Automated Analysis of Social Choice Problems: Approval Elections with Small Fields of Candidates

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

Talk Outline

- The Field: *Computational Social Choice*
- The Problem: strategic behaviour in *Approval Voting*
- The Approach: *automated analysis* of all relevant cases
- The Result: *positive*, (only) under very specific conditions

Computational Social Choice

Social choice theory asks: how should we aggregate the preferences of the members of a group to obtain a "social preference"?

Social choice is useful for AI/CS:

- to aggregate beliefs / coordinate actions in a *multiagent system*
- to aggregate the output of several Internet *search engines*

AI/CS is useful for social choice:

- to compactly model preferences (*knowledge representation*)
- to verify correctness of formal results (*automated reasoning*)
- to understand limitations (*complexity theory*)

COMSOC specifically focusses on such computational concerns.

The Problem

Set of *voters*. Set of *candidates*. Each voter has a *preference order* (= linear order) over the candidates.

Approval Voting: each voter *approves* of some of the candidates; candidate with the most approvals wins the election.

Call a voter's ballot *sincere* *iff* she really prefers everyone she approves of over everyone she disapproves of [compare to *truthfulness*].

Suppose a voter has found out how the others will vote. Can we be sure she has not incentive to vote insincerely?

An election could be tied (*several winners*). So, to define "incentives", we must extend our voter's preferences over individual candidates to *preferences over sets of candidates*. $\rightsquigarrow$ *preference extension principles*

Example

Election with 4 *candidates:* a, b, c, d. My preference: $a > b > c > d$.
Suppose I'm subject to the *Gärdenfors principle* of preference extension:

- I prefer set $X \cup \{y\}$ to X if I prefer candidate y to every $x \in X$.
- I prefer set X to $X \cup \{y\}$ if I prefer every $x \in X$ to candidate y.

Now, suppose I know how the others voted. Before I vote, the scores are:

$$a\text{: } 9 \quad b\text{: } 10 \quad c\text{: } 9 \quad d\text{: } 10$$

What are my options?

- Have $\{a, b, d\}$ win by voting $\{a\}$ (sincere)
- Have $\{b\}$ win by voting, say, $\{a, b\}$ (sincere)
- Have $\{a, b, c, d\}$ win by voting $\{a, c\}$ (*insincere*)
- Have $\{b, d\}$ win by voting, say, $\{a, b, c, d\}$ (sincere)
- Have $\{b, c, d\}$ win by voting $\{c\}$ (insincere)
- Have $\{d\}$ win by voting, say, $\{d\}$ (insincere)

Under Gärdenfors, one of first three must be the best (that's all we know!).
If $\{a, b, c, d\}$ is *actually* the best, I have an incentive to vote insincerely!

The Approach

For a fixed number of candidates (and any number of voters), the number of "situations" to check is (large but) finite.

Explore all of them systematically using a computer program!

Discussion:

- Acceptable as a method of proof if the correctness of the program can be verified easily (it's a small logic program: so, ok).
- Interesting side issue is how to implement a "theorem prover" for deciding $X \succ Y$, given a preference extension principle and a voter's preferences over individual candidates.

The Result

We found that when the *number of candidates is small*, then even for very *weak preference extension principles*, it is the case that voters have *no (or few) incentives to vote insincerely*. For instance:

- No voter conforming to the *Gärdenfors principle* can benefit from voting insincerely in elections with $\leqslant 3$ *candidates*.
- For the *Gärdenfors principle* and 4 *candidates* our earlier example is one of *only two* problematic cases (out of 65 "situations").

This complements earlier results for *arbitrary numbers of candidates* and *strong preference extension principles*.

U. Endriss. Sincerity and Manipulation under Approval Voting. *Theory and Decision*. In press (2012).

Conclusion

- Better understanding of *incentives to strategise* in *approval voting*. Results depend on *preference extension principle*.
- Our simple *method* of automated analysis of voting situations complements existing (but so far still rather scarce) work on using *automated reasoning* in *computational social choice*.