

Social Choice Theory as a Foundation for Multiagent Systems

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Outline

Why *social choice theory* as a foundation for *multiagent systems*?

Outline of the argument:

- MAS = group of agents who coordinate, cooperate, compete.
- Thus: it's all about group decision making.
- The classical discipline studying group decision making is SCT.

Outline of the talk:

- Resource allocation
- Voting and elections
- Judgment aggregation

For each of them: examples, basic concepts, and a theorem

Resource Allocation

Many of the *applications* studied in MAS are about *resource allocation*.

And even when not, agents tend to first have to agree on an allocation of the resources available, before tackling the problem at hand.

Formal Model

Allocating goods to agents who value them:

- set of *agents* $N = \{1, \dots, n\}$
- set of (indivisible) *goods* $G = \{g_1, \dots, g_m\}$
- each agent $i \in N$ has a *utility function* $u_i : 2^G \rightarrow \mathbb{R}$
- an *allocation* is a function $A : N \rightarrow 2^G$ 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.

Notions of Social Optimality

When is an allocation socially optimal? Social choice theory and welfare economics offer many possible answers:

- *Utilitarian social welfare*: maximal sum of utilities
- *Egalitarian social welfare*: maximal minimum of utilities
- *Nash product*: maximal product of utilities
- *Pareto optimality*: cannot improve utility of any agent without decreasing utility of some other agent
- *Lorenz optimality*: cannot improve sum of utilities for k poorest agents without decreasing sum of utilities for k' poorest agents
- *Envy-freeness*: no agent wants to swap with any other agent

MAS designers need to make right choice for application at hand!

Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J.A. Rodríguez-Aguilar and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.

Another Convergence Theorem

Under stronger assumptions, we can get similar convergence theorems for more interesting notions of social optimality. Example:

Theorem 1 (Chevaleyre et al., 2007) *If utilities are **submodular** and agents **start equally well off**, then any sequence of **individually rational** deals with **uniform payments** will result in an **envy-free** allocation, even on a **social network** (constraining both deals and envy).*

Y. Chevaleyre, U. Endriss and N. Maudet. Allocating Goods on a Graph to Eliminate Envy. Proc. AAI-2007.

Voting

Resource allocation is a very specific social choice problem with lots of *internal structure* and, typically, preferences modelled as *utility functions*.

Sometimes we just want to choose an *alternative* (not an *allocation*) and we only want to commit to *ordinal preferences*.

Then the right model to work with is classical voting theory.

Three Voting Rules

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 (i.e., each voter assigns 1 point to an alternative of her choice, and the alternative receiving the most points wins)
- *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?

Condorcet Consistency

A desirable property of a voting rule is *Condorcet consistency*: elect the winner of all pairwise majority contests whenever there is one.

Example: beer is a *Condorcet winner* for the profile below.

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

Of the rules we have seen, only Copeland is Condorcet-consistent.

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*”).

What can we say about the meta voting rule thus obtained? Example:

Theorem 2 (Reijngoud and Endriss, 2012) *If all voters are k -pragmatists, then Condorcet consistency is preserved under iteration.*

A. Reijngoud and U. Endriss. Voter Response to Iterated Poll Information. Proc. AAMAS-2012.

Judgment Aggregation

Preferences are not the only type of information we may wish to aggregate within a multiagent system.

Example

Three agents hold different views on the truth of the propositions p , q , and $p \rightarrow q$ (e.g., p might stand for “the temperature is below 16°C ” and q for “we should switch off the air conditioning”).

	p	$p \rightarrow q$	q
Agent 1:	Yes	Yes	Yes
Agent 2:	Yes	No	No
Agent 3:	No	Yes	No

What should be the *collective decision* of the group?

Safety of the Agenda

As we have seen, judgment aggregation can lead to *paradoxes*.

Suppose we don't know what *aggregation rule* our agents use, but we do know some of its *properties*. When can we be certain there won't be any paradox for a given *agenda* (set of formulas to be judged)?

Theorem 3 (Endriss et al., 2012) *An agenda Φ is safe for all anonymous, unanimous, independent, complete, and complement-free aggregation rules iff every inconsistent subset of Φ has a subset of the form $\{\varphi, \neg\varphi\}$. Deciding this is coNP^{NP} -complete.*

Thus: only simplistic agendas are safe, yet checking safety is hard!

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

Last Slide

We have seen examples, basic concepts, and one theorem each for:

- Resource allocation
- Voting and elections
- Judgment aggregation

My claim is that these fundamental ideas are helpful in designing and analysing multiagent systems.

For more information on *computational social choice*, have a look at our introductory chapter on the topic.

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