Computational Social Choice 2023

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

[http://www.illc.uva.nl/~ulle/teaching/comsoc/2023/]

Plan for Today

Complexity theory is a core tool in computational social choice.

Today we briefly review some representative examples of its use:

- Winner determination
- Strategic manipulation
- Possible winners

Winner Determination

The most fundamental algorithmic problem in voting theory is that of determining the winners under a given voting rule for a given profile.

Here's what we would expect to find:

- Winner determination is easy (polynomial-time computable) for rules such as Plurality, Veto, Borda, Copeland, ...
- Winner determination seems difficult (at least NP-hard) for rules such as Slater and Kemeny

Exercise: *How would you define this winner determination problem?*

The Winner Determination Problem

For a basic complexity analysis, best to state this as a decision problem:

WDP(F)
Input: Profile $\boldsymbol{R} \in \mathcal{L}(A)^n$ of preferences and alternative $x^\star \in A$
Question: Is it the case that $x^\star \in F(\boldsymbol{R})$?

Exercise: *Explain how this relates to the problem of computing* $F(\boldsymbol{R})$.

Exercise: *Why is asking "*$X^\star = F(\boldsymbol{R})$*?" for* $X^\star \subseteq A$ *a bad idea?*

Examples

Let's briefly look a three examples highlighting interesting features of the winner determination problem ...

The topic is treated in depth in Chapters 3–5 of the *Handbook*.

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

Complexity of the Banks Rule

Under *Banks*, an alternative wins if it is a top element in an inclusion-maximal *acyclic subgraph* of the *majority graph* induced by the profile.

We state this (not very surprising) result without proof:

Theorem (Woeginger, 2003): $\mathrm{WDP}(\mathit{Banks})$ *is NP-complete.*

NP-membership obvious. NP-hardness from $\textsc{Graph 3-Colouring}$.

G.J. Woeginger. Banks Winners in Tournaments are Difficult to Recognize. *Social Choice and Welfare*, 2003.

Easiness of Computing Some Banks Winner

We have seen that checking whether x is a Banks winner is NP-hard. So computing *all* Banks winners is also NP-hard.

But computing just *some* Banks winner is easy! Algorithm:

(1) Let $S := \{a_1\}$ and $i := 1$ (alternatives $A = \{a_1, \ldots, a_m\}$)

(2) While $i < m$, repeat:

- Let $i := i + 1$
- If the majority graph restricted to $S \cup \{a_i\}$ is acyclic, then let $S := S \cup \{a_i\}$

(3) Return the top element in S (it is a Banks winner)

O. Hudry. A Note on "Banks Winners in Tournaments are Difficult to Recognize" by G.J. Woeginger. *Social Choice and Welfare*, 2004.

Complexity of Ranked Pairs

Under *Ranked Pairs*, we order the ranked pairs in $A \times A$ by the number of voters supporting them; then lock in pairs in this order but skip pairs that would create cycles; and finally elect the top alternative.

Need to specify how we *break ties* in this order on pairs.

Fact: *For lexicographic tie-breaking,* $\mathrm{WDP}($*RankedPairs*$)$ *is in P.*

Observe that lexicographic tie-breaking violates neutrality. A more principled approach would be *parallel-universe tie-breaking*. But:

Theorem (Brill and Fisher, 2012): *For parallel-universe tie-breaking,* $\mathrm{WDP}($*RankedPairs*$)$ *is NP-complete.*

NP-memb. obvious (witness: order on pairs). NP-hardness from SAT.

M. Brill and F. Fischer. The Price of Neutrality for the Ranked Pairs Method. AAAI-2012.

Complexity of the Dodgson Rule

Dodgson (a.k.a. Lewis Carroll, author of *Alice in Wonderland*) proposed:

> *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.*

Well over 100 years after its invention, the WDP of this rule turned out to be compete for a complexity class thought to lack natural problems:

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

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.

Complexity as a Barrier to Manipulation

Every voting rule can be manipulated in some profiles. But even when it is *possible* to manipulate, maybe actually doing so is *difficult*?

If manipulation is computationally intractable for F, then F might be considered *resistant* (albeit still not *immune*) to manipulation.

Most interesting when winner determination is tractable. At the very least, we would like to see a *complexity gap* between manipulation (undesired behaviour) and winner determination (desired functionality).

V. Conitzer and T. Walsh. Barriers to Manipulation in Voting. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. CUP, 2016.

Classical Results

The seminal paper by Bartholdi, Tovey and Trick (1989) starts by showing that manipulation is in fact *easy* for a range of commonly used voting rules, and then presents one system (a variant of the Copeland rule) for which manipulation is NP-complete. Next:

- We first present a couple of these easiness results, namely for *plurality* and for the *Borda rule*.
- We then mention a result from a follow-up paper by Bartholdi and Orlin (1991): the manipulation of *STV* is *NP-complete*.

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare*, 1989.

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

Manipulability as a Decision Problem

Need to formulate manipulability as a decision problem:

> $\textsc{Manip}(F)$
> **Input:** Ballots for all but one voter and alternative $x^\star$.
> **Question:** Is there a ballot for the final voter such that $x^\star$ wins?

To find out what the best winner achievable for the manipulator is, she has to solve $\textsc{Manip}(F)$ for all $x^\star$, in order of her preference.

Remark: This formulation of the *decision problem* cannot be used to solve the *search problem* of computing the manipulating ballot. As our focus here is on intractability results, this is ok.

Remark: We assume that the *manipulator knows* all the other ballots. This unrealistic assumption is intentional: if manipulation is hard even under such favourable conditions, then all the better.

Manipulating the Plurality Rule

Recall: under *plurality*, the alternative(s) ranked first most often win(s).

The plurality rule is easy to manipulate (trivial):

- Simply vote for $x^\star$, the alternative to be made winner by means of manipulation. If manipulation is possible at all, this will work. Otherwise manipulation is not possible.

Thus: $\textsc{Manip}(\textit{plurality})$ can be decided in *polynomial* time.

General: $\textsc{Manip}(F) \in \mathrm{P}$ for any rule F with polynomial winner determination problem and polynomial number of ballots.

Manipulating the Borda Rule

Recall: under *Borda*, you submit a ranking of all alternatives and thereby award $m-k$ points to the alternative ranked in position k.

Remark: We now have superpolynomially-many possible ballots.

But Borda still is easy to manipulate. Use a *greedy algorithm:*

- Place $x^\star$ (the alternative to be made winner through manipulation) at the top of your ballot.
- Then inductively proceed as follows: Check if any of the remaining alternatives can be put next on the ballot without preventing $x^\star$ from winning. If yes, do so. (If no, manipulation is impossible.)

After convincing ourselves that this algorithm is indeed correct, we see that $\textsc{Manip}(\mathit{Borda})$ can be decided in *polynomial* time.

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare*, 1989.

Intractability of Manipulating STV

Recall: *Single Transferable Vote* (STV) works by eliminating plurality losers until an alternative is ranked first by $> 50\%$ of the voters.

Theorem (Bartholdi and Orlin, 1991): $\textsc{Manip}(STV)$ *is NP-compl.*

Proof: Omitted. But try to get an intuition for why this is intractable.

For example, it is often not optimal to put the alternative x you want to win at the top of your ballot (by ranking y at the top, you may be able to eliminate z, which may be a stronger competitor than y).

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

Coalitional Manipulation

It rarely is the case that a *single* voter really can make a difference.
So we should look into *manipulation by a coalition* of voters.

Variants of the problem:

- Ballots may be *weighted* or *unweighted*.

  Examples: countries in the EU, shareholders of a company

- Manipulation may be *constructive* (making alternative $x^\star$ win) or *destructive* (ensuring $x^\star$ does not win).

Decision Problems

Next, we consider two decision problems, for a given voting rule F:

ConstructiveManipulability(F)
Input: List of weighted ballots; set of weighted manipulators; $x^\star \in A$.
Question: Are there ballots for the manipulators such that $x^\star$ wins?

DestructiveManipulability(F)
Input: List of weighted ballots; set of weighted manipulators; $x^\star \in A$.
Question: Are there ballots for the manipulators such that $x^\star$ loses?

Constructive Manipulation under Borda

In the context of coalitional manipulation with weighted voters, we can get hardness results for elections with small numbers of alternatives:

Theorem (Conitzer et al., 2007): *For Borda, constructive coalitional manipulation with weighted voters is NP-complete for $\geqslant 3$ alternatives.*

Proof: We have to prove NP-membership and NP-hardness:

- NP-membership: easy (if you guess ballots for the manipulators, we can check that it works in polynomial time)
- NP-hardness: for three alternatives by reduction from PARTITION (next slide); hardness for more alternatives follows

V. Conitzer, T. Sandholm, and J. Lang. When are Elections with Few Candidates Hard to Manipulate? *Journal of the ACM*, 2007.

Proof of NP-hardness

We use a reduction from the NP-complete Partition problem:

Partition
Input: $(w_1, \ldots, w_n) \in \mathbb{N}^n$
Question: Is there a set $S \subseteq \{1, \ldots, n\}$ s.t. $\sum_{i \in S} w_i = \frac{1}{2} \sum_{i=1}^{n} w_i$?

Let $K := \sum_{i=1}^{n} w_i$. Given an instance of Partition, we construct an election with $n + 2$ weighted voters and three alternatives:

- two voters with weight $\frac{1}{2}K - \frac{1}{4}$, voting $(a \succ b \succ c)$ and $(b \succ a \succ c)$
- a coalition of n voters with weights $w_1, \ldots, w_n$ who want c to win

Clearly, each manipulator should vote either $(c \succ a \succ b)$ or $(c \succ b \succ a)$. Suppose there does exist a partition. Then they can vote like this:

- manipulators corresponding to elements in S vote $(c \succ a \succ b)$
- manipulators corresponding to elements outside S vote $(c \succ b \succ a)$

Scores: $2K$ for c; $\frac{1}{2}K + (\frac{1}{2}K - \frac{1}{4}) \cdot (2 + 1) = 2K - \frac{3}{4}$ for both a and b
If there is no partition, then either a or b will get at least 1 point more.
Hence, manipulation is feasible iff there exists a partition. ✓

Destructive Manipulation under Borda

Theorem (Conitzer et al., 2007): *For Borda, destructive coalitional manipulation with weighted voters is in P.*

Proof: Let $x^\star$ be the alternative the manipulators want to lose.

For every $y \neq x^\star$, simply try everyone ranking y at the top and $x^\star$ at the bottom. If none of these $m - 1$ attempts work, nothing will. ✓

V. Conitzer, T. Sandholm, and J. Lang. When are Elections with Few Candidates Hard to Manipulate? *Journal of the ACM*, 2007.

Critique of the Approach

Such complexity results provide interesting insights into the dynamics of strategic manipulation. *But do they really offer protection?*

NP-hardness is a *worst-case* notion and cannot rule out the possibility that problem instances enountered in practice are easy to solve.

Research suggests that it might be impossible to find a voting rule that is *usually* hard to manipulation—for a some definition of "usual".

V. Conitzer and T. Walsh. Barriers to Manipulation in Voting. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. CUP, 2016.

Possible Winners

Suppose we have not yet elicited all individual preferences in full.

Possible Winner Problem: *Is there a way in which to complete the preferences so that $x^\star$ wins under rule F?*

Important problem due to the rich variety of possible interpretations:

- Some *individuals* did not vote yet.
- Some new *alternatives* have entered the field (say, new proposals).
- When *eliciting* preferences (costly!), we can stop once *possible* = *necessary winners* (alternatives winning for *all* completions).

Exercise: *How does this relate to coalitional manipulation?*

K. Konczak and J. Lang. *Voting Procedures with Incomplete Preferences*. Proc. Multidisciplinary Workshop on Advances in Preference Handling 2005.

C. Boutilier and J.S. Rosenschein. Incomplete Information and Communication in Voting. In F. Brandt et al. (eds.), *Handbook of COMSOC*. CUP, 2016.

Complexity of the Possible Winner Problem

Rich literature on the complexity of the possible winner problem.

$\textsc{PossWin}(F)$
Input: Profile $\boldsymbol{R}$ of partial ballots and alternative $x^\star$.
Question: Is $x^\star$ a possible winner for $\boldsymbol{R}$ under voting rule F?

Two examples (without proof) to give an impression:

Theorem (Betzler and Dorn, 2010): $\textsc{PossWin}$(*plurality*) *is in P.*

Theorem (Xia and Conitzer, 2008): $\textsc{PossWin}$(*Borda*) *is NP-compl.*

N. Betzler and B. Dorn. Towards a Dichotomy for the Possible Winner Problem in Elections Based on Scoring Rules. *J. Computer and System Sciences*, 2010.

L. Xia and V. Conitzer. Determining Possible and Necessary Winners under Common Voting Rules Given Partial Orders. AAAI-2008.

Summary

We saw some representative examples for the use of complexity theory in the analysis of voting rules, for three different scenarios:

- Winner determination
- Strategic manipulation
- Possible winners

It's generally *good practice*, for any new problem you put forward, to also analyse its computational complexity.

At the same time, it's important not to read too much into (negative) complexity results. *Is worst case relevant? Might heuristics work?*

What next? Social choice in richer models of decision making.