Agent-Based Systems

Michael Rovatsos
mrovatso@inf.ed.ac.uk

Lecture 9 – Social Choice

Where are we?

Last time . . .

- Discussed simple, abstract models of multiagent encounters
- Utilities, preferences and outcomes
- Game-theoretic models and solution concepts
- Examples: Prisoner's Dilemma, Coordination Game
- Axelrod's tournament its conclusions and critique

Today . . .

- **Social Choice**

Making Group Decisions

- Previously we looked at agents acting strategically
- Outcome in normal-form games follows immediately from agents' choices
- Often a mechanism for deriving group decision is present
- What rules are appropriate to determine the joint decision given individual choices?
- **Social Choice Theory** is concerned with group decision making (basically analysis of mechanisms for voting)
- Basic setting:
  - Agents have preferences over outcomes
  - Agents vote to bring about their most preferred outcome

Preference Aggregation

- Setting:
  - $Ag = \{1, \ldots, n\}$ **voters** (finite, odd number)
  - $\Omega = \{\omega_1, \omega_2, \ldots\}$ possible **outcomes** or **candidates**
  - $\varpi_i \in \Pi(\Omega)$, preference ordering for agent i (e.g. $\omega \succ_i \omega'$)
- Preference Aggregation:

  How do we combine a collection of potentially different preference orders in order to derive a group decision?

- Voting Procedures:
  - **Social Welfare Function**: $f : \Pi(\Omega) \times \ldots \times \Pi(\Omega) \to \Pi(\Omega)$
  - **Social Choice Function**: $f : \Pi(\Omega) \times \ldots \times \Pi(\Omega) \to \Omega$
- Task is either to derive a globally acceptable preference ordering, or determine a winner

Plurality

- Voters submit preference orders
- The outcome that appears first in most preference orders wins
- Only submission of the highest-ranked candidate is required
- **Simple majority voting** when $|\Omega| = 2$
- Advantages: simple to implement and easy to understand
- Problems:
    - Tactical voting
    - Strategic manipulation
    - Condorcet's paradox

UK Politics Example

- Outcomes: $\Omega = \{\omega_L, \omega_D, \omega_C\}$, where ω_L represents the Labour Party, ω_D the Liberal Democrats and ω_C the Conservative Party
- Voters:
  - 43% of $|Ag|$ are left-wing voters: $\omega_L \succ \omega_D \succ \omega_C$
  - 12% of $|Ag|$ are centre-left voters: $\omega_D \succ \omega_L \succ \omega_C$
  - 45% of $|Ag|$ are right-wing voters: $\omega_C \succ \omega_D \succ \omega_L$
- ω_C wins with 45%

Anomalies with Plurality

- Despite not securing majority, ω_C wins with 45%
- Even worse: ω_C is the **least preferred option** for 55% of voters
- **Tactical Voting**:

  Centre-left candidates may do better by voting for ω_L instead of their actual preference

- **Strategic manipulation**: misrepresenting your preferences to bring about a more preferred outcome
- But is lying bad? Not in principle, but it favours computationally stronger voters, and wastes computational resources

Condorcet's Paradox

- Outcomes: $\Omega = \{\omega_1, \omega_2, \omega_3\}$
- Voters: $Ag = 1, 2, 3$ with preference orders
  - $\omega_1 \succ_1 \omega_2 \succ_1 \omega_3$
  - $\omega_3 \succ_2 \omega_1 \succ_2 \omega_2$
  - $\omega_2 \succ_3 \omega_3 \succ_3 \omega_1$
- With plurality voting, we obtain a tie
- For every candidate, $\frac{2}{3}$ of the voters prefers another outcome
- **Condorcet's Paradox**:

  There are scenarios in which no matter which outcome we choose the majority of voters will be unhappy with the outcome chosen

Sequential Majority Elections

- Instead of one-step protocol, voting can be done in several steps
- Candidates face each other in **pairwise elections**, the winner progresses to the next election
- **Election agenda** is the ordering of these elections (e.g. $\omega_2, \omega_3, \omega_4, \omega_1$)
- Can be organised as a **binary voting tree**

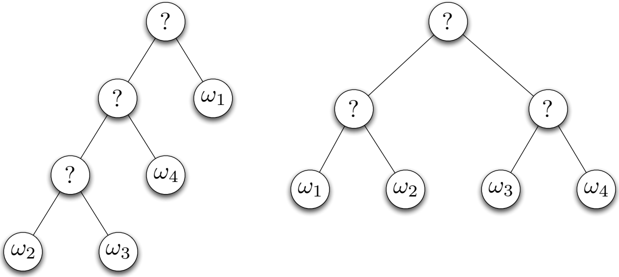

- Key Problem: The final outcome depends on the election agenda

Majority Graphs (I)

- Need to introduce better tools for discussing sequential voting
- A **majority graph** is a succinct representation of voter preferences
- Nodes correspond to outcomes, e.g. $\omega_1, \omega_2, \ldots$
- There is an edge from ω to ω' if a majority of voters rank ω above ω'

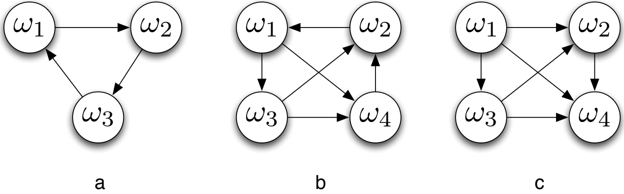

Majority Graphs (II)

- **Tournament:** complete, assymetric and irreflexible majority graph (produced with odd number of voters)
- **Possible winner:** There is an agenda that leads the outcome to win
  - Every outcome in graphs *a* and *b*
- **Condorcet winner:** overall winner for every possible agenda
  - Outcome ω_1 in graph *c*
- **Strategic manipulation:** fixing the election agenda

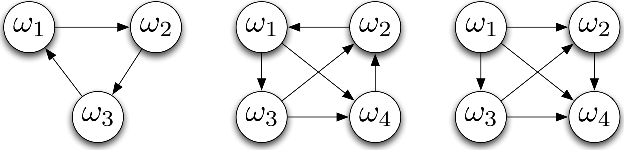

a b c

The Borda Count

- In simple mechanisms above, only top-ranked candidate taken into account, rest of orderings disregarded
- **Borda count** looks at entire preference ordering, counts the strength of opinion in favour of a candidate
- For all preference orders and outcomes ($|\Omega = k|$)

  if ω_i is lth in a preference ordering, increment its strength by $k - l$
- Politics example:
  - 43 of $|Ag|$ are left-wing voters: $\omega_L \succ \omega_D \succ \omega_C$
  - 12 of $|Ag|$ are centre-left voters: $\omega_D \succ \omega_L \succ \omega_C$
  - 45 of $|Ag|$ are right-wing voters: $\omega_C \succ \omega_D \succ \omega_L$

$$\omega_L: \; 43 * (3-1) + 12 * (3-2) + 45 * (3-3) = 86 + 12 = 98$$
$$\omega_D: \; 43 * (3-2) + 12 * (3-1) + 45 * (3-2) = 43 + 24 + 45 = 112$$
$$\omega_C: \; 43 * (3-3) + 12 * (3-3) + 45 * (3-1) = 90$$

The Slater Ranking

- Idea: how can we minimise disagreement between the majority graph and the social choice?
- For each possible ordering measure the degree of disagreement with the majority graph
- Degree of disagreement = edges that need to be flipped (NP-hard to compute)
- Example:

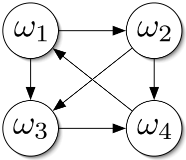

Consider $\omega_1 \succ^* \omega_2 \succ^* \omega_4 \succ^* \omega_3$

cost is 2, we have to flip the edges (ω_3, ω_4) and (ω_4, ω_1)

Consider $\omega_1 \succ^* \omega_2 \succ^* \omega_3 \succ^* \omega_4$

cost is 1, we have to flip the edge (ω_4, ω_1)

this is the ordering with the lowest disagreement

Desirable Properties (I)

- **Pareto Condition**
  - If every voter ranks ω_i above ω_j then $\omega_i \succ^* \omega_j$
  - Satisfied by plurality and Borda, but not by sequential majority
- **Condorcet winner condition**
  - The outcome would beat every other outcome in a pairwise election
  - Satisfied only by sequential majority elections

Desirable Properties (II)

- **Independence of irrelevant alternatives (IIA)**
  - The social ranking of two outcomes ω_i and ω_j should exclusively depend on their relevant ordering in the preference orders
  - Plurality, Borda and sequential majority elections do not satisfy IIA
- **Non-Dictatorship**
  - A social welfare function f is a **dictatorship** if for some voter i

$$f(\varpi_1, \ldots, \varpi_n) = \varpi_i$$

  - Dictatorships satisfy Pareto condition and IIA

Arrow's Theorem

- Overall vision in social choice theory: identify "good" social choice procedures
- Unfortunately, a fundamental theoretical result gets in the way
- **Arrow's Theorem:**

  For elections with more than two outcomes, the only voting procedures that satisfy the Pareto condition and IIA are dictatorships

- Disappointing, basically means we can never achieve combination of good properties without dictatorship

Strategic Manipulation

- As stated above, while lying could be allowed as part of rational behaviour, it is unfair and wasteful
- Can we engineer voting procedures immune to manipulation?
- A social choice function f is **manipulable** if, for a collection of preference profiles there exists ϖ'_i such that

$$f(\varpi_1, \ldots, \varpi'_i, \varpi_n) \succ_i f(\varpi_1, \ldots, \varpi_i, \varpi_n)$$

- **Gibbard-Satterthwaite Theorem:**

  For elections with more than two outcomes, the only non-manipulable voting method satisfying the Pareto property is a dictatorship

Complexity of Manipulation

- So we have another negative result: strategic manipulation is possible in principle in all desirable mechanisms
- But how easy is it to manipulate effectively?
- Distinction between being **easy to compute** and **easy to manipulate**
- Mechanisms can be designed for which manipulation is very computationally complex (but often only in the worst case)
- Are there non-dictactorial voting procedures that are easy to compute but not easy to manipulate?
- Yes, for example **second-order Copeland**

Summary

- Discussed procedures for making group decisions
- Plurality, Sequential Majority Elections, Borda Count, Slater Ranking
- Desirable properties
- Dictatorships
- Strategic manipulation and computational complexity
- Next time: **Coalition Formation**