Making Collective Choices: Guest Lecture for Logic, Language and Computation

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

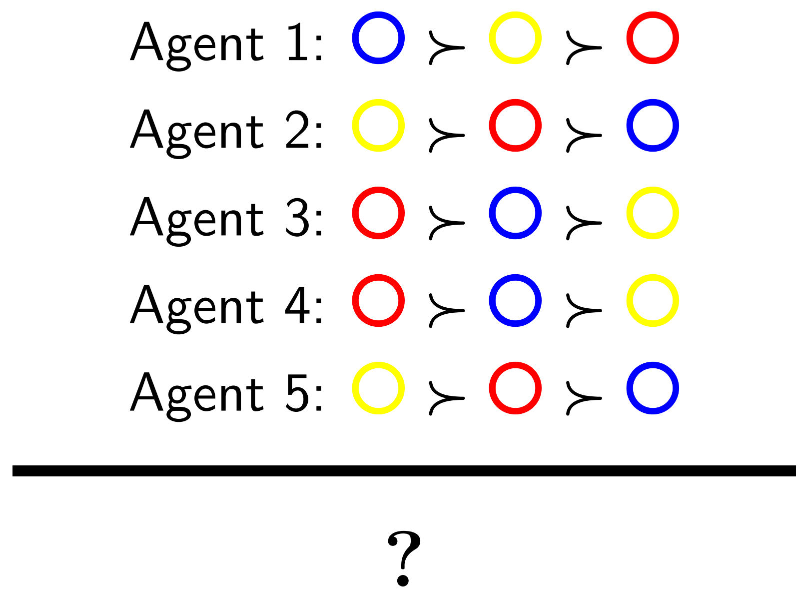
Agent 1:
Agent 2:
Agent 3:
Agent 4:
Agent 5:
?

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"?

Agent 1: ${\color{blue}\bigcirc} \succ {\color{yellow}\bigcirc} \succ {\color{red}\bigcirc}$

Agent 2: ${\color{yellow}\bigcirc} \succ {\color{red}\bigcirc} \succ {\color{blue}\bigcirc}$

Agent 3: ${\color{red}\bigcirc} \succ {\color{blue}\bigcirc} \succ {\color{yellow}\bigcirc}$

Agent 4: ${\color{red}\bigcirc} \succ {\color{blue}\bigcirc} \succ {\color{yellow}\bigcirc}$

Agent 5: ${\color{yellow}\bigcirc} \succ {\color{red}\bigcirc} \succ {\color{blue}\bigcirc}$

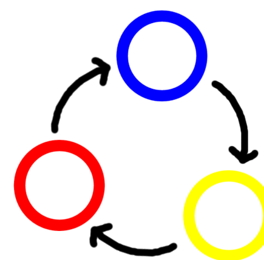

Marie Jean Antoine Nicolas de Caritat (1743–1794), better known as the **Marquis de Condorcet**: Highly influential Mathematician, Philosopher, Political Scientist, Political Activist. Observed that the *majority rule* may produce inconsistent outcomes ("Condorcet Paradox").

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 16,287 citations of the thesis.

Research Area: Computational Social Choice

The philosophical and mathematical study of different methods for *collective decision making* is known as *social choice theory*.

Classical SCT is mostly about political decision making. But in fact, the basic principles are relevant to all these questions:

- How to choose a president given people's preferences?
- How to combine the website rankings of multiple search engines?
- How to divide a cake between several children?
- How to assign bandwidth to competing processes on a network?
- How to assign student doctors to hospitals?
- How to aggregate the views of different judges in a court case?
- How to extract information from noisy crowdsourced data?

Computational social choice, my main area of research, emphasises the fact that any method of decision making is ultimately an *algorithm*.

Outline of Rest of Talk

- More *examples* for challenges when making collective choices
- Identification of a *common pattern*
- Understanding links between properties of *aggregation rules* and properties of the *space of feasible choices*

	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?*
Councillor 1:	Yes	Yes	No
Councillor 2:	Yes	No	Yes
Councillor 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 "$p \to q$" 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?

Classical perspective (following Arrow):

- specify formal *axioms* for intuitively "good" *aggregators*
  Examples: anonymity, neutrality, unanimity, monotonicity, . . .

Alternative perspective:

- specify *expressive power* of language for *integrity constraints*
  Example: $\neg(\texttt{museum} \wedge \texttt{school} \wedge \texttt{metro})$

Example for a result:

Theorem 1 *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 majority paradoxes?

We saw that the majority rule generates paradoxes: it doesn't always lift the integrity constraint from the individual to the collective level.

So: For *which* integrity constraints can this (not) happen?

The following result holds for *odd* numbers of agents:

Theorem 2 *The majority rule lifts an integrity constraint if and only if it can be written in 2-CNF (as a conjunction of clauses of length $\leqslant 2$).*

Indeed, all the paradoxes we saw involved IC's with 3-clauses.

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. 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 there aggregators that lift *all* integrity constraints? *Yes!*

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

To be sure, this includes some pretty *bad* aggregators:

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

But also some that look fairly *interesting:*

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

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

U. Endriss and U. Grandi. Binary Aggregation by Selection of the Most Representative Voter. Proc. AAAI-2014.

More Examples of Recent Local Research

Complexity Theory: How difficult is it to check whether a given aggregation problem is paradox-safe? Relevant parameters?

Game Theory: What does it mean for an agent or a group of agents to strategically manipulate the aggregation by providing false input? Which aggregators are immune to such attacks?

Linguistics: What aggregation methods are appropriate for extracting a high-quality annotation of a linguistic corpus from many low-quality annotations obtained by means of crowdsourcing?

U. Endriss, R. de Haan, and S. Szeider. Parameterized Complexity Results for Agenda Safety in Judgment Aggregation. Proc. AAMAS-2015.

S. Botan, A. Novaro, and U. Endriss. Group Manipulation in Judgment Aggregation. Proc. AAMAS-2016.

C. Qing, U. Endriss, R. Fernández, and J. Kruger. Empirical Analysis of Aggregation Methods for Collective Annotation. Proc. COLING-2014.

Last Slide

I have tried to offer a glimpse at *computational social choice* and presented one particular line of research in this broad field:

- many paradoxes of collective choice have a common structure
- useful general model: *binary aggregation with integrity constraints*
- necessary and sufficient conditions for paradox-free aggregation

COMSOC is a booming field of research with lots of opportunities.

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/`

Or you could read the *Handbook*. Or take my course next year.

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