Computational Social Choice: Spring 2017

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

Opening Example

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

Agent 1:	$\triangle \succ \bigcirc \succ \square$
Agent 2:	$\bigcirc \succ \square \succ \triangle$
Agent 3:	$\square \succ \triangle \succ \bigcirc$
Agent 4:	$\square \succ \triangle \succ \bigcirc$
Agent 5:	$\bigcirc \succ \square \succ \triangle$

?

Computational Social Choice

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

Its methodology ranges from the *philosophical* to the *mathematical*.

Its findings are relevant to all of these *applications:*

- How to divide a cake between several children?
- How to assign bandwidth to competing processes on a network?
- How to choose a president given people's preferences?
- How to combine the website rankings of multiple search engines?
- How to decide who should get married to whom?
- 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, 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.

- Organisational matters: planning, expectations, assessment, ...
- Examples for problems and techniques in COMSOC research:
  - fair allocation of goods
  - preference modelling
  - voting in elections
  - coalition formation and matching
  - judgment aggregation

Organisational Matters

Prerequisites: This is an advanced course: 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.

Assessment: Homework (50%) + mini-project (50%)

Commitment: Be prepared to invest around 20h/week. You *should* usually be present; you *must* be for all presentations in the exam week.

Website: Lecture slides, literature, homework, project ideas, and other important information will get posted on the course website:

`http://www.illc.uva.nl/~ulle/teaching/comsoc/2017/`

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

Homework

Regular homework during the first part of the course, less of it later on.

- Each assignment will be graded on the usual 1–10 scale.
- I'll disregard the assignment you did worst on.
- Submission is via Blackboard.
- All deadlines are strict.

The usual rules on permissible *collaboration* apply: discussing with others to improve your *understanding* is fine (indeed, it is encouraged), but producing your *solutions* is something you do by yourself.

Whenever additional collaboration is permitted, I will say so explicitly.

Requirements for Homework Solutions

Most questions will be of the problem-solving sort, requiring:

- a good understanding of the topic to see what the question is
- some creativity to find the solution
- mathematical maturity, to write it up correctly, often as a proof
- good taste, to write it up in a reader-friendly manner

Solutions must be typed up professionally (LaTeX strongly preferred).

Of course, solutions should be *correct*. But just as importantly, they should be *short* and *easy to understand*. (This is the advanced level: it's not anymore just about you getting it, it now is about your reader!)

Also: a small number of (optional) *programming assignments*.

Mini-Projects

During the second part of the course you'll work on your mini-project in a small group. Possible types of projects include:

- identify an interesting paper on voting not covered in class and fill in some gaps, or come up with an extension or a generalisation
- apply an algorithmic technique to a problem that to date has only been considered by economists/political scientists/philosophers
- explore an application domain for voting: could be a literature review, an idea for a new application, or an experimental study
- ...

The purpose of this is to provide you with some research experience.

Deliverables: *Presentation* (exam week) + *paper* (by end of block)

Activities: Sessions on *how to write a paper* and *how to give a talk*, and one individual *project meeting* with each group.

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, 2009/2010.

Cut-and-Choose

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

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

The cut-and-choose procedure 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.
- Discussion: In fact, the first agent (if she is risk-averse) will receive exactly $1/2$, while the second will usually get more.

What if there are *more than two* agents?

The Banach-Knaster Last-Diminisher Procedure

In the original paper on fair division, Steinhaus (1948) reports on a *proportional* procedure 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 the agents. Each agent either lets it pass (if she considers it too small) or she trims it down 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. Play cut-and-choose once $n = 2$. ✓

Each agent is guaranteed a *proportional* piece. Requires $O(n^2)$ cuts. May not be *contiguous* (unless you always trim "from the right").

H. Steinhaus. The Problem of Fair Division. *Econometrica*, 16:101–104, 1948.

The Even-Paz Divide-and-Conquer Procedure

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

(1) Ask each agent to cut the cake at her $\lfloor \frac{n}{2} \rfloor : \lceil \frac{n}{2} \rceil$ mark.

(2) Associate the union of the leftmost $\lfloor \frac{n}{2} \rfloor$ pieces with the agents who made the leftmost $\lfloor \frac{n}{2} \rfloor$ cuts, and the rest with the others.

(3) Recursively apply the same procedure to each of the two groups, until only a single agent is left. ✓

Each agent is guaranteed a *proportional* piece. Takes $O(n \log n)$ cuts.

Woeginger and Sgall (2007) later showed that we cannot do much better: $\Omega(n \log n)$ is a lower bound on the number of queries for any proportional procedure producing contiguous pieces.

S. Even and A. Paz. A Note on Cake Cutting. *Discrete Applied Mathematics*, 7(3):285–296, 1984.

G.J. Woeginger and J. Sgall. On the Complexity of Cake Cutting. *Discrete Optimization*, 4(2):213–220, 2007.

Preferences

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

- preferences are modelled as *utility functions*
- 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* we need to exchange when eliciting her preferences.

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*).

The Kannai-Peleg Theorem

Rather surprisingly, our requirements are impossible to satisfy:

Theorem 1 (Kannai and Peleg, 1984) *When there are $\geqslant$ 6 objects, no weak order $\hat{\succcurlyeq}$ satisfies both independence and dominance.*

<u>Proof:</u> 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 $A \setminus \{\max(A)\} \hat{\succ} \{\min(A)\}$ from (DOM), and then $A \hat{\succcurlyeq} \{\max(A), \min(A)\}$ from (IND).
> Show $\{\max(A), \min(A)\} \hat{\succcurlyeq} A$ analogously. ✓

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

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

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

Y. Kannai and B. Peleg. A Note on the Extension of an Order on a Set to the Power Set. *Journal of Economic Theory*, 32(1):172–175, 1984.

Automated Discovery of Theorems

A major challenges 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 ($\sim$Łoś-Tarski Theorem).
- 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. *J. Artif. Intell. Res.*, 40:143–174, 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

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 Frenchmen:	Wine $\succ$ Beer $\succ$ Milk
4 Dutchmen:	Milk $\succ$ Beer $\succ$ Wine

Which beverage *wins* the election for

- the plurality rule?
- plurality with runoff?
- the Borda rule?

Properties of Voting Rules

So how do you choose the right voting rule for your problem?

This really depends on your problem. No perfect *one-fits-all* solution.

- *Axioms:* When using a voting rule F as a *compromise-seeking* tool, we might want F to satisfy certain normative desiderata.
- *Truth:* When using a voting rule F as a *truth-finding* tool to aggregate the opinions of several experts, we may want F to maximise the probability of recovering some ground truth.
- *Complexity:* We might like a voting rule that is easy to compute.

Examples: The Normative Perspective

Recall that we had seen definitions of three different voting rules: *plurality*, *plurality with runoff*, and *Borda*.

What do you think, which one of them satisfy the following axioms?

- *Anonymity:* All voters should have the same "weight".
- *Monotonicity:* If the winning alternative $x^\star$ receives additional support (if some voters move $x^\star$ up in their preference orders), then $x^\star$ should still win the election.
- *Condorcet:* If alternative $x^\star$ is preferred to every other alternative by some strict majority, then $x^\star$ should win the election.
- *Reinforcement:* If alternative $x^\star$ wins in two disjoint districts, then $x^\star$ should also win when we join those two districts into one.
- *Strategyproofness:* No voter should ever have an incentive to lie about her preferences.

Coalition Formation and Matching

Agents may also have *preferences over each other* and we may be interested in finding the best way of partitioning the group.

Properties of interest include *stability*, *strategyproofness*, *fairness*.

These topics are covered in the course on Game Theory (so not here). If this is new to you, read the classic paper by Gale and Shapley.

D. Gale and L.S. Shapley. College Admissions and the Stability of Marriage. *American Mathematical Monthly*, 69:9–15, 1962.

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 $p \to q$.

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

- *What should be the collective decision?*

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 18(1):89–110, 2002.

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
- *coalition formation:* agents with preferences over each other
- *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)

▶ Read the introductory chapter of the *Handbook* to get a feel for the history and scope of the discipline.

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

Plan for the Rest of the Course

We'll focus on *voting theory*, cover this particular form of aggregation in depth, and see many of the techniques used in COMSOC exemplified in this specific domain ($\sim$10 lectures). Topics:

- formal framework of voting theory, and many voting rules
- axiomatic method: characterisation and impossibility results
- voting as truth-tracking (probabilistic methods)
- strategic behaviour (in the sense of game theory)
- complexity analysis of problems arising in voting theory
- research topics in voting inspired by computer science and AI

Plus a few one-off lectures on *other COMSOC topics*, probably just fair allocation of goods and judgment aggregation.

Plus a *tutorial on complexity theory* (for those who need it).

What next? Lots of voting rules and how to categorise them ...