Multi-Agent Systems

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

Material from Marie desJardin

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

3

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, so 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)

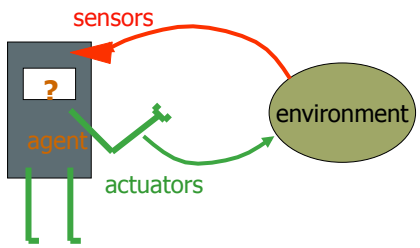

What's autonomy, anyway?

- Jennings and Wooldridge, p. 4:
  - "[In contrast with objects, we] think of **agents as encapsulating *behavior***, in addition to state.
  - An object does not encapsulate behavior: it has no control over the execution of methods – if an object *x* invokes a method *m* on an object *y*, then *y* has no control over whether *m* is executed or not – it just *is*.
  - In this sense, object *y* 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 **1-5**
- Sort the candy from **best to worst**
- Fill this out:

Multi-Agent Systems

Multi-agent systems

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

Aspects of multi-agent systems

- Cooperative vs. competitive
- Homogeneous vs. heterogeneous
- Macro vs. micro

- Interaction protocols and languages
- Organizational structure
- Mechanism design / market economics
- Learning

Topics in multi-agent systems

- 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

Typical cooperative MAS domains

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

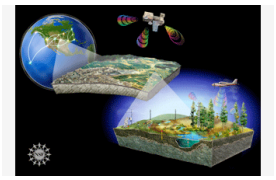

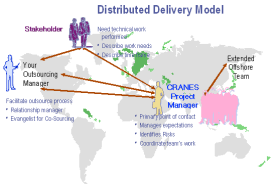

NSF; www.linkedin.com/pulse/3g4g-gps-vehicle-cctv-systems-taxi-bus-truck-kinds-ellies-w; www.cranessoftware.com/alliances/fluid/offshore-dev.php

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

Distributed rationality

- **Techniques to 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
  - Fairness
  - Stability
  - Cheating and lying

Pareto optimality

- An outcome is **Pareto optimal** if **there is no other outcome that all players would prefer**.
- 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, or global utility:
  - Sum of all agents' utility
  - If S maximizes social welfare, it is also Pareto-optimal (but not vice versa)

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

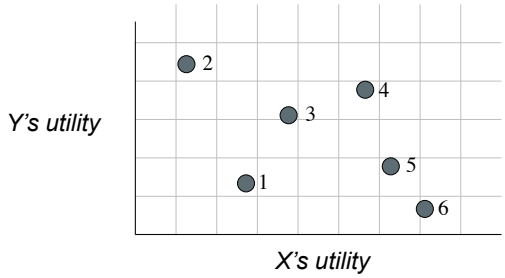

Which solutions are Pareto-optimal?

Which solutions maximize global utility (social welfare)?

Stability

- If an agent can always maximize its utility with a particular strategy (regardless of other agents' behavior) then that strategy is **dominant**
  - Strategy *s* dominates *s'* if **outcome for *s* is better than the outcome for *s'* in every case**
  - (for player *p*)
- 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

A \ B	Cooperate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

Let's play!

Prisoner's Dilemma: Analysis

- Pareto-optimal and social welfare maximizing solution: Both agents cooperate
- Dominant strategy and Nash equilibrium: Both agents defect

A \ B	Cooperate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

- Why?

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

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 **1-5**
- Sort the candy from **best to worst**
- Fill this out:

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 to investigate how teams of agents can learn individually and as groups.**
- **Distributed reinforcement learning**: Behave as an individual, receive team feedback, and learn to individually contribute to team performance.
- **Distributed reinforcement learning**: Iteratively allocate "credit" for group performance to individual decisions.
- **Genetic algorithms**: Evolve a society of "fittest" agents.
  - In practice: a very cool idea that is very hard to make work.
- **Strategy learning**: In market environments, learn other agents' strategies.

Conclusions and directions

- "Agent" means many different things
- 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, …)