Computational Social Choice 2026

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

http://www.illc.uva.nl/~ulle/teaching/comsoc/2026/

Opening Example

Five *voters* express their *preferences* over three *alternatives*. We need to find a *good* ranking of the alternatives to reflect this information:

Voter 1:	$a \succ b \succ c$
Voter 2:	$b \succ c \succ a$
Voter 3:	$c \succ a \succ b$
Voter 4:	$c \succ a \succ b$
Voter 5:	$b \succ c \succ a$
	?

What is Computational Social Choice?

Social choice theory is about methods for *collective decision making*, such as *political* decision making by groups of *economic* agents.

Its methodology ranges from the *philosophical* to the *mathematical*. It is traditionally studied in *Economics* and *Political Science* and it is a close cousin of both *decision theory* and *game theory*.

Its findings are relevant to multiple *applications*, such as these:

- How to fairly allocate resources to the members of a society?
- How to fairly divide computing time between several users?
- How to elect a president given people's preferences?
- How to combine the website rankings of multiple search engines?
- How to aggregate the views of different judges in a court case?
- How to extract information from noisy crowdsourced data?

Computational social choice, the topic of this course, emphasises the fact that any method of decision making is ultimately an *algorithm*.

Plan for Today

The purpose of today's lecture is to give you enough information to decide whether you want to take this course.

- Examples for domains (and techniques) in COMSOC research
- Organisational matters: planning, expectations, assessment

Three Voting Rules

Suppose n *voters* choose from a set of m *alternatives* by stating their preferences in the form of *strict rankings* of the 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
- *Borda:* each voter (implicitly) gives $m-1$ points to the alternative she ranks first, $m-2$ to the alternative she ranks second, etc.; and the alternative with the most points wins

Exercise: *Do you know real-world elections where these rules are used?*

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 French people:	Wine	$\succ$	Beer	$\succ$ Milk
4 Dutch people:	Milk	$\succ$	Beer	$\succ$ Wine

Recall that we saw three different voting rules:

- Plurality
- Plurality with runoff
- Borda

Exercise: *For each of the rules, which beverage wins the election?*

The Axiomatic Method

So many voting rules! *How do you choose the right one?*

One approach is to use the *axiomatic method:*

- identify so-called *axioms:* properties of rules we care about
- give precise definitions of axioms and explore their consequences

Results you might obtain:

- *Separation:* one rule satisfies the axiom and another does not
- *Characterisation:* one rule is the only one satisfying certain axioms
- *Impossibility:* no rule can possibly satisfy a certain set of axioms

Anonymity and Neutrality

Two *axioms* encoding fairness we might want our voting rule to satisfy:

- *Anonymity:* Treat all voters the same.
- *Neutrality:* Treat all alternatives the same.

Exercise: *How would you go about formalising these axioms?*

Resolute and Irresolute Voting Rules

A voting rule is called *resolute* if it returns a single winning alternative for every possible profile of preferences. Very useful to have.

Exercise: *Give an example for a voting rule that is resolute!*

For the remainder of this slide, let us restrict attention to voting rules for scenarios with *two voters* ($n = 2$) and *two alternatives* ($m = 2$).

Exercise: *Show that there exists no resolute voting rule that is 'fair' in the sense of being both anonymous and neutral.*

Exercise: *But there still are some irresolute voting rules that are both anonymous and neutral. Give some examples!*

Epistemic Social Choice

The axiomatic method is adequate when we're after fair voting rules that duly respect individual preferences. *Not always what we need!*

Sometimes we are looking for an *objectively best alternative* and treat voters as experts who report noisy versions of the objective ranking.

Condorcet Jury Theorem (1795): *Suppose a jury of n voters need to select the better of two alternatives and each voter independently makes the correct decision with the same probability $p > \frac{1}{2}$. Then the probability that the simple majority rule returns the correct decision increases monotonically in n and approaches 1 as n goes to infinity.*

Exercise: *Explain why this is true!*

Writings of the Marquis de Condorcet. In I. McLean and A. Urken (eds.), *Classics of Social Choice*. University of Michigan Press, 1995.

Different Models

Modelling voter preferences as *strict rankings* and looking for rules that return a *single winning alternative* is but one model of voting.

Exercise: *Can you think of other models of voting?*

Participatory Budgeting

People vote on grassroots projects (park benches, cultural activities) and then the most popular projects receive public funding.

Used to allocate many millions each year. Common approach:

- Ballot format: ask each voter to *approve* some projects
- Voting rule: *Greedy Approval* (fund projects with most approvals)

Greedy Approval is not *proportional:* majority might get everything.

Exercise: *Describe a situation (approval profile) where this happens!*

Lots of recent research on designing better voting rules for PB ...

Working with Data

Some phenomena cannot be adequately studied by analytical methods. Sometimes we need to run experiments involving preference data.

Example (MoL individual project by Pelle Nelissen, 2023):

> *For PB elections in Amsterdam, replacing Greedy Approval by a modern proportional rule would cut in half the number of residents not getting any project funded they approved.*

You can explore PB data at `pabulib.org` and `pabuviz.org`.

P. Faliszewski, J. Flis, D. Peters, G. Pierczyński, P. Skowron, D. Stolicki, S. Szufa, and N. Talmon. Participatory Budgeting: Data, Tools, and Analysis. IJCAI-2023.

P. Nelissen. An Empirical Analysis of Participatory Budgeting in Amsterdam. ArXiv, 2023 (`arxiv.org/abs/2310.18033`).

Beyond Political Elections

Social choice is relevant also beyond political elections. Examples:

- *AI alignment:* Aligning GenAI systems to (diverse) human values might be modelled as a problem of social choice and techniques developed in COMSOC might help improve *RLHF*.
- *Consensus clustering:* Computing a consensus clustering of a dataset (an equivalence relation) given individual clusterings produced by several algorithms is similar to preference aggregation.

V. Conitzer, R. Freedman, J. Heitzig, W.H. Holliday, B.M. Jacobs, N. Lambert, M. Mossé, E. Pacuit, S. Russell, H. Schoelkopf, E. Tewolde, and W.S. Zwicker. Social Choice Should Guide AI Alignment in Dealing with Diverse Human Feedback. Position paper presented at ICML-2024.

U. Endriss and U. Grandi. Graph Aggregation. *Artificial Intelligence*, 2017.

Fair Allocation

We need to divide a number of *goods* between a number of *agents*.
Each agent has a *utility function* mapping *bundles* to real numbers.
For today, focus on scenarios where all utility functions are *additive*.

We are looking for a *fair* allocation of the goods to the agents.

Exercise: *What would you consider fair in this context?*

Envy-Freeness

One nice fairness notion is *envy-freeness:* in the chosen allocation, nobody should want to swap with any of the other agents.

Not always achievable if we insist on allocating all goods! (*Why?*)

A weaker requirement is *envy-freeness up to one good* (EF1): any existing envy can be eliminated by removing (at most) one good from the bundle given to the agent being envied.

This is always achievable!

Proof: Let agents pick goods in a *round robin* fashion. ✓

R. Lipton, E. Markakis, E. Mossel, and A. Saberi. On Approximately Fair Allocations of Indivisible Goods. EC-2004.

Assumptions on Preferences

Ultimately, agents have *preferences over allocations*. So could think of fair allocation as a *voting problem* (alternatives = allocations).

For voting we modelled preferences as *linear orders* on alternatives.
For fair allocation we instead used *utility functions*. *Good idea?*

- Using utility functions presupposes that *preference intensity* makes sense ("I like A *twice as much* as B", not just "*more than*").
- Also: *interpersonal comparison* ("I like A *as much* as you like B").

This might not always be appropriate. *So the choice of model matters!*

Ranking Sets of Objects

Suppose we know your preferences $\succcurlyeq$ over a finite number of *objects:*

$$a_m \succ a_{m-1} \succ \cdots \succ a_3 \succ a_2 \succ a_1$$

When you compare *sets of objects*, representing *opportunities*, what can we say about your preferences $\hat{\succcurlyeq}$ over sets of objects?

- It seems uncontroversial that $\{a_3\} \hat{\succ} \{a_1, a_2\}$.
- It seems impossible infer anything regarding $\{a_2\}$ and $\{a_1, a_3\}$.
- We might be willing to infer $\{a_1, a_3, a_4\} \hat{\succ} \{a_1, a_2, a_4\}$. (*How?*)

Suppose we accept the following two *axioms* for preference extensions:

- *Independence:* $A \hat{\succ} B$ and $c \notin A \cup B$ imply $A \cup \{c\} \hat{\succcurlyeq} B \cup \{c\}$
- *Dominance:* $b \succ a$ for all $a \in A$ implies $A \cup \{b\} \hat{\succ} A$
  and similarly $b \succ a$ for all $b \in B$ implies $B \hat{\succ} B \cup \{a\}$

Of course, we also want $\hat{\succcurlyeq}$ to be transitive and complete (*weak order*).

Exercise: *Can you think of a way of defining $\hat{\succcurlyeq}$ that works?*

The Kannai-Peleg Theorem

Rather surprisingly, our requirements are impossible to satisfy:

Kannai-Peleg Theorem: *If there are $\geqslant$ 6 objects, then no weak order on sets of objects satisfies both independence and dominance.*

Proof: We first show that $A \mathrel{\hat{\sim}} \{\max(A), \min(A)\}$ for any set A.

> Clear for $|A| \leqslant 2$. For $|A| \geqslant 3$, get $A \setminus \{\max(A)\} \mathrel{\hat{\succ}} \{\min(A)\}$ from (DOM), and then $A \mathrel{\hat{\succcurlyeq}} \{\max(A), \min(A)\}$ from (IND).
> Show $\{\max(A), \min(A)\} \mathrel{\hat{\succcurlyeq}} A$ analogously. ✓

Now suppose $a_6 \succ a_5 \succ a_4 \succ a_3 \succ a_2 \succ a_1$. Show $\{a_2, a_5\} \mathrel{\hat{\succcurlyeq}} \{a_4\}$:

> Assume not: $\{a_4\} \mathrel{\hat{\succ}} \{a_2, a_5\}$. By (IND): $\{a_1, a_4\} \mathrel{\hat{\succcurlyeq}} \{a_1, a_2, a_5\}$.
> By above lemma: $\{a_1, a_2, a_3, a_4\} \succcurlyeq \{a_1, a_2, a_3, a_4, a_5\}$. ↯

Thus also: $\{a_2, a_5\} \mathrel{\hat{\succ}} \{a_3\}$, and by (IND): $\{a_2, a_5, a_6\} \mathrel{\hat{\succcurlyeq}} \{a_3, a_6\}$.
By above lemma: $\{a_2, a_3, a_4, a_5, a_6\} \mathrel{\hat{\succcurlyeq}} \{a_3, a_4, a_5, a_6\}$. ↯

Y. Kannai and B. Peleg. A Note on the Extension of an Order on a Set to the Power Set. *Journal of Economic Theory*, 1984.

Automated Discovery of Theorems

A major challenge in COMSOC is to facilitate *automated verification*, and possibly even the *automated discovery*, of theorems.

Works for *ranking sets of objects* (Christian Geist's MoL thesis, 2010):

- Devise logic for expressing axioms (many-sorted FOL).
- Find syntactic conditions on axioms under which any impossibility for k objects generalises to $k' > k$ objects.
- For any fixed k, axioms can be expressed in propositional logic, and impossibility for small fixed k can be checked by a SAT solver.
- Search over all combinations of axioms from some set (we used 20) and all values of k up to some bound (we used 8) to discover all impossibilities (we found 84 impossibility theorems).

C. Geist and U. Endriss. Automated Search for Impossibility Theorems in Social Choice Theory: Ranking Sets of Objects. *Journal of AI Research*, 2011.

Judgment Aggregation

Suppose three robots are in charge of climate control for this building. They need to make judgments on p (*the temperature is below 17°C*), q (*we should switch on the heating*), and the 'policy' $p \to q$.

	p	$p \to q$	q
Robot 1	Yes	Yes	Yes
Robot 2	No	Yes	No
Robot 3	Yes	No	No

Exercise: *Should we switch on the heating?*

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 2002.

Algorithm Design and Complexity Theory

So the *majority rule* is not a good choice for judgment aggregation, given that it will sometimes return *inconsistent judgment sets*.

We could design more sophisticated rules to avoid this. Example:

> Compute the *majority judgment set*. Then, from amongst all *consistent judgment sets*, return the one that is *closest* to that majority outcome, where the *distance* between two judgment sets is the number of propositions on which they differ.

Exercise: *Can you design an algorithm for this? What's its runtime?*

U. Endriss, R. de Haan, J. Lang, and M. Slavkovik. The Complexity Landscape of Outcome Determination in Judgment Aggregation. *Journal of AI Research*, 2020.

Summary

COMSOC is all about *aggregating* information supplied by *individuals* into a *collective* view. Different *domains* of aggregation:

- *fair allocation:* preferences over highly structured alternatives
- *voting:* ordinal preferences over alternatives w/o internal structure
- *judgment aggregation:* assignments of truth values to propositions

Different *techniques* used to analyse them, such as:

- axiomatic method: philosophical and mathematical
- logical modelling, automated theorem proving
- algorithm design and complexity analysis
- probability theory (e.g., for truth-tracking)
- new questions in view of applications beyond politics

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

Organisational Matters

Prerequisites: This is an advanced course, so I assume mathematical maturity, we'll move fast, and we'll often touch upon recent research. On the other hand, almost no specific background is required.

Information: Website for slides, homework assignments, and readings. Canvas for assignment submission, announcements, and rules.

Rules: Read the course manual, especially on collaboration and GenAI.

Assessment: Homework (50%) and mini-project (50%).

Commitment: Be ready to invest ∼20h/week. Heavy HW regime for the first three weeks; after that the focus is on the projects.

Schedule: Two or three meetings per week. No tutorials.

Attendance: You should usually be present at all meetings.

Research Seminars: Sometimes attending research seminars should be normal for MSc students. The COMSOC Seminar is one option.

Assignment

Form a *group of two* and prepare a brief presentation on your assigned voting rule for tomorrow (no slides, both of you talk):

- 60 seconds on how the rule works
- 30 seconds on a positive property of the rule
- 30 seconds on a negative property of the rule

Be ready to answer a few questions about the material you present.

Outlook

Plan for the rest of the course:

- focus on voting theory
- axiomatic method and impossibility results
- strategic behaviour and strategyproofness
- then: broad exploration of methodological tools
- finally: various other topics

Prepare your presentation + read the non-technical parts of Chapter 1 of the *Handbook* to learn about the history of the field of COMSOC.

What next? Definition, classification, and properties of voting rules.