Collective Decision Making in Combinatorial Domains

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

Talk Outline

Introduction to *computational social choice*, with some examples:

- *logical modelling* of social choice problems
- *computational complexity* of strategic behaviour in elections
- choosing from huge numbers of alternatives (*combinatorial domains*)

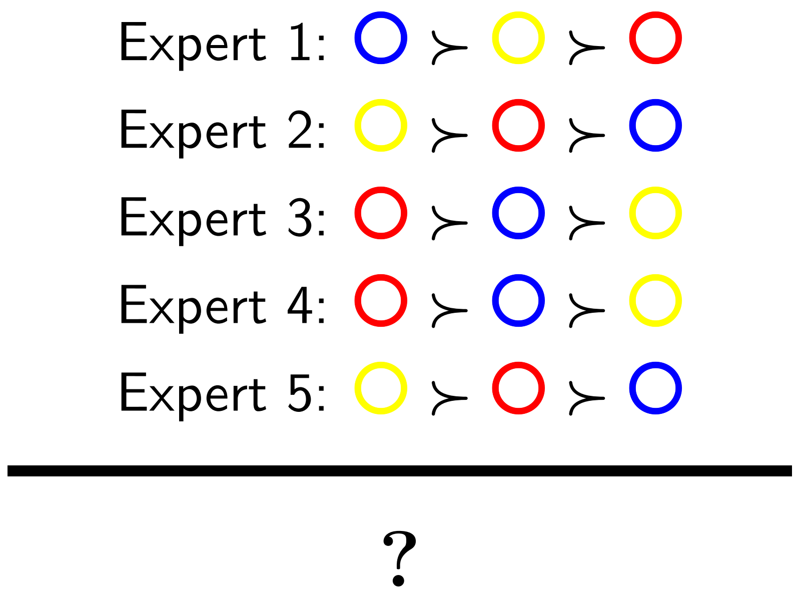
Expert 1: ≻ ≻
Expert 2: ≻ ≻
Expert 3: ≻ ≻
Expert 4: ≻ ≻
Expert 5: ≻ ≻
?

Social Choice and the Condorcet Paradox

Social Choice Theory asks: how should we aggregate the preferences of the members of a group to obtain a "social preference"?

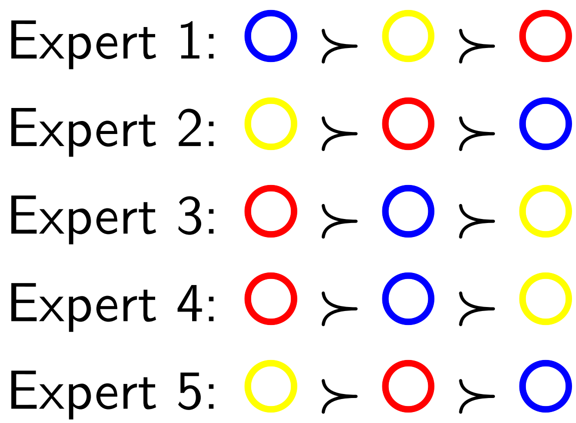

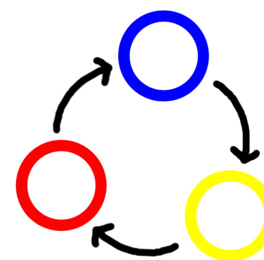

Marie Jean Antoine Nicolas de Caritat (1743–1794), better known as the **Marquis de Condorcet**: Highly influential Mathematician, Philosopher, Political Scientist, Political Activist. Observed that the *majority rule* may produce inconsistent outcomes ("Condorcet Paradox").

A Classic: Arrow's Impossibility Theorem

In 1951, K.J. Arrow published his famous *Impossibility Theorem:*

Any preference aggregation mechanism for *three* or more alternatives that satisfies the axioms of *unanimity* and *IIA* must be *dictatorial*.

- Unanimity: if everyone says $A \succ B$, then so should society.
- Independence of Irrelevant Alternatives (IIA): if society says $A \succ B$ and someone changes their ranking of C, then society should still say $A \succ B$.

Kenneth J. Arrow (born 1921): American Economist; Professor Emeritus of Economics at Stanford; Nobel Prize in Economics 1972 (youngest recipient ever). His 1951 PhD thesis started modern Social Choice Theory. Google Scholar lists 12,792 citations of the thesis.

Social Choice and Computer Science

Social choice theory has natural applications in computer science:

- *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
- *Multiagent Systems:* to aggregate the beliefs + to coordinate the actions of groups of autonomous software agents

Vice versa, techniques from computer science are useful for advancing the state of the art in social choice theory . . .

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

Logical Modelling

What kind of features do we need in a *logic* to be able to reason about problems in social choice?

Example for a result:

Theorem 1 *The first-order theory* T_{ARROW} *has no finite model.*

Also of interest:

- use of automated theorem provers to confirm results
- automated search for new results with variants of axioms
- model checking to assess concrete algorithms for voting rules

U. Grandi and U. Endriss. First-Order Logic Formalisation of Impossibility Theorems in Preference Aggregation. *Journal of Philosophical Logic*. In press (2012).

Example: Strategic Manipulation

Remember Florida 2000 (simplified):

49%:	Bush $\succ$ Gore $\succ$ Nader
20%:	Gore $\succ$ Nader $\succ$ Bush
20%:	Gore $\succ$ Bush $\succ$ Nader
11%:	Nader $\succ$ Gore $\succ$ Bush

Questions:

- Who wins?
- What would your advice to the Nader-supporters have been?

Complexity as a Barrier against Manipulation

By the classical *Gibbard-Satterthwaite Theorem*, any voting rule for $\geqslant 3$ candidates can be manipulated (unless it is dictatorial).

Idea: So it's always *possible* to manipulate, but maybe it's *difficult*! Tools from *complexity theory* can be used to make this idea precise.

- For *some* procedures this does *not* work: if I know all other ballots and want X to win, it is *easy* to compute my best strategy.
- But for *others* it does work: manipulation is *NP-complete*.

Recent work in COMSOC has expanded on this idea:

- NP is a worst-case notion. What about average complexity?
- Also: complexity of winner determination, control, bribery, ...

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Soc. Choice and Welfare*, 6(3):227–241, 1989.

P. Faliszewski, E. Hemaspaandra, and L.A. Hemaspaandra. Using Complexity to Protect Elections. *Communications of the ACM*, 553(11):74–82, 2010.

Multi-issue Elections

Suppose 13 voters are asked to each vote *yes* or *no* on three issues; and we use the simple majority rule for each issue independently:

- 3 voters each vote for YNN, NYN, NNY.
- 1 voter each votes for YYY, YYN, YNY, NYY.
- No voter votes for NNN.

But then NNN wins! (on each issue, 7 out of 13 vote *no*)

What to do instead? The number of candidates is *exponential* in the number of issues (e.g., $2^3 = 8$), so even just representing the voters' preferences is a challenge ($\leadsto$ *knowledge representation*).

S.J. Brams, D.M. Kilgour, and W.S. Zwicker. The Paradox of Multiple Elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

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.

Paradox?

	p	$p \to q$	q
Judge 1:	True	True	True
Judge 2:	True	False	False
Judge 3:	False	True	False

?

	fund museum?	*fund school?*	*fund metro?*
Voter 1:	Yes	Yes	No
Voter 2:	Yes	No	Yes
Voter 3:	No	Yes	Yes

?

[Constraint: we have money for *at most two projects*]

General Perspective

We can view many of our problems as problems of *binary aggregation:*

Do you rank option ○ *above option* ○ *?*	Yes/No
Do you believe formula "$p \rightarrow q$" is true?	Yes/No
Do you want the new school to get funded?	Yes/No

Each problem domain comes with its own *integrity constraints:*

Rankings should be transitive and not have any cycles.

The accepted set of formulas should be logically consistent.

We should fund at most two projects.

The *paradoxes* we have seen show that the *majority rule* does not *lift* our integrity constraints from the *individual* to the *collective* level.

Characterisation Results

So: Which aggregation rules lift which integrity constraints?

Example for a result:

Theorem 2 *An aggregator F will lift all integrity constraints that can be expressed as a conjunction of literals if and only if F is unanimous.*

U. Grandi and U. Endriss. Lifting Integrity Constraints in Binary Aggregation. *Artificial Intelligence*, 199–200:45–66, 2013.

Can we avoid all paradoxes?

That is: Are the aggregators that lift *all* integrity constraints? *Yes!*

Theorem 3 *An aggregator F will lift all integrity constraints if and only if F is a generalised dictatorship (that is, if F is defined by a function g from profiles to agents via $F(B_1, \ldots, B_n) = B_{g(B_1,\ldots,B_n)}$).*

This includes some pretty *bad* aggregators:

- proper (Arrovian) *dictatorships:* $g \equiv i$ (dictator fixed in advance)

And some that look at least *interesting:*

- return the individual vector closest to the *majority* vector
- return the individual vector closest to the *average* vector

U. Grandi and U. Endriss. Lifting Integrity Constraints in Binary Aggregation. *Artificial Intelligence*, 199–200:45–66, 2013.

Voting as Choosing the Most Representative Voter

Somewhat surprisingly, this *majority-voter rule* and *average-voter rule* have excellent properties:

- *no paradoxes* (outcomes are always consistent)
- *low complexity* (MVP slightly lower than AVP)
- *2-approximations* of the (intractable) distance-based rule returning the consistent vector closest to the profile (AVP slightly better)
- satisfaction of *choice-theoretic axioms* (except for independence): anonymity, neutrality, unanimity (MVP also reinforcement)

That is, our method of seeking to characterise aggregators via the IC's they lift has helped to identify useful practical methods . . .

U. Endriss and U. Grandi. Binary Aggregation by Selection of the Most Representative Voter. *Proc. MPREF-2013*.

Last Slide

I have tried to offer a glimpse at *computational social choice*.

Examples discussed:

- *logical modelling* in social choice (Arrow's Theorem in FOL)
- *computational hardness* as a barrier against strategic behaviour
- *choice-theoretic* and *algorithmic* challenges in multi-issue elections
- *characterisation* of aggregation rules in terms of the IC's lifted

COMSOC is a booming field of research with lots of opportunities (and links to your favourite topic in computation yet to be discovered).

To find out more about the field, you could have a look at this website (biannual workshop series, PhD theses, mailing list):

`http://www.illc.uva.nl/COMSOC/`