Introduction to Computational Social Choice

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

Social Choice Theory

SCT studies collective decision making: how should we aggregate the preferences of the members of a group to obtain a "social preference"?

$$\triangle \succ_1 \bigcirc \succ_1 \square$$

$$\square \succ_2 \triangle \succ_2 \bigcirc$$

$$\bigcirc \succ_3 \square \succ_3 \triangle$$

$$?$$

SCT is traditionally studied in Economics and Political Science, but now also by "us": *Computational Social Choice*.

Talk Outline

- Computational Social Choice: *research area* and *community*
- Examples for typical research questions (mostly Amsterdam)
- Conclusions and how to find out more

Computational Social Choice

Social choice theory studies mechanisms for collective decision making, such as voting procedures or protocols for fair division.

- Precursors: Condorcet, Borda (18th century) and others
- serious scientific discipline since 1950s
- Classics: Black, Arrow, May, Sen, Gibbard, Satterthwaite, ...

Computational social choice adds a computational perspective to this, and also explores the use of concepts from social choice in computing.

- "classical" papers: ∼1990 (Bartholdi et al.)
- active research area with regular contributions since ∼2002
- name "COMSOC" and biannual workshop since 2006

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. *A Short Introduction to Computational Social Choice*. Proc. SOFSEM-2007.

The COMSOC Research Community

- International Workshop on Computational Social Choice:
  - 1st edition: COMSOC-2006 in Amsterdam, December 2006
    48 paper submissions and 80 participants (14 countries)
  - 2nd edition; COMSOC-2008 in Liverpool, September 2008
    55 paper submissions and ∼80 participants (∼20 countries)
  - 3rd edition: COMSOC-2010 in Düsseldorf, September 2010
    Paper submission deadline: 15 May 2010
- Special issues in international journals:
  - *Mathematical Logic Quarterly*, vol. 55, no. 4, 2009
  - *Journal of Autonomous Agents and Multiagent Systems*, 2010
- Journals and conferences in AI, MAS, TCS, Logic, Econ, . . .
- COMSOC website: `http://www.illc.uva.nl/~ulle/COMSOC/`
  (workshop proceedings, related events, mailing list, etc.)

Example from Voting

Suppose the *plurality rule* is used to decide an election: the candidate receiving the highest number of votes wins.

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

$$\begin{array}{ll} 49\%: & \text{Bush} \succ \text{Gore} \succ \text{Nader} \\ 20\%: & \text{Gore} \succ \text{Nader} \succ \text{Bush} \\ 20\%: & \text{Gore} \succ \text{Bush} \succ \text{Nader} \\ 11\%: & \text{Nader} \succ \text{Gore} \succ \text{Bush} \end{array}$$

So even if nobody is cheating, Bush will win this election. <u>But:</u>

- In a *pairwise contest*, Gore would have defeated anyone.
- It would have been in the interest of the Nader supporters to *manipulate*, i.e., to misrepresent their preferences.

Is there a better voting procedure that avoids these problems?

Some Voting Procedures

- *Plurality:* elect the candidate ranked first most often
- *Borda:* each voter gives $m-1$ points to the candidate they rank first, $m-2$ to the candidate they rank second, etc., and the candidate with the most points wins
- *Copeland:* award 1 point to a candidate for each pairwise majority contest won and $\frac{1}{2}$ points for each draw, and elect the candidate with the most points
- *Single Transferable Vote (STV)*: keep eliminating the plurality loser until someone has an absolute majority
- *Approval:* voters can approve of as many candidates as they wish, and the candidate with the most approvals wins

Complexity as a Barrier against Manipulation

By the *Gibbard-Satterthwaite Theorem*, any voting rule for choosing between ≥ 3 candidates can be manipulated (unless it is dictatorial).

Idea: So it's always *possible* to manipulate, but maybe it's *difficult*! Tools from *complexity theory* can be used to make this idea precise.

- For the *plurality rule* this does *not* work: if I know all other ballots and want X to win, it is *easy* to compute my best strategy.
- But for *single transferable vote* it does work. Bartholdi and Orlin showed that manipulation of STV is *NP-complete*.

Recent work in COMSOC has expanded on this idea:

- NP is a worst-case notion. What about average complexity?
- Also: complexity of winner determination, control, bribery ...

J.J. Bartholdi III and J.B. Orlin. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

Preferences and Ballots

Two common assumptions in voting theory:

- Voters have *preferences* that are *total orders* over candidates.
- Voters vote by submitting a structure just like their preferences, truthfully or not (*ballots* and preferences have *the same* structure).

We may want to drop these assumptions, because:

- For lack of information or processing resources, voters may be *unable to rank* all candidates (in their mind or on the ballot sheet).
- To reduce *complexity of communication*, we may want to design voting rules that work with ballots of bounded size.
- For *approval voting*, ballots cannot be encoded using total orders.

Beyond Classical Voting Theory

In recent work we have proposed a model where:

- *preferences* and *ballots* can be different structures; and
- a notion of *sincerity* replaces the standard notion of *truthfulness* (because the ballot language may *not allow* you to be truthful).
- Now you can get positive results for certain combinations:
  - Under *approval voting* with standard preferences, you can never benefit from not voting sincerely.
  - If you have *dichotomous preferences*, you can never benefit from not voting sincerely for a wide range of voting procedures.
  - Voting sincerely and effectively is *computationally tractable* in above scenarios.

U. Endriss, M.S. Pini, F. Rossi, and K.B. Venable. *Preference Aggregation over Restricted Ballot Languages: Sincerity and Strategy-Proofness*. Proc. IJCAI-2009.

Arrow's Impossibility Theorem

It seems reasonable to require a *social welfare function* (SWF), mapping profiles of individual preference orderings to a social preference ordering, to satisfy the following axioms:

- *Unanimity* (UN): if every individual prefers alternative x over alternative y, then so should society
- *Independence of Irrelevant Alternatives* (IIA): social preference of x over y should only depend on individual pref's over x and y
- *Non-Dictatorship* (ND): no single individual should be able to impose a social preference ordering

Theorem 1 (Arrow, 1951) *For three or more alternatives, there exists no SWF that satisfies all of (UN), (IIA) and (ND).*

K.J. Arrow. *Social Choice and Individual Values*. 2nd edition, Wiley, 1963.

Formal Verification of Arrow's Theorem

Logic has long been used to *formally specify* computer systems, facilitating formal or even *automatic verification* of various properties. Can we apply this methodology also to *social choice* mechanisms?

Tang and Lin (2009) show that the *"base case"* of Arrow's Theorem with 2 agents and 3 alternatives can be fully modelled in *propositional logic:*

- Automated theorem provers can verify $\textsc{Arrow}(2,3)$ to be correct in < 1 second — that's $(3!)^{3! \times 3!} \approx 10^{28}$ SWFs to check
- Opens up opportunities for quick sanity checks of hypotheses regarding new possibility and impossibility theorems.

Our own work using *first-order logic* tries to go beyond such base cases.

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

U. Grandi and U. Endriss. *First-Order Logic Formalisation of Arrow's Theorem.* Proc. LORI-2009.

Social Choice in Combinatorial Domains

Many social choice problems have a *combinatorial structure:*

- Elect a *committee* of k members from amongst n candidates.
- Find a fair *allocation* of n indivisible goods to agents.

Seemingly small problems generate huge numbers of alternatives:

- Number of 3-member committees from 10 candidates: $\binom{10}{3} = 120$ (i.e., $120! \approx 6.7 \times 10^{198}$ possible rankings)
- Allocating 10 goods to 5 agents: $5^{10} = 9765625$ allocations and $2^{10} = 1024$ bundles for each agent to think about

Conclusion: We need good *languages* for representing preferences!

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

Weighted Goals

A compact representation language for modelling utility functions (cardinal preferences) over products of binary domains —

<u>Notation:</u> finite set of propositional letters PS; propositional language $\mathcal{L}_{PS}$ over PS to describe requirements, e.g.:

$$p, \quad \neg p, \quad p \wedge q, \quad p \vee q$$

A *goalbase* is a set $G = \{(\varphi_i, \alpha_i)\}_i$ of pairs, each consisting of a (consistent) propositional formula $\varphi_i \in \mathcal{L}_{PS}$ and a real number α_i. The utility function u_G generated by G is defined by

$$u_G(M) \quad = \quad \sum \{\alpha_i \mid (\varphi_i, \alpha_i) \in G \text{ and } M \models \varphi_i\}$$

for all models $M \in 2^{PS}$. G is called the *generator* of u_G.

Different syntactic restrictions give different representation languages.

Some Results

Examples from our research on weighted goal languages:

- *Expressivity:* If all formulas and weights are positive, then we can express all monotonic utility function, and only those.
- *Succinctness:* Conjunctions of literals can express the same functions as general formulas, but do so strictly less succinctly.
- *Complexity:* Finding the most preferred model is NP-hard in general, but in $O(n \log n)$ if all formulas are literals.
- Applications: *combinatorial auctions* and *expressive voting*

J. Uckelman. *More than the Sum of its Parts: Compact Preference Representation over Combinatorial Domains*. PhD thesis, ILLC, University of Amsterdam, 2009.

J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang. Representing Utility Functions via Weighted Goals. *Mathematical Logic Quarterly*, 55(4):341–361, 2009.

Judgment Aggregation

Preferences are not the only structures that we may wish to aggregate. JA studies the aggregation of judgments on related propositions.

	p	$p \to q$	q
Agent 1:	Yes	Yes	Yes
Agent 2:	No	Yes	No
Agent 3:	Yes	No	No
Majority:	Yes	Yes	No

While each individual set of judgments is logically consistent, the collective judgment produced by the majority rule is not.

Research issues: impossibility theorems; characterisation of admissible agendas; proposals for "good" aggregation procedures; . . .

C. List and C. Puppe. Judgment Aggregation: A Survey. *Handbook of Rational and Social Choice*. Oxford University Press, 2009.

Complexity of Judgment Aggregation

What about computational considerations in JA?

In recent work we address the following questions:

- *Safety of the Agenda:* Given an agenda Φ (set of propositions), can we guarantee that any aggregation procedure belonging to a given class of procedures (characterised via some axioms) will never "produce a paradox"?
- What is the computational complexity of deciding SoA? (turns out to be Π_2^p-complete for all interesting axioms)

U. Endriss, U. Grandi, and D. Porello. *Complexity of Judgment Aggregation*. Working Paper, ILLC, University of Amsterdam, 2009.

Other Topics

Social choice theory is not just about voting and preferences.
Aggregation also plays a role in other domains, e.g.:

- Multiagent Resource Allocation and Fair Division
- Mechansim Design
- Stable Matchings

Computational Social Choice

We have seen several examples for work in COMSOC.

Research can be broadly classified along two dimensions —

The kind of *social choice problem* studied, e.g.:

- aggregating individual judgements into a collective verdict
- electing a winner given individual preferences over candidates
- fairly dividing a cake given individual tastes

The kind *computational technique* employed, e.g.:

- algorithm design to implement complex mechanisms
- complexity theory to understand limitations
- logical modelling to fully formalise intuitions
- knowledge representation techniques to compactly model problems
- deployment in a multiagent system

Last Slide

- COMSOC is an exciting area of research bringing together ideas from mathematical economics (particularly social choice theory) and computer science (including logic, AI, MAS, TCS).
- COMSOC website: `http://www.illc.uva.nl/~ulle/COMSOC/`
- You are welcome to attend the COMSOC Seminar at the ILLC (typically, every 2-3 Fridays at 4pm)
- There's a Workshop on Preferences at the ILLC *this* Friday at 2pm (speakers: Jérôme Lang, Francesca Rossi, Mike Wooldridge et al.)
- Papers are on my website (including the surveys below).

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. *A Short Introduction to Computational Social Choice*. Proc. SOFSEM-2007.

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