Computational Social Choice 2024

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

[http://www.illc.uva.nl/~ulle/teaching/comsoc/2024/]

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:
  - fair allocation of goods
  - preference modelling
  - voting in elections
  - judgment aggregation
- Organisational matters: planning, expectations, assessment, ...

Cake Cutting

A classical example for a problem of collective decision making:

> *We have to divide a cake with different toppings amongst n agents by means of parallel cuts. Agents have different preferences regarding the toppings (additive utility functions).*

```
|----------------------|
0                      1
```

The exact details of the formal model are not important for this short exposition. You can look them up in my lecture notes (cited below).

<u>Exercise:</u> *Can you think of a suitable method for $n = 2$ agents?*

U. Endriss. *Lecture Notes on Fair Division*. Institute for Logic, Language and Computation, University of Amsterdam, 2009.

Cut-and-Choose

The classical approach for dividing a cake between *two agents:*

- *One agent cuts the cake into two pieces (of equal value to her), and the other chooses one of them (the piece she prefers).*

The cut-and-choose protocol is *fair* in the sense of guaranteeing a property known as *proportionality:*

- Each agent is *guaranteed* at least one half (general: $1/n$), according to her own valuation, however the other one plays.
- Discussion: In fact, the first agent (if she is risk-averse) will receive exactly $1/2$, while the second will usually get more.

Exercise: *What about three agents? Or more?*

The Banach-Knaster Last-Diminisher Protocol

In the first ever paper on fair division, Steinhaus (1948) reports on a *proportional* protocol for n *agents* due to Banach and Knaster.

(1) Agent 1 cuts off a piece (that she considers to represent $1/n$).

(2) That piece is passed around. Each agent either lets it pass (if she finds it too small) or trims it further (to what she considers $1/n$).

(3) After the piece has made the full round, the last agent to cut something off (the "last diminisher") is obliged to take it.

(4) The rest (including the trimmings) is then divided amongst the remaining $n-1$ agents. Last agent takes what's left. ✓

Observe that each agent is guaranteed a *proportional* piece. (*Why?*)

Exercise: *How many cuts do we require in the worst case?*

H. Steinhaus. The Problem of Fair Division. *Econometrica*, 1948.

The Even-Paz Divide-and-Conquer Protocol

So the last-diminisher protocol requires $O(n^2)$ cuts. *Can we do better?*

Even and Paz (1984) introduced the divide-and-conquer protocol:

(1) Ask each agent to put a (cut) *mark* on the cake.

(2) *Cut* the cake at the $\lfloor \frac{n}{2} \rfloor$*th mark* (counting from the left).
Associate the agents who made the *leftmost* $\lfloor \frac{n}{2} \rfloor$ *marks* with the *lefthand part*, and the *remaining agents* with the *righthand part*.

(3) *Repeat* for each group, until only one agent is left.

Also here, each agent is guaranteed a *proportional* piece. (*Why?*)

Exercise: *How many cuts/marks do we need now (Big-O notation)?*

S. Even and A. Paz. A Note on Cake Cutting. *Discrete Appl. Mathematics*, 1984.

Preference Modelling

For the cake-cutting scenario, we made some very specific assumptions regarding the *preferences* of individual agents:

- preferences are modelled as *utility functions* (so: using numbers)
- those preferences are *additive* (severe restriction)

Discussion: *cardinal* utility function *vs.* *ordinal* preference relation

We also did not worry about what formal *language* to use to *represent* an agent's preferences, e.g., to be able to say *how much information* you need to exchange when eliciting an agent's preferences.

Preference representation is an interesting research field in its own right. A possible starting point is the survey cited below.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 2008.

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 $A \,\hat{\sim}\, \{\max(A), \min(A)\}$ 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.

Three Voting Rules

Suppose n *voters* choose from a set of m *alternatives* by stating their preferences in the form of *linear orders* over the 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
- *Borda:* each voter 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?*

Special Case: Two Alternatives

For the special case of $m = 2$ alternatives, all three voting rules we saw reduce to the *same* rule (*why?*), known as the *simple majority rule*.

Intuitively, this is "the right" way to run an election when $m = 2$.

Exercise: *Can you make this precise? Why is there no better rule?*

The Condorcet Jury Theorem

We'll mostly focus on axiomatic arguments, but there are also others.

The simple majority rule for two alternatives is *epistemically* attractive, in terms of *tracking the truth* (assuming there is a "correct" choice):

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.*

Proof sketch: By the law of large numbers, the number of voters making the correct choice approaches $p \cdot n > \frac{1}{2} \cdot n$. ✓

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

Example: 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" $p \rightarrow q$.

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

Exercise: *Should we switch on the heating?*

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

Differen *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 regulations.

Schedule: Usually two meetings per week, but sometimes more.

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.

Attendance: You should usually be present at all meetings.

Research Seminars

Make it a habit to regularly attend research seminars in multiple fields.

For COMSOC, we have a local option (irregular meetings):

COMSOC Seminar at the ILLC
(`bit.ly/comsoc-seminar-amsterdam`)

There also is an international online seminar happening once a month:

COMSOC Video Seminar
(`comsocseminar.org`)

Outlook

Plan for the rest of the course:

- application focus: voting in elections
- axiomatic method and impossibility results
- strategic behaviour and strategyproofness
- methodological focus: representation + reasoning, SAT technology
- various minor topics

Assignments for tomorrow:

- Read the non-technical parts of Chapter 1 of the *Handbook* to get an understanding of the nature and history of the field.
- Prepare a 90-second presentation about your assigned voting rule (definition + one positive property + one negative property).

What next? Classifying and axiomatically characterising voting rules.