Challenges

- How to break down and resynthesize the problem among agents
- Communication/interaction protocols
- Maintain coherence, stability: guarantees?
  - Coherence is a global property
- Representation by agents of each other and interactions

Issues and Challenges

- How to break down and resynthesize the problem among agents
- Communication/interaction protocols
- Maintain coherence, stability: guarantees?
  - Coherence is a global property
- Representation by agents of each other and interactions
- Reconciling different points of view

Issues and Challenges

- How to break down and resynthesize the problem among agents
- Communication/interaction protocols
- Maintain coherence, stability: guarantees?
  - Coherence is a global property
- Representation by agents of each other and interactions
- Reconciling different points of view
- Engineering

Dimensions and issues

- cooperative vs. competitive
- communication
- trust
- recursive modeling
- coalititions
- game theory

Dimensions and issues

- cooperative vs. competitive
- communication
- trust
- recursive modeling
- coalititions
- game theory

Convoy example

Conflicts, Resources

- Omniscience for one agent creates bottleneck

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility
  - Good for global performance?

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility
  - Good for global performance?(invisible hand)

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility
  - Good for global performance?(invisible hand)
  - Pitfall: tragedy of the commons

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility
  - Good for global performance?(invisible hand)
  - Pitfall: tragedy of the commons
  - Pitfall: no stability

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility
  - Good for global performance?(invisible hand)
  - Pitfall: tragedy of the commons
  - Pitfall: no stability
  - Pitfall: lying

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility
  - Good for global performance?(invisible hand)
  - Pitfall: tragedy of the commons
  - Pitfall: no stability
  - Pitfall: lying
- Market-based methods/auctions

Conflicts, Resources

- Omniscience for one agent creates bottleneck
- Self-interested agents: each agent maximizes own local utility
  - Good for global performance?(invisible hand)
  - Pitfall: tragedy of the commons
  - Pitfall: no stability
  - Pitfall: lying
- Market-based methods/auctions
- Negotiation, game theory

Multiagent Planning

- Complex individual agents
- Teamwork modeling
  - Modeling of teammates and opponents
- Recent: emphasis on flexibility in dynamic environments

Communication

- Middle agents (brokers)
- Standard languages
- Ontologies

More next week

Individual Agents

- Purely reactive agents have disadvantages
  - Can't react to nonlocal info or predict effects on global behavior
  - hard to engineer
- Hybrid approach better
- Hard to evaluate agent architecture against one another