Computational Social Choice 2022

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

[http://www.illc.uva.nl/~ulle/teaching/comsoc/2022/]

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 problems and techniques in COMSOC research:
  - fair allocation of goods
  - 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. 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 one 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. ✓

Each agent is guaranteed a *proportional* piece. Requires $O(n^2)$ cuts.

Exercise: *How many cuts exactly?*

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

The Even-Paz Divide-and-Conquer Protocol

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

(1) Ask each agent to put a *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.

Exercise: *How many cuts do we need this time (Big-O notation)?*

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

Preferences

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

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

<ins>Exercise:</ins> *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?*

Axiomatic Method

So how do you decide which is the right voting rule to use?

The classical approach is to use the so-called *axiomatic method*:

- identify normatively appealing properties of rules
- cast those properties into mathematically rigorous definitions
- explore the consequences of the thus defined "axioms"

The definitions on the following slide are only sketched, but can be made mathematically precise (see the paper cited below for how).

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*. College Publications, 2011.

May's Theorem

When there are only *two alternatives*, then all the voting rules we have seen coincide. This is usually called the *simple majority rule* (SMR).

Intuitively, it does the "right" thing. Can we make this precise? *Yes!*

Theorem 1 (May, 1952) *A voting rule for two alternatives satisfies anonymity, neutrality, and positive responsiveness iff it is the SMR.*

Meaning of these *axioms:*

- *anonymity* = voters are treated symmetrically
- *neutrality* = alternatives are treated symmetrically
- *positive responsiveness* = if x is the (sole or tied) winner and one voter switches from y to x, then x becomes the sole winner

Exercise: *One direction is easy. Which one? Prove it!*

K.O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 1952.

Proof Sketch

We want to prove:

> *A voting rule for two alternatives satisfies anonymity, neutrality, and positive responsiveness iff it is the SMR.*

($\Leftarrow$) Clearly, the simple majority rule has all three properties. ✓

($\Rightarrow$) Assume *#voters* is *odd* (other case: similar) $\rightsquigarrow$ no ties

Let a A be the set of voters voting $a \succ b$ and B those voting $b \succ a$.
Anonymity $\rightsquigarrow$ only *number* of ballots of each type matters. Cases:

- If $|A| = |B| + 1$, then *only a wins*. $\rightsquigarrow$ By *PR*, a wins for $|A| > |B|$. By *neutrality*, b wins otherwise. So we get the SMR. ✓
- There exist A, B with $|A| = |B|+1$ yet *b wins*. $\rightsquigarrow$ Let one a-voter switch to b. By *PR*, now *only* b wins. But now $|B'| = |A'| + 1$, which is symmetric to the first situation, so by *neutrality* a wins. $\lightning$

The Condorcet Jury Theorem

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

Theorem 2 (Condorcet, 1785) *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.

More Voting Rules: Positional Scoring Rules

We can generalise the idea underlying the Borda rule as follows:

A *positional scoring rule* (PSR) is defined by a so-called *scoring vector* $\boldsymbol{s} = (s_1, \ldots, s_m) \in \mathbb{R}^m$ with $s_1 \geqslant s_2 \geqslant \cdots \geqslant s_m$ and $s_1 > s_m$.

Each voter submits a ranking of the m alternatives. Each alternative receives s_j points for every voter putting it at the jth position.

The alternative(s) with the highest score (sum of points) win(s).

Examples:

- *Borda rule* $=$ PSR with scoring vector $(m-1, m-2, \ldots, 0)$
- *Plurality rule* $=$ PSR with scoring vector $(1, 0, \ldots, 0)$
- *Antiplurality* (or *veto*) *rule* $=$ PSR with scoring vector $(1, \ldots, 1, 0)$
- For any $k < m$, *k-approval* $=$ PSR with $(\underbrace{1, \ldots, 1}_{k}, 0, \ldots, 0)$

Exercise: *Name the rule induced by $\boldsymbol{s} = (9, 7, 5)$. General idea?*

Another Axiom: The Condorcet Principle

Another idea going back to Condorcet: an alternative beating all other alternatives in pairwise majority contests is a *Condorcet winner*.

Sometimes there is no Condorcet winner (*Condorcet paradox*):

Ann:	$a \succ b \succ c$
Bob:	$b \succ c \succ a$
Cindy:	$c \succ a \succ b$

But if a Condorcet winner exists, then it must be *unique*.

A voting rule satisfies the *Condorcet Principle*, if it elects (only) the Condorcet winner whenever one exists.

Positional Scoring Rules and the Condorcet Principle

Consider this example with three alternatives and seven voters:

$$\begin{array}{ll} \text{3 voters:} & a \succ b \succ c \\ \text{2 voters:} & b \succ c \succ a \\ \text{1 voter:} & b \succ a \succ c \\ \text{1 voter:} & c \succ a \succ b \end{array}$$

So a is the *Condorcet winner:* a beats both b and c (by 4 to 3).
But any *positional scoring rule* makes b win (because $s_1 \geqslant s_2 \geqslant s_3$):

$$\begin{array}{ll} a\text{:} & 3 \cdot s_1 + 2 \cdot s_2 + 2 \cdot s_3 \\ b\text{:} & 3 \cdot s_1 + 3 \cdot s_2 + 1 \cdot s_3 \\ c\text{:} & 1 \cdot s_1 + 2 \cdot s_2 + 4 \cdot s_3 \end{array}$$

Thus, *no positional scoring rule* for three (or more) alternatives can possibly satisfy the *Condorcet Principle*.

Dodgson's Rule and its Complexity

Here is a rule that satisfies the Condorcet Principle. It was proposed by C.L. Dodgson (a.k.a. Lewis Carroll, author of *Alice in Wonderland*).

> *If a Condorcet winner exists, elect it. Otherwise, for each alternative x compute the number of adjacent swaps in the individual preferences required for x to become a Condorcet winner. Elect the alternative(s) that minimise that number.*

But this voting rule is particularly hard to compute:

Theorem 3 (Hemaspaandra et al., 1997) *Winner determination for Dodgson's rule is complete for parallel access to NP.*

Writings of C.L. Dodgson. In I. McLean and A. Urken (eds.), *Classics of Social Choice*, University of Michigan Press, 1995.

E. Hemaspaandra, L. Hemaspaandra and J. Rothe. Exact Analysis of Dodgson Elections: Lewis Carroll's 1876 Voting System is Complete for Parallel Access to NP. *Journal of the ACM*, 1997.

Example: Strategic Manipulation

Recall that under the *plurality rule* (used in most political elections) the candidate ranked first most often wins the election.

Assume the preferences of the people in, say, Florida are as follows:

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

So even if nobody is cheating, Bush will win this election.

> It would have been in the interest of the Nader supporters to *manipulate*, i.e., to misrepresent their preferences.

Exercise: *Is there a better voting rule that avoids this problem?*

The Gibbard-Satterthwaite Theorem

Answer to the previous question: *No!* — surprisingly, not only the plurality rule, but *all* "reasonable" rules have this problem.

Theorem 4 (Gibbard-Satterthwaite) *All resolute and surjective voting rules for $\geqslant 3$ alternatives are manipulable or dictatorial.*

Meaning of the terms mentioned in the theorem:

- *resolute* = the rule always returns a single winner (no ties)
- *surjective* = each alternative can win for *some* way of voting
- *dictatorial* = the top alternative of some fixed voter always wins

So this is seriously bad news.

A. Gibbard. Manipulation of Voting Schemes. *Econometrica*, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. *Journal of Economic Theory*, 1975.

Logic for Social Choice Theory

Nowadays, the (omitted) proof of the Gibbard-Satterthwaite Theorem is well understood, but after people developed good intuitions that something like G-S must be the case in the 1960s, it still took around a decade before someone was able to prove it. So this is not trivial!

Idea: Cast this in a suitable *logic* and use *automated theorem provers!*

- Nipkow (2009) *verified* a known proof for G-S in ISABELLE.
- *SAT solvers* have been used not only to *re-prove* known theorems but also to *discover* new ones (lots of activity in recent years).

T. Nipkow. Social Choice Theory in HOL. *J. Automated Reasoning*, 2009.

P. Tang and F. Lin. Computer-aided Proofs of Arrow's and other Impossibility Theorems. *Artificial Intelligence*, 2009.

C. Geist and D. Peters. Computer-Aided Methods for Social Choice Theory. In U. Endriss (ed.), *Trends in Computational Social Choice*. AI Access, 2017.

Complexity as a Barrier against Manipulation

By the 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*!

Theorem 5 (Bartholdi and Orlin, 1991) *The manipulation problem for the rule known as single transferable vote (STV) is NP-complete.*

STV is (roughly) defined as follows:

> *Proceed in rounds. In each round, eliminate the current plurality loser. Stop once only one alternative is left.*

Discussion: NP is a worst-case notion. What about average complexity?

J.J. Bartholdi III and J.B. Orlin. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare*, 1991.

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 above 22°C), q (we should switch on the AC), and $p \rightarrow q$.

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

Exercise: *What should be the collective decision?*

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.

Relationship with AI

Ideas from Economics entered AI when it became clear that we can use them to study interaction between agents in a *multiagent system*. Nowadays, the study of so-called *economic paradigms* is all over AI.

The influential *One Hundred Year Study on Artificial Intelligence* (2016) singles out the following eleven "*hot topics*" in AI:

large-scale machine learning | deep learning | reinforcement learning
robotics | computer vision | natural language processing
collaborative systems | crowdsourcing and human computation
algorithmic game theory and *computational social choice*
internet of things | neuromorphic computing

And indeed, while COMSOC transcends several disciplines, about half of it gets published in AI conference proceedings and journals.

P. Stone et al. "Artificial Intelligence and Life in 2030". One Hundred Year Study on Artificial Intelligence. Stanford, 2016.

Plan for the Rest of the Course

We'll focus on *judgment aggregation*, cover this form of aggregation in depth, and see many of the techniques used in the field of COMSOC exemplified in this specific domain (8–10 lectures). Topics:

- formal models for judgment aggregation
- design of aggregation methods based on various principles
- axiomatic method: characterisation and impossibility results
- (briefly) truth-tracking (probabilistic methods)
- strategic behaviour (in the sense of game theory)
- complexity analysis of problems arising in judgment aggregation
- embedding of preference aggregation into judgment aggregation

Plus a few one-off lectures on various *other COMSOC topics*, such as the fair allocation of goods and voting in elections.

Plus a *crash course in complexity theory*, if needed (it probably is).

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 and announcements.

Schedule: Usually two meetings per week.

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

Option to cancel your worst HW grade in return for a project review.

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

Attendance: You should usually be present at all meetings.

Seminars: There occasionally are seminar talks at the ILLC that are relevant to the course and that you are welcome to attend.