Computational Social Choice: Spring 2015

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

Collective Decision Making

Social choice theory is the philosophical and mathematical study of methods for *collective decision making*.

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

- 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.

- Examples for problems and techniques in COMSOC research:
  - fair allocation of goods
  - voting in elections
  - two-sided matching
  - 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, 2009/2010.

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 first ever 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 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* 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*, 29(4):37–46, 2008.

Three Voting Rules

In voting, 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?

Axiomatic Method

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

The classical approach is to use the *axiomatic method:*

- identify good axioms: normatively appealing high-level properties
- give mathematically rigorous definitions of these axioms
- explore the consequences of the 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 A is the (sole or tied) winner and one voter switches from B to A, then A becomes the sole winner

K.O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 20(4):680–684, 1952.

Proof Sketch

We want to prove:

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

<u>Proof:</u> Clearly, the simple majority rule has all three properties. ✓

Other direction: assume *#voters* is *odd* (other case: similar) $\rightsquigarrow$ no ties

Let a $\mathcal{A}$ be the set of voters voting $A \succ B$ and $\mathcal{B}$ those voting $B \succ A$. *Anonymity* $\rightsquigarrow$ only number of ballots of each type matters. <u>Two cases:</u>

- Whenever $|\mathcal{A}| = |\mathcal{B}| + 1$ then *only A wins*. Then, by *PR*, A wins whenever $|\mathcal{A}| > |\mathcal{B}|$ (which is exactly the simple majority rule). ✓
- There exist $\mathcal{A}$, $\mathcal{B}$ with $|\mathcal{A}| = |\mathcal{B}| + 1$ but *b wins*. Let one A-voter switch to B. By *PR*, now only B wins. But now $|\mathcal{B}'| = |\mathcal{A}'| + 1$, which is symmetric to the first situation, so by *neutrality* A wins. ↯

The Condorcet Jury Theorem

The simple majority rule for two alternatives is attractive also in terms of *truth-tracking* (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 $p \cdot n > \frac{1}{2} \cdot n$. ✓

For a modern exposition, see Young (1995).

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

H.P. Young. Optimal Voting Rules. *J. Economic Perspectives*, 9(1):51–64, 1995.

Positional Scoring Rules

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

A *positional scoring rule* is given by a *scoring vector* $s = \langle s_1, \ldots, s_m \rangle$ 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_i points for every voter putting it at the ith position.

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

Examples:

- *Borda rule* = PSR with scoring vector $\langle m{-}1, m{-}2, \ldots, 0 \rangle$
- *Plurality rule* = PSR with scoring vector $\langle 1, 0, \ldots, 0 \rangle$
- *Antiplurality rule* = PSR with scoring vector $\langle 1, \ldots, 1, 0 \rangle$
- For any $k \leqslant m$, *k-approval* = PSR with $\langle \underbrace{1, \ldots, 1}_{k}, 0, \ldots, 0 \rangle$

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.

All PSR's Violate Condorcet

Consider the following example:

3 voters: $A \succ B \succ C$
2 voters: $B \succ C \succ A$
1 voter: $B \succ A \succ C$
1 voter: $C \succ A \succ B$

A is the *Condorcet winner*; she beats both B and C $4:3$. But any *positional scoring rule* makes *B* win (because $s_1 \geqslant s_2 \geqslant s_3$):

$$A: \quad 3 \cdot s_1 + 2 \cdot s_2 + 2 \cdot s_3$$
$$B: \quad 3 \cdot s_1 + 3 \cdot s_2 + 1 \cdot s_3$$
$$C: \quad 1 \cdot s_1 + 2 \cdot s_2 + 4 \cdot s_3$$

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*, 44(6):806–825, 1997.

Example: Strategic Manipulation

Suppose the *plurality rule* is used to decide an election: the candidate ranked first most often wins. Recall Florida in 2000 (simplified):

49%:	Bush $\succ$ Gore $\succ$ Nader
20%:	Gore $\succ$ Nader $\succ$ Bush
20%:	Gore $\succ$ Bush $\succ$ Nader
11%:	Nader $\succ$ <u>Gore</u> $\succ$ <u>Bush</u>

Bush will win this election. It would have been in the interest of the Nader supporters to pretend that they like Gore the most.

Thus, the plurality is subject to *strategic manipulation:* sometimes, some voters can get a better outcome by lying about their preferences.

▶ 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: A General Result. *Econometrica*, 41(4):587–601, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. *Journal of Economic Theory*, 10:187–217, 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 1960's, 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!*

Indeed, this works to some extent (but is still an underdeveloped area):

- Nipkow (2009) *verified* a known proof for G-S in ISABELLE.
- For related results, proofs have also been *derived* automatically, and some simpler results even have been *discovered* automatically.

T. Nipkow. Social Choice Theory in HOL. *J. Autom. Reas.*, 43(3):289–304, 2009.

P. Tang and F. Lin. Computer-aided Proofs of Arrow's and other Impossibility Theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.

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.

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*, 8(4):341–354, 1991.

The Stable Marriage Problem

We are given:

- n *men* and n *women*
- each has a linear *preference* ordering over the opposite sex

We seek:

- a *stable* matching of men to women: no man and woman should want to divorce their assigned partners and run off with each other

The Gale-Shapley Algorithm

Theorem 6 (Gale and Shapley, 1962) *There exists a stable matching for any combination of preferences of men and women.*

The *Gale-Shapley "deferred acceptance" algorithm* for computing a stable matching works as follows:

- In each round, each man who is not yet engaged proposes to his favourite amongst the women he has not yet proposed to.
- In each round, each woman picks her favourite from the proposals she's receiving and the man she's currently engaged to (if any).
- Stop when everyone is engaged.

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

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 $p \to q$.

	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?

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
- *matching:* two groups of 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)

F. Brandt, V. Conitzer, and U. Endriss. Computational Social Choice. In G. Weiss (ed.), *Multiagent Systems*. MIT Press, 2013.

Plan for the Rest of the Course

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

- formal frameworks for judgment aggregation, comparison
- design of aggregation methods based on various principles
- axiomatic method: characterisation and impossibility results
- (maybe) 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 *other COMSOC topics*, e.g.:
fair allocation of goods, voting in elections, two-sided matching, ...

Plus one *tutorial on complexity theory*, if needed (it probably is).

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%) and mini-project (50%).

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

`http://www.illc.uva.nl/~ulle/teaching/comsoc/2015/`

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

Homework

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

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!)

Submission: Type up your solutions (LaTeX strongly preferred). Hand them in on paper (not electronically) before the start of class.

Grading: Each question is graded out of 10 (usual Dutch scale). I'll disregard the two questions you did worst on.

Deadlines: They are strict.

Mini-Projects

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

- identify an interesting paper on JA not covered in class and fill in some gaps or come up with an extension or a generalisation
- come up with a new rule / axiomatisation / complexity result / ...
- explore an application domain for JA: could be a literature review, an idea for a new application, or an experimental study
- show how to embed another social choice framework into JA
- implement one or several aggregation rules

The purpose of this is to give you 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.

Time Commitment

You will need to be able to devote an average of *20 hours/week*, over the full eight weeks, to this course.

Usually *two classes per week*, sometimes three (see website).

No formal *attendance* requirements, but in practice you are expected to show up and to participate actively.

There'll be homework *almost every week*, but it'll be very light during the second half of the course, to leave time for work on the projects.

You need to be available during the *exam week* for the presentations.

The deadline for the paper will at the *end of the block*.

What next?

- Look over the course *website* and read through some of the suggested *literature* references.
- The next lecture will be a systematic introduction to basic *judgment aggregation*, focussing on the *axiomatic method*.