Multi-Agent Systems

Overview and Research Directions
AI Class 12 (Ch. 17.5–17.6)

Bookkeeping

- Project teams
  - If you aren't part of a 2-4 person team OR would like additional members, please talk to me after class

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

What's an Agent?

What's An Agent?

- 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."
- Rosenschein and Zlotkin, p. 4:
  - "The more complex the considerations that [a] machine takes into account, the more justified we are in considering our computer an 'agent,' who **acts as our surrogate** in an automated encounter." *[emph. mine]*

What's An Agent? II

- Ferber, p. 9:
  - "An agent is a physical or virtual entity [which]
    a) Is capable of acting in an **environment**,
    b) Can **communicate** directly with other agents,
    c) Is driven by a set of **tendencies**…,
    d) Possesses **resources** of its own,
    e) Is capable of **perceiving** its environment…,
    f) Has only a **partial representation** of this environment…,
    g) Possesses **skills** and can offer **services**,
    h) May be able to **reproduce** itself,
    i) 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."

OK, What's An Environment?

- Isn't any system that has inputs and outputs **situated** in an **environment** of sorts?
  - We've also said world
  - Or world state (a snapshot of an environment)

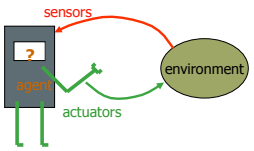

What's Autonomy?

- Jennings and Wooldridge, p. 4:
  - "[In contrast with objects] … **agents as 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?

So Now What?

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

A Pause to Vote… (more on which later)

- For Tic-Tacs, lemon drops, licorice, gummi bears:

- Which of these is **best**?
- Rank each candy on a scale from **1-10**
- Sort the candy from **best to worst**

MULTI-AGENT SYSTEMS

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?

Aspects of MAS

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

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

Some Cooperative MAS Domains

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

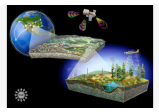

NSF; www.linkedin.com/pulse/3gdg-gps-vehicle-cctv-systems-taxi-bus-truck-kinds-ellies-w; www.cranessoftware.com/alliances/fluid/offshore-dev.php

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

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

Competitive Multi-Agent Systems

Games and Game Theory

- Much effort to develop 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 mixture of cooperative and competitive behavior
- Applies to zero-sum and non-zero-sum games

Basic Ideas

- Game theory studies how strategic interactions among **rational players** produce **outcomes** with respect to the players' **preferences** (or utilities)
  - Outcomes might not have been intended
- Offers a general theory of strategic behavior
- Generally depicted in mathematical form
- Plays important role in economics, decision theory and multi-agent systems

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 (simplified)*
- 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

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
  - "Fair games" remain fair (given optimal play)

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

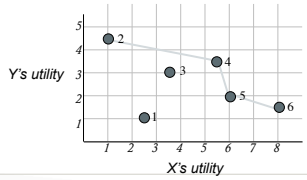

Which solutions are Pareto-optimal?

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

Nash Equilibrium

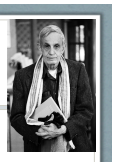

- Occurs when each player's strategy is optimal, **given** strategies of the other players
- No player benefits by **unilaterally** changing strategy while others stay fixed
- Every finite game has at least one Nash equilibrium in either pure or mixed strategies (proved by John Nash)
  - J. F. Nash.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

24

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 locally optimal, given the other agents' strategies
  - No agent has an incentive to change strategies
  - Hence this set of strategies is **locally stable**

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

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 light sentence and the other gets a heavy sentence
  - If **both confess**, both get moderate sentences
- What should Bonnie do?
- What should Clyde do?

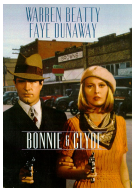

Group Work: Prisoner's Dilemma

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

B \ C	Confesses	Denies
Confesses	(3, 3)	(0, 5)
Denies	(5, 0)	(1, 1)

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

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	(3, 3)	(0, 5)
Denies	(5, 0)	(1, 1)

Bonnie's Decision Tree

There are two cases

No wonder Economics is called "the dismal science"

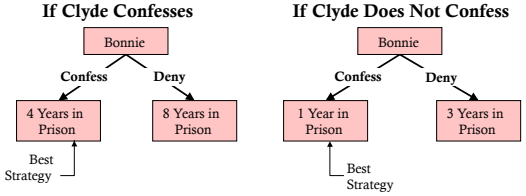

Dominant strategy **for Bonnie** is to confess because no matter what Clyde does she is better off confessing.

Iterated Prisoner's Dilemma

- Rational players should always defect 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?

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
  - (CD=5,CC=3,DD=1,DC=0)
- The simplest program won!

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

Voting

- 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

Voting protocols

- **Plurality voting:**
  - The outcome with the highest number of votes wins
  - Irrelevant alternatives can change the outcome (e.g., Gary Johnson)
- **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
- **Binary voting**:
  - Agents rank sequential pairs of choices ("elimination voting")
  - Irrelevant alternatives can still change the outcome
  - Very order-dependent

Voting…

- For Tic-Tacs, lemon drops, licorice, gummi bears:

- Which of these is **best**?
- Rank each candy on a scale from **1-10**
- Sort the candy from **best to worst**

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 the *range votes* directly to select a winner:
  - Add the range votes
    - Different people use different "widths/ranges" – 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?

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!)

Learning in MAS

- Emerging field: How can teams of agents learn? Individually? As groups?
- **Distributed Reinforcement Learning** (next slide)
- **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

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.

Conclusions and Directions

- Different types of "multi-agent systems":
  - Cooperative vs. competitive
  - Heterogeneous vs. homogeneous
  - Micro vs. macro
- 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