Fair Allocation of Indivisible Goods

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

Collective Decision Making

How should we aggregate the views of several agents to help them take a collective decision? Examples:

- *voting:* e.g., for candidates in political elections
- *fair allocation of goods:* e.g., computing-resources to users
- *two-sided matching:* e.g., junior doctors to hospitals
- *judgment aggregation:* e.g., regarding annotated data in linguistics

This is *social choice theory*, traditionally studied in economics and political science, but now also by "us": *computational social choice*.

Plan for this talk:

- a few remarks about *computational social choice* in general
- examples for research questions regarding *fair allocation problems*

Social Choice and Computer Science (1)

Social choice theory has natural applications in computing:

- *Multiagent Systems:* to aggregate the beliefs + to coordinate the actions of groups of autonomous software agents
- *Search Engines:* to determine the most important sites based on links ("votes") + to aggregate the output of several search engines
- *Recommender Systems:* to recommend a product to a user based on earlier ratings by other users

But not all of the classical assumptions will fit these new applications. So we need to develop *new models* and *ask new questions*.

Social Choice and Computer Science (2)

Vice versa, computational techniques are useful for advancing the state of the art in social choice:

- *Algorithms and Complexity:* to develop algorithms for (complex) voting procedures + to understand the hardness of "using" them
- *Knowledge Representation:* to compactly represent the preferences of individual agents over large spaces of alternatives
- *Logic and Automated Reasoning:* to formally model problems in social choice + to automatically verify (or discover) theorems

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia (eds.), *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Fair Allocation of Indivisible Goods

Consider a set of agents and a set of goods. Each agent has her own preferences regarding the allocation of goods to agents. Examples:

- allocation of resources amongst members of our society
- allocation of bandwith to processes in a communication network
- allocation of compute-time to scientists on a super-computer
- ...

We will focus on *indivisible* objects (as opposed to *divisible* "cakes").

U. Endriss. *Lecture Notes on Fair Division*. ILLC, University of Amsterdam, 2009.

S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Notation and Terminology

Notation and terminology:

- Set of *agents* $\mathcal{N} = \{1, \ldots, n\}$ and finite set of *objects* $\mathcal{O}$.
- An *allocation* A is a partitioning of $\mathcal{O}$ amongst the agents in $\mathcal{N}$.
  Example: $A(i) = \{a, b\}$ — agent i owns items a and b
- Each agent $i \in \mathcal{N}$ has got a *utility function* $u_i : 2^{\mathcal{O}} \to \mathbb{R}$.
  Example: $u_i(A) = u_i(A(i)) = 577.8$ — agent i is pretty happy

How can we find a socially optimal allocation of objects?

Social Objectives

There are many possible definitions for social optimality:

- *Pareto optimality:* no (weak) improvements for all possible
- maximal *utilitarian social welfare:* $\sum_{i \in \mathcal{N}} u_i(A(i))$
- maximal *egalitarian social welfare:* $\min_{i \in \mathcal{N}} u_i(A(i))$
- maximal *Nash social welfare:* $\prod_{i \in \mathcal{N}} u_i(A(i))$
- *equitability:* $u_i(A(i)) = u_j(A(j))$ for all $i, j \in \mathcal{N}$
- minimal *inequality*, e.g., in terms of the *Gini index*
- *proportionality:* $u_i(A(i)) \geqslant \frac{1}{n} \cdot \max_{S \subseteq \mathcal{O}} u_i(S)$
- *envy-freeness:* $u_i(A(i)) \geqslant u_i(A(j))$ for all $i, j \in \mathcal{N}$
- minimal *degree of envy* for some way of aggregating envy pairs
- and more

How to pick the right objective is a major concern of classical SCT ("*axiomatic method*"). CS applications suggest new perspectives.

Preference Representation

Example: Allocating 10 goods to 5 agents means $5^{10} = 9765625$ allocations and $2^{10} = 1024$ bundles for each agent to think about.

So we need to choose a good *language* to compactly represent preferences over such large numbers of alternative bundles, e.g.:

- Logic-based languages (weighted goals)
- Bidding languages for combinatorial auctions (OR/XOR)
- Program-based preference representation (straight-line programs)
- CP-nets and CI-nets (for ordinal preferences)

The choice of language affects both *algorithm design* and *complexity*.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 29(4):37–46, 2008.

Computing Socially Optimal Allocations

Suppose all agents have sent us their preferences, expressed in a suitable representation language, and we have picked a social objective.

How can we compute the social optimum? Specifically:

- What are useful *algorithms*? [not today]
- What is the *computational complexity* of this problem?
- How much easier does it get for *restricted preferences*?
- Can we *distribute* computation over the agents ("*negotiation*")?
- How do we deal with *strategic behaviour*? [not today]

Welfare Optimisation

How hard is it to find an allocation with maximal social welfare?
Rephrase this *optimisation problem* as a *decision problem:*

WELFARE OPTIMISATION (WO)
Instance: agents with utility functions over goods, and $K \in \mathbb{Q}$
Question: Is there an allocation A such that $\mathrm{SW}_{\mathrm{util}}(A) > K$?

Unfortunately, the problem is intractable:

Theorem 1 WELFARE OPTIMISATION *is NP-complete, even when every agent assigns nonzero utility to just a single bundle.*

Proof sketch: Language not important (single-bundle assumption).
In NP: routine. NP-hardness: reduction from SET PACKING. ✓

This seems to have first been stated by Rothkopf et al. (1998).

M.H. Rothkopf, A. Pekeč, and R.M. Harstad. Computationally Manageable Combinational Auctions. *Management Science*, 44(8):1131–1147, 1998.

Welfare Optimisation under Additive Preferences

Sometimes we can reduce complexity by restricting attention to problems with certain types of preferences.

A utility function $u : 2^{\mathcal{O}} \to \mathbb{R}$ is called *additive* if for all $S \subseteq \mathcal{O}$:

$$u(S) \quad = \quad \sum_{x \in S} u(\{x\})$$

For this restriction, we get a positive result:

Proposition 2 Welfare Optimisation *is in P in case all individual utility functions are additive.*

Exercise: Why is this true?

Distributed Approach

Instead of computing a socially optimal allocation in a centralised manner, we now want agents to negotiate amongst themselves.

- We are given some *initial allocation* A_0.
- A *deal* $\delta = (A, A')$ is a pair of allocations (before/after).
- A deal may come with a number of *side payments* to compensate some of the agents for a loss in utility. A *payment function* is a function $p : \mathcal{N} \to \mathbb{R}$ with $p(1) + \cdots + p(n) = 0$.

  Example: $p(i) = 5$ and $p(j) = -5$ means that agent i *pays* €5, while agent j *receives* €5.

Negotiating Socially Optimal Allocations

The main question of interest concerns the relationship between:

- the *local view:* what deals are agents willing to make?
- the *global view:* what allocations do we consider socially optimal?

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

The Local/Individual Perspective

A rational agent (who does not plan ahead) will only accept deals that improve her individual welfare:

- A deal $\delta = (A, A')$ is called *individually rational* (IR) if there *exists* a payment function p such that $u_i(A') - u_i(A) > p(i)$ for all $i \in \mathcal{N}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.

That is, an agent will only accept a deal if it results in a gain in utility (or money) that strictly outweighs a possible loss in money (or utility).

The Global/Social Perspective

As system designers, we are interested in *utilitarian social welfare:*

$$\mathrm{SW}_{\mathrm{util}}(A) \quad = \quad \sum_{i \in \mathcal{N}} u_i(A(i))$$

While the local perspective is driving the negotiation process, we use the global perspective to assess how well we are doing.

Exercise: How well/badly do you expect this to work?

Example

Let $\mathcal{N} = \{ann, bob\}$ and $\mathcal{O} = \{chair, table\}$ and suppose our agents use the following utility functions:

$$
\begin{array}{rclrcl}
u_{ann}(\emptyset) & = & 0 & u_{bob}(\emptyset) & = & 0 \\
u_{ann}(\{chair\}) & = & 2 & u_{bob}(\{chair\}) & = & 3 \\
u_{ann}(\{table\}) & = & 3 & u_{bob}(\{table\}) & = & 3 \\
u_{ann}(\{chair, table\}) & = & 7 & u_{bob}(\{chair, table\}) & = & 8
\end{array}
$$

Furthermore, suppose the initial allocation of objects is A_0 with $A_0(ann) = \{chair, table\}$ and $A_0(bob) = \emptyset$.

Social welfare for allocation A_0 is 7, but it could be 8. By moving only a *single* item from agent ann to agent bob, the former would lose more than the latter would gain (not individually rational).

The only possible deal would be to move the whole *set* $\{chair, table\}$.

Convergence

The good news:

Theorem 3 (Sandholm, 1998) *Any sequence of IR deals will eventually result in an allocation with maximal social welfare.*

Discussion: Agents can act *locally* and need not be aware of the global picture (convergence is guaranteed by the theorem).

Discussion: Other results show that (a) arbitrarily complex deals might be needed and (b) paths may be exponentially long. Still NP-hard!

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

So why does this work?

The key to the proof is the insight that IR deals are exactly those deals that increase social welfare:

- **Lemma 4** *A deal $\delta = (A, A')$ is individually rational if and only if $\mathrm{SW}_{\mathrm{util}}(A) < \mathrm{SW}_{\mathrm{util}}(A')$.*

  Proof: ($\Rightarrow$) Rationality means that overall utility gains outweigh overall payments (which are $= 0$).

  ($\Leftarrow$) The social surplus can be divided amongst all agents by using, say, the following payment function:

$$p(i) \;=\; u_i(A') \;-\; u_i(A) \;-\; \underbrace{\frac{\mathrm{SW}_{\mathrm{util}}(A') - \mathrm{SW}_{\mathrm{util}}(A)}{|\mathcal{N}|}}_{>0} \qquad \checkmark$$

Thus, as SW increases with every deal, negotiation must *terminate*. Upon termination, the final allocation A must be *optimal*, because if there were a better allocation A', the deal $\delta = (A, A')$ would be IR.

Related Work

Many ways in which this can be (and has been) taken further:

- other social objectives? / other local criteria?
- what types of deals needed for what utility functions?
- path length to convergence?
- other types of goods: sharable, nonstatic, ... ?
- negotiation on a social network?

For several combinations of the above there still are open problems.

Last Slide

We have seen that finding a fair/efficient allocation of indivisible goods to agents gives rise to a combinatorial optimisation problem.

Two approaches:

- *Centralised:* Give a complete specification of the problem to an optimisation algorithm. Often *intractable*.
- *Distributed:* Try to get the agents to solve the problem. For certain fairness criteria and certain assumptions on agent behaviour, we can predict *convergence* to an optimal state.

All of this is part of *computational social choice*, which is also studying other types of collective decision making scenarios, using methods that include game theory, logic, knowledge representation, statistics, algorithms, complexity theory, . . .