Computational Models of Group Decision Making

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Plan for this Talk

My interest is in *group decision making*, studied from the perspectives of computer science, logic, and economics.

To focus on phenomena at the group level, we use very *simple models* of individual decision makers. (Q: Is this still relevant to *cognition*?)

Today I will present three such models for group decision making and highlight one result for each one of them:

- *negotiation:* selfish behaviour *vs.* social welfare
- *voting:* effects of iterated strategic manipulation
- *judgment aggregation:* complexity of paradox avoidance

Model 1: Allocating Goods to Agents

Allocating goods to agents who value them:

- set of *agents* $N = \{1, \ldots, n\}$
- set of (indivisible) *goods* $G = \{g_1, \ldots, g_m\}$
- each agent $i \in N$ has a *utility function* $u_i : \mathcal{P}(G) \to \mathbb{R}$
- an *allocation* is a function $A : N \to \mathcal{P}(G)$ with $A(i) \cap A(j) = \emptyset$

We want to find the *best* allocation:

- find A maximising *utilitarian social welfare* $\sum_{i \in N} u_i(A(i))$

Not easy: ... NP-hard ... we can't order agents what to do ... etc.

Negotiation Protocol and Agent Behaviour

Consider this (very liberal!) *negotiation protocol:*

- groups can agree on arbitrary *deals* (exchanges of goods)
- deals can be coupled with *side-payments* (adding up to 0)

And make this (rather simplistic!) assumption on *agent behaviour:*

- agents are *myopic* and only compare current and next allocation
- agents are *individually rational:* accept a deal *iff* utility gain outweighs monetary loss (or: monetary gain outweighs utility loss)

Can we make any predictions about the quality of allocations reached?

Convergence Theorem

We can do much better than you might expect!

As first noted by Sandholm (1998):

> *Any sequence of individually rational deals will converge to an allocation with maximal utilitarian social welfare.*

T. Sandholm. Contract Types for Satisficing Task Allocation: I Theoretical Results. Proc. AAAI Spring Symposium 1998.

Proof and Discussion

Why is this true?

> *Any sequence of individually rational deals will converge to an allocation with maximal utilitarian social welfare.*

Main insights:

- individually rational deal = deal increasing in social welfare
- *finite* space of possible allocations

Very nice result, but there are issues:

- you might need *many* and/or *complex deals* (it's still NP-hard)
- is *utilitarian social welfare* really what we want?

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of Artif. Intelligence Research*, 25:315–348, 2006.

Model 2: Voting with Ordinal Preferences

How should n *voters* choose from a set of m *alternatives*?

Here are three *voting rules* (there are many more):

- *Plurality:* elect the alternative ranked first most often
- *Plurality with runoff:* run a plurality election and retain the two front-runners; then run a majority contest between them
- *Copeland:* elect the alternative winning the most pairwise majority contests (awarding half a point for each draw)

Example: Choosing a Beverage for Lunch

Consider this election with nine *voters* having to choose from three *alternatives* (namely what beverage to order for a common lunch):

2 Germans:	Beer	$\succ$	Wine	$\succ$	Milk
3 Frenchmen:	Wine	$\succ$	Beer	$\succ$	Milk
4 Dutchmen:	Milk	$\succ$	Beer	$\succ$	Wine

Which beverage *wins* the election for

- the plurality rule?
- plurality with runoff?
- the Copeland rule?

So which rules are *Condorcet-consistent* (i.e., which rules guarantee to elect the winner of all pairwise majority contests when there is one)?

Iterated Voting

Suppose voters *update their ballots* again and again, after observing the election outcome. Suppose they do so by moving their favourite amongst the k front-runners to the top position ("k*-pragmatism*").

Example for a nice result:

> *If all voters are* k*-pragmatists, then Condorcet consistency is preserved under iteration.*

Also: we have observed *improved Condorcet efficiency* under iteration, for rules that are not Condorcet consistent.

A. Reijngoud and U. Endriss. Voter Response to Iterated Poll Information. Proc. 11th Internat. Conf. on Auton. Agents and Multiagent Systems (AAMAS-2012).

Model 3: Judgment Aggregation

Three judges hold different views on the propositions p, q, and $p \to q$ (e.g., p might stand for "the temperature is below 17°C" and q for "we should switch on the heating").

	p	$p \to q$	q
Judge 1:	Yes	Yes	Yes
Judge 2:	Yes	No	No
Judge 3:	No	Yes	No

Avoiding Paradoxes

Paradox: Each individual judge provides a logically consistent opinion, yet the majority outcome is logically inconsistent.

Agenda = set of formulas to vote on (closed under complementation)

Nice (but daunting) characterisation result:

> *An agenda Φ is safe (guarantees consistent outcomes) for the majority rule iff Φ has no minimally inconsistent set of size $\geqslant 3$. Checking whether this is the case is coNP^{NP}-complete.*

K. Nehring and C. Puppe. The Structure of Strategy-proof Social Choice. Part I: General Characterization and Possibility Results on Median Space. *Journal of Economic Theory*, 135(1):269–305, 2007.

U. Endriss, U. Grandi, and D. Porello. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research*, 45:481–514, 2012.

Last Slide

We saw three examples for *simple models* of *group decision making*:

- *negotiation* (sample result: convergence to social optimum)
- *voting* (sample result: iterated voting can improve outcomes)
- *judgment aggregation* (sample result: paradoxes hard to avoid)

Simple models are useful to clearly bring out the 'group effect', but questions regarding their adequacy are justified ...

For more information on the wider field, *computational social choice*, have a look at our introductory chapter on the topic (on my website).

F. Brandt, V. Conitzer, and U. Endriss. Computational Social Choice. In G. Weiss (ed.), *Multiagent Systems*. MIT Press, 2013.