Multi-Agent Systems

Overview and Research Directions (Ch. 17.5–17.6)

1

Bookkeeping

- HW 3 out; please start it early (earlier)
  - A note: late penalty will be strict for HW 3-5
  - **Summary: if Blackboard says it's late, it's late**
- Midterm: currently scheduled for 10/18 (Tuesday).
- Would we prefer 10/13 (the preceding Thursday)?
  - We would not. 10/18 it is.

2

Today's Class

- What's an agent?
- Multi-Agent Systems
  - Cooperative multi-agent systems
  - Competitive multi-agent systems
    - Game time!
- MAS Research Directions
  - Organizational structures
  - Communication limitations
  - Learning in multi-agent systems

3

What's an Agent?

4

What's An Agent? (Redux)

- Weiss, p. 29 [after Wooldridge and Jennings]:
  - "An **agent** is a computer system that is **situated** in some **environment**, and that is capable of **autonomous action** in this environment in order to meet its design objectives."
- Russell and Norvig, p. 7:
  - "An agent is just something that perceives and acts."

5

What's An Agent? II

- Ferber, p. 9: "An agent is a physical or virtual entity [which]
  - Is capable of **acting** in an **environment**,
  - Can **communicate** directly with other agents,
  - Is driven by a set of **tendencies**…,
  - Possesses **resources** of its own,
  - Is capable of **perceiving** its environment…,
  - Has only a **partial representation** of this environment…,
  - Possesses **skills** and can offer **services**,
  - May be able to **reproduce** itself,
  - Whose behavior tends towards **satisfying its objectives**, taking account of the resources and skills available to it and depending on its perception, its representations and the communications it receives."

6

OK, What's An Environment?

- Any system that has inputs and outputs is **situated** in an **environment** of sorts
- We've also said **world**
  - **World state**: a snapshot of an environment

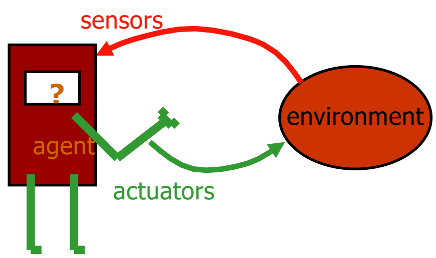

7

What's Autonomy?

- Jennings and Wooldridge, p. 4:
  - "[In contrast with objects] … **agents encapsulate *behavior***, in addition to state.
  - An **object** does not encapsulate behavior: it has no control over the execution of methods – if an object *A* invokes a method *m* on an object *B*, then *B* has no control over whether *m* is executed or not – it just is.
  - In this sense, object *B* is not autonomous, as it has no control over its own actions.
  - Because of this …, we do not think of agents as invoking methods (actions) on agents – rather, we tend to think of them **requesting** actions to be performed."
- Is an if-then-else statement autonomous?

8

So Now What?

- If those definitions aren't useful, is there a useful definition?
- Should we bother trying to create "agents" at all?

9

A Pause to Vote… (more on which later)

- For Tic-Tacs, Skittles, licorice, gummi bears:

- Do you **prefer** Skittles or gummi bears?
- Do you **prefer** licorice or Tic-Tacs?
- Which of these is **best**?
- Sort the candy from **best to worst (1 = best; no ties)**

10

Multi-Agent Systems

11

Multi-Agent Systems

- Jennings et al.'s key properties:
  - Situated [existing in relation to some environment]
  - Autonomous
  - Flexible:
    - Responsive to dynamic environment
    - Pro-active / goal-directed
    - **Social interactions with other agents and humans**
- Research questions: How do we design agents to:
  - **Interact effectively…**
  - …To solve a wide range of problems…
  - …In many different environments?

12

What are Multi-Agent Systems?

- Thus a multiagent system contains a number of agents that:
  - Interact through communication;
  - Are able to act in an environment;
  - Have different "spheres of influence" (which may coincide)
  - And will be linked by other (organizational) relationships

Slide: M. J. Wooldridge & S. Parsons

13

Why Multi-Agent Systems?

- Some of the trends in computing:
  - Ubiquity, interconnection, intelligence, delegation
  - The Internet of Things, self-steering cars, home automation devices
- What advantages does it offer over the alternatives?
- In what circumstances is it useful?

Slide: Matthew E. Taylor, with thanks

14

Possible Motivations for MAS

- Task is too complex for one agent
- Task is inherently distributed
  - Ex. soccer (goalie, striker, defenders…)
- Several resource-bound agents are cheaper (or more feasible to build) or faster than a single, more capable agent
- Multiple agents are more robust because they offer redundancy
  - Ex. one fails, others take its place; **graceful degradation**

Slide adapted from Introduction to AI Robotics, R. Murphy (MIT Press 2000) for 2nd edition

15

Aspects of Multi-Agent Systems

- Cooperative vs. competitive
- Homogeneous vs. heterogeneous
- Interaction protocols and languages
- Organizational structure
- Mechanism design / market economics
- Learning

16

Topics in MAS

- Cooperative MAS:
  - Distributed problem solving: Less autonomy
    - (At least in a certain sense)
  - Distributed planning: Models for cooperation and teamwork
- Competitive or self-interested MAS:
  - Distributed rationality: Voting, auctions
  - Negotiation: Contract nets
  - Strictly adversarial interactions ← least complex

17

Some Cooperative MAS Domains

- Distributed sensor network establishment
- Distributed vehicle monitoring
- Distributed delivery

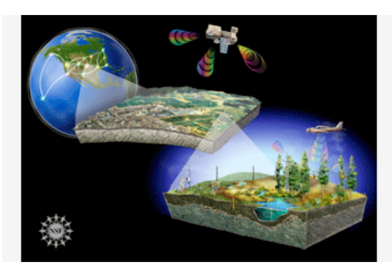

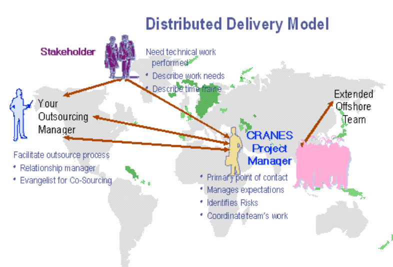

NSF; www.linkedin.com/pulse/3g4g-gps-vehicle-cctv-systems-taxi-bus-truck-kinds-ellies-w; www.cranessoftware.com/alliances/fluid/offshore-dev.php

18

Distributed Sensing & Monitoring

- Distributed sensing:
  - Distributed sensor network establishment:
    - Locate sensors to provide the best coverage
    - Centralized vs. distributed solutions
  - Track vehicle/other movements using multiple sensors
- Distributed vehicle monitoring:
  - Control sensors and integrate results to track vehicles as they move from one sensor's "region" to another's
  - Centralized vs. distributed solutions

19

Distributed Delivery

- Logistics problem: move goods from original locations to destination locations using multiple delivery resources (agents)
- Dynamic, partially accessible, nondeterministic environment (goals, situation, agent status)
- Centralized vs. distributed solution

20

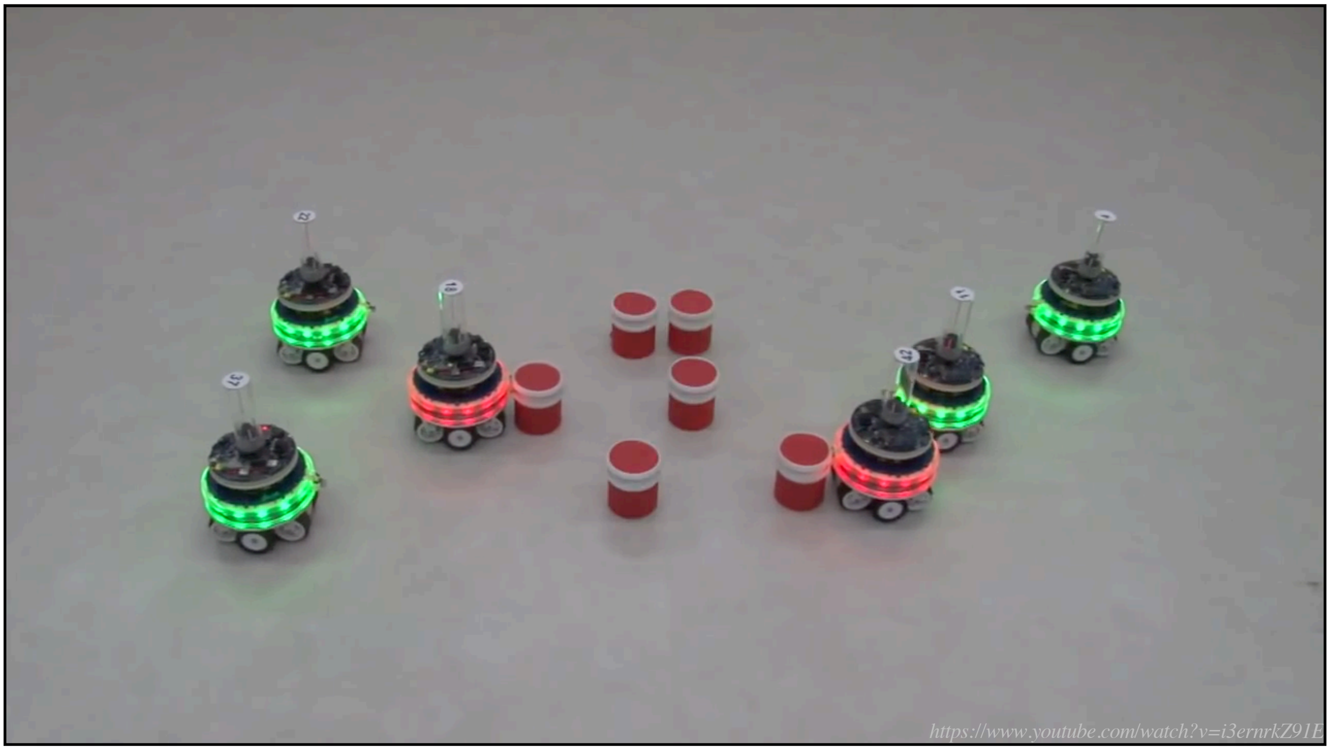

21

Competitive Multi-Agent Systems

22

Games and Game Theory

- Much effort on developing programs for artificial games like chess or poker, played for entertainment
- Larger issue: account for, model, and predict how agents (human or artificial) **interact with other agents**
- **Game theory** accounts for a mixture of cooperative and competitive behavior
- Applies to zero-sum and non-zero-sum games

23

Game Theory

- Defined by von Neumann & Morgenstern
  - von Neumann, J., and Morgenstern, O., (1947).
    The Theory of Games and Economic Behavior.
- Covers both cooperative and non-cooperative situations
- Developed and used in economics, now used to model artificial agents
- Provides a powerful model and practical tools to think about interactions among a set of autonomous agents

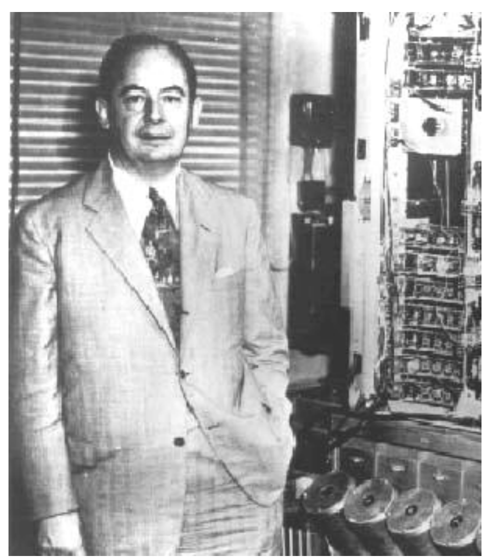

24

Basic Ideas of Game Theory

- **Game theory** studies how strategic interactions among rational players produce outcomes with respect to the players' preferences (utilities)
  - Outcomes might not have been intended
- Game theory offers a general theory of strategic behavior
- Generally depicted in mathematical form
- Plays important role in economics, decision theory and multi-agent systems
- So how do we describe the utility of states across many agents?
  - Social welfare; Pareto optimality; Nash equilibria; other equilibria

25

Social Welfare

- **Social welfare**, or global utility:
  - Sum of all agents' utility
  - If state *s* maximizes social welfare, it is also Pareto-optimal (but not vice versa)
- Somewhat poorly named
  - Sum ≠ average
  - Allocation of resources typically affects influence
    - e.g., you get to take 1 turn per point accrued

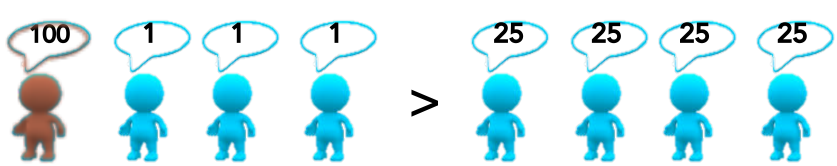

26

Pareto Optimality

- An outcome is **Pareto optimal** if there is **no other outcome that all players would prefer.**
  - "a state … from which it is impossible to [change] so as to make any one individual better off without making at least one individual worse off." – *Wikipedia*
  - There can be multiple Pareto optimal solutions.
- If 2 students on a project both have Bs, would they switch to an A and a C?
  - No – if those are the choices, B/B is Pareto optimal
  - If A/A is an option, B/B is no longer Pareto optimal

27

Pareto Optimality

- *s* is a Pareto-optimal solution iff
  - $\forall s' (\exists x\ U_x(s') > U_x(s) \rightarrow \exists y\ U_y(s') < U_y(s))$
  - I.e., if X is better off in *s'*, then some Y must be worse off
  - There is no other outcome that **all** players would prefer

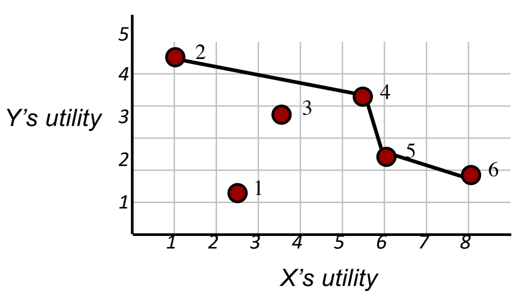

Which solutions are Pareto-optimal?

Which solution(s) maximize global utility (social welfare)?

State 1 = 1.25+2.5 = 3.75
State 2 = 1+4.25 = 5.25
State 3 = 3.5+3.75 = 7.25
State 4 = 5.5+3.25 = 8.75
State 5 = 6+2.25 = 8.25
State 6 = 8+1.75=9.75

28

Nash Equilibrium

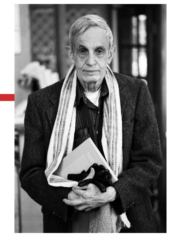

- Occurs when each player's strategy is optimal, **given** strategies of the other players
- No player benefits by **unilaterally** changing strategy while others stay fixed
- Every finite game has at least one Nash equilibrium in either pure or mixed strategies (proved by John Nash)
  - J. F. Nash. 1950. Equilibrium Points in *n*-person Games. Proc. National Academy of Science, 36
  - Nash won 1994 Nobel Prize in economics for this work
  - Read *A Beautiful Mind* by Sylvia Nasar (1998) and/or see the 2001 film

29

Stability

- If an agent can always maximize its own utility with a particular strategy (regardless of other agents' behavior) then that strategy is **dominant**
  - Strategy *s* **dominates** *s′* iff:
  - Outcome (for player *p*) of *s* is better than the outcome of *s′* **in every case**
- A **set** of agent strategies is in **Nash equilibrium** if each agent's strategy S_i is <u>locally optimal</u>, *given* the other agents' strategies
  - No agent has an incentive to change strategies
  - Hence this set of strategies is **locally stable**

30

Prisoner's Dilemma

- Famous example of game theory
- Will two prisoners cooperate to minimize total loss of liberty or will one of them betray the other so as to go free?
- Strategies must be undertaken without full knowledge of what other players will do
- Players adopt dominant strategies, but they don't necessarily lead to the best outcome
- Rational behavior leads to a situation where everyone is worse off

31

Bonnie & Clyde

- Bonnie and Clyde are arrested. They're questioned separately, unable to communicate. They know the deal:
  - If **both** proclaim innocence (deny involvement), they will both get short sentences
  - If **one confesses** and the **other doesn't**, the confessor gets a heavy sentence and the denier gets a light sentence
  - If **both confess**, both get moderate sentences
- What should Bonnie do?
- What should Clyde do?

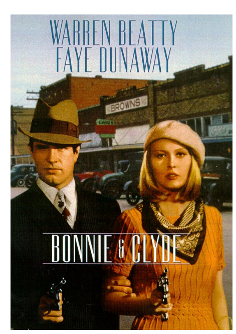

32

Exercise: Prisoner's Dilemma

- <Bonnie's sentence, Clyde's sentence>

B \ *C*	Confesses	Denies
Confesses	(4, 4)	(8, 1)
Denies	(1, 8)	(3, 3)

- Play 1 round – what are results?
- Switch partners
- Play 5 rounds, keeping track of total years

33

Prisoner's Dilemma: Analysis

- Pareto-optimal and social welfare maximizing solution: **Both agents deny**
- Dominant strategy and Nash equilibrium: **Both agents confess**
- **Why?**

B \ *C*	Confesses	Denies
Confesses	(4, 4)	(8, 1)
Denies	(1, 8)	(3, 3)

- Because if you deny and the other confesses, it's the worst case
  - No agent has incentive to change strategies, **given that others do not change**

34

Bonnie's Decision Tree

No wonder Economics is called "the dismal science"

There are two cases to consider:

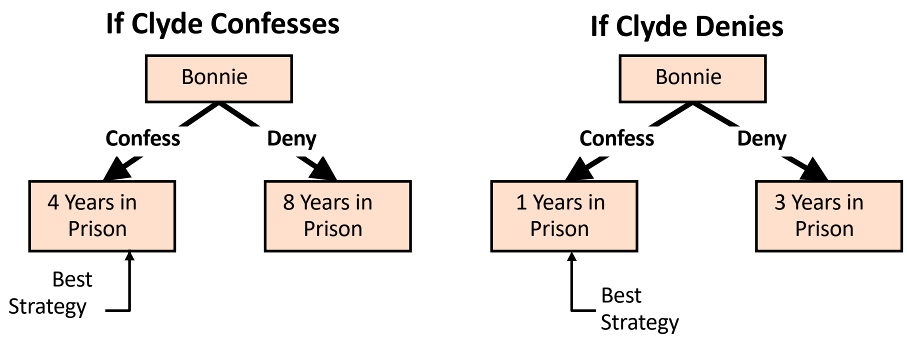

Dominant strategy **for Bonnie** is to confess because no matter what Clyde does she is better off confessing.

35

Iterated Prisoner's Dilemma

- Rational players should always confess in a PD situation
- In real situations, people don't always do this
- Why not? Possible explanations:
  - People aren't rational
  - Morality
  - Social pressure
  - Fear of consequences
  - Evolution of species-favoring genes
- Which make sense? How can we formalize?

36

Iterated PD

- **Key idea**: We often play more than one "game" with someone
- Players have complete knowledge of past games, including their choices **and other players' choices**
- Can choose based on whether they've been cooperative in past
- Simulation was first done by Robert Axelrod (Michigan) where programs played in a round-robin tournament

37

Strategies

- Simple single-agent strategies:
  - Always deny
  - Always cooperate
  - Randomly choose
  - Pavlovian (win-stay, lose-switch): Start by always cooperating, switch to always denying when "punished" by other's denial, switch back and forth at every such punishment
  - Tit-for-tat (TFT): "Be nice, but punish any defections": Starts by cooperating and after that always does what the other player did on the previous round
  - Joss: A sneaky TFT that defects 10% of the time
- In an idealized (noise free) environment, TFT is both a very simple and a very good strategy
- In the IPD tournament, one winner used an initial pattern of choices to identify other players with the same programming, and then switched to all denying

38

Distributed Rationality

How can we encourage/coax/force self-interested agents to play *fairly* in the sandbox?

- **Voting**: Everybody's opinion counts (but how much?)
- **Auctions**: Everybody gets a chance to earn value (but fairly?)
- **Contract nets**: Work goes to the highest bidder
- **Issues**:
  - Global utility
  - Stability
  - Fairness
  - Cheating and lying

39

Voting: It's Not Easy

- How should we **rank** the possible outcomes, given individual agents' preferences (votes)?
- Six desirable properties which can't all be satisfied:
  - Every **combination of votes** should lead to a **ranking**
  - Every **pair of outcomes** should have a **relative ranking**
  - The ranking should be **asymmetric** and **transitive**
  - The ranking should be **Pareto-optimal**
  - **Irrelevant alternatives** shouldn't influence the outcome
  - **Share the wealth**: No agent should always get their way

40

Voting Protocols

- **Plurality voting:**
  - The outcome with the highest number of votes wins
  - Irrelevant alternatives can change the outcome (e.g., third-party candidates "splitting" the vote)
- **Binary voting**:
  - Agents rank sequential pairs of choices ("elimination voting")
  - Irrelevant alternatives can still change the outcome
  - Very order-dependent
- **Borda voting**:
  - Agents' rankings are used as weights, which are summed across all agents
  - Agents can "spend" high rankings on losing choices, making their remaining votes less influential

41

Voting…

- Plurality results:
  - Most #1 votes: ____
- Elimination (bracket) results:

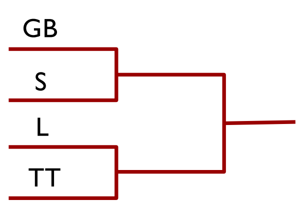

- Borda results:
  - Score(candy) = [4 × (# 1st place votes)] + [3 × (# 2nd)] + [2 × (# 3rd)] + [1 × (# 4th)]
  - TT = ____, S ____, L ____, GB ____

42

Voting game

Discuss… did we achieve global social welfare? Fairness? Were there interesting dynamics?

- Using *plurality (1/0)* voting to select a winner:
  - The winner is the candidate with the most votes
  - The naive strategy is to vote for your top choice – is that best?
- Using *elimination* to select a winner:
  - 1st and 2nd choices can compete, so 3rd or 4th choice comes in 2nd
    - Different people use different strategies – how does that change it?
- Using *Borda (1..k)* voting:
  - Everybody ranks the k candidates that are running in that round
  - Your top choice receives k votes; your second choice, k-1, etc.
  - The winner is the candidate with the most votes
  - Borda voting is often used in combination with a runoff
    - Eliminate the lowest-ranked candidates and try again – how does that change it?

43

Auctions

- Many different types and protocols
- All of the common protocols yield Pareto-optimal outcomes
- **But**… bidders can agree to artificially lower prices in order to cheat the auctioneer
- What about when the colluders cheat each other?
  - (Now that's really not playing nicely in the sandbox!)

44

Learning in MAS

- Emerging field: How can teams of agents learn? Individually? As groups?
- **Genetic algorithms**:
  - Evolve a society of "fittest" agents
  - In practice: a cool idea that is very hard to make work
- **Strategy learning**:
  - In market environments, learn other agents' strategies
- **Distributed Reinforcement Learning** (next slide)

45

MAS RL

- **Distributed Reinforcement Learning**
  - Behave as an individual
  - Receive team feedback
  - Learn to individually contribute to team performance
- How?
  - Iteratively allocate "credit" for group performance to individual decisions

46

Conclusions and Directions

- Different types of "multi-agent systems":
  - Cooperative vs. competitive
  - Heterogeneous vs. homogeneous
- Lots of interesting/open research directions:
  - Effective cooperation strategies
  - "Fair" coordination strategies and protocols
  - Learning in MAS
  - Resource-limited MAS (communication, …)
- Economics: agents are human players with resources

47