

Realistic Assumptions for Attacks on Elections

Zack Fitzsimmons

College of Computing and Information Sciences
Rochester Institute of Technology
Rochester, NY 14623, USA
zmf6921@rit.edu

Abstract

We must properly model attacks and the preferences of the electorate for the computational study of attacks on elections to give us insight into the hardness of attacks in practice. Theoretical and empirical analysis are equally important methods to understand election attacks. I discuss my recent work on domain restrictions on partial preferences and on new election attacks. I propose further study into modeling realistic election attacks and the advancement of the current state of empirical analysis of their hardness by using more advanced statistical techniques.

Models for Election Attacks

In computational social choice we apply techniques from computer science to understand social choice problems such as elections, which are a way to reach a fair decision when presented with the preferences of several agents. However, elections can be vulnerable to voters misrepresenting their preferences (manipulation) or even attacks on the structure of the election itself (control).

It is obvious to suggest that we should use election systems where these attacks are not possible. Unfortunately, a crucial negative result, the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975), states that every reasonable election system is manipulable. While it is not possible to design an election system that is impossible to manipulate it may be computationally infeasible to determine if a manipulation exists. Bartholdi et al. (1989) introduced the concept of measuring the resistance of elections to manipulation using computational complexity. Since then there has been a focus on determining the worst-case complexity of manipulation and other attacks on different election systems (see, e.g., Faliszewski et al. (2010)) and more recently how hard they are in practice through experimental means (see, e.g., Walsh (2011)).

The problems of manipulation and control have been well-studied individually. However, to better model real-world scenarios, other attacks can and have been developed that have aspects of manipulation and control or are logical extensions of them. One extension is to explore multiple attacks happening in the same election, which is more

likely in practice than one isolated attack with all other voters and election organizers acting honestly. This was the topic of the paper that I presented at IJCAI-13 (Fitzsimmons, Hemaspaandra, and Hemaspaandra 2013). This paper explored elections where there is both an election chair controlling the election and a subset of voters manipulating the election. The chair and the manipulators may act either collaboratively or competitively and this has interesting effects on the corresponding worst-case complexity.

In another paper we explored the complexity of manipulating two-stage elections when the same election system is used at each stage (Fitzsimmons, Hemaspaandra, and Hemaspaandra 2014). The main motivation was that when there are multiple winners in an election it is reasonable to assume that a runoff election would be held among the winners. Like the aforementioned combination of manipulation and control, the aim of this research is to present a model that is likely to occur in real-world situations.

Election attacks are not only affected by the structure of the attack, but also the behavior of the voters. One common assumption is that a given electorate satisfies a domain restriction such as single peakedness.

The notion of single-peaked preferences introduced by Black (1948) is the most commonly studied domain restriction on voters' preferences in an election. Single-peaked preferences model the preferences of voters with respect to a polarizing issue where the candidates can be arranged with respect to a one-dimensional axis where the leftmost and rightmost positions of the axis represent the extremes of the issue. When an election has voters with single-peaked preferences the worst-case complexity of the manipulation and control problems often decreases (Faliszewski et al. 2011).

In real-world elections voters often have some degree of partial preference and this should be properly considered when determining if a collection of voters satisfies a certain domain restriction. Recently single-peaked preferences were examined for preference profiles of partial votes in an existential model (Lackner 2014). I expanded on this work in a recent technical report (Fitzsimmons 2014) and showed that single-peaked consistency for weak orders in this existential model is in P, solving the main open problem in Lackner (2014). Additionally, I showed that the two other definitions for single-peaked consistency for weak orders are each also in P (Fitzsimmons 2014) and returned to more estab-

lished models of single-peakedness for partial preferences.

The current models for manipulation, control, and assumptions on the structure of the preferences of the voters in an election must be extended to better model real-world scenarios, but must be within a reasonable scope to study.

Analysis of Election Attacks

Theoretical and empirical analysis of hard election problems must both be done to better understand the worst-case and in-practice hardness of attacks on elections.

Theoretical results are often either polynomial-time algorithms or proofs of NP-hardness, while a large proportion of the empirical studies published in computational social choice have similar design to the important early work by Walsh (2011). Heuristic and/or approximation algorithms are run a large number of times on votes sampled from different statistical cultures or from real-world data. Descriptive statistics are then gathered such as the average time required and the sample probability that manipulation is possible for different combinations of the voter distribution, the size of the candidate and voters sets, and the number of manipulators. Graphs are generated that illustrate general trends and some conclusions are drawn. This experimental approach fails to assess the significance of the predictor variables (e.g., voter distribution, size of the candidate and voter sets) with respect to the response (either time required or manipulation possible/not possible). If a more rigorous experimental design is used then stronger conclusions can be made. For example, instead of being limited to stating that the time required by a given algorithm increases in an experiment as the size of the candidate set increases, we would be able to test if the size of the candidate set causes a statistically significant difference in the time required by the algorithm in general.

We will expand on the current state of empirical analysis of heuristic algorithms for different attacks on elections in several ways. We will use experimental design techniques such as analysis of variance (ANOVA) to determine which values for our predictors results in a significant difference in the time required by a given algorithm. In the case of interpreting the frequency that an attack is possible, we can use multinomial regression models or other classification methods such as decision trees to fit the data.

These suggested techniques will result in better inference into the behavior of a given algorithm when presented with different inputs compared to simple descriptive statistics. When the inputs are real-world data, or generated from reasonable assumptions, we gain inference into the in-practice hardness and the frequency that a certain election attack is possible.

Research Plan

I see worst-case complexity analysis and empirical analysis of attacks on elections as two aspects that should be examined side-by-side for new and existing problems. I intend on continuing to examine how statistical analysis applied to the experimental study of elections could lead to a deeper understanding of what factors influence the in-practice hardness of

attacks on elections. I will analytically examine the effect of voter distributions, models for preferences, and the relative vulnerability that an election has to different attacks. New variants of election attacks can each be applied to known election systems or variants of these election systems to discover where each of these cases are computationally easy or hard in the worst case and in practice.

Before AAAI-15 I intend on exploring partial single-peakedness in the datasets containing partial votes found on PREFLIB (Mattei and Walsh 2013) and expanding the computational study of domain restrictions from social choice literature, focusing on partial preferences. Additionally, I intend on exploring new variants of election attacks both theoretically and empirically using the proposed methods.

The theoretical and empirical analysis of hard election problems under realistic assumptions will advance our understanding of the in-practice hardness of attacks on elections and result in a better understanding of the conditions that make these problems hard.

Acknowledgments: This work was supported in part by NSF grant no. CCF-1101452 and the NSF Graduate Research Fellowship under NSF grant no. DGE-1102937.

References

- Bartholdi, III, J.; Tovey, C.; and Trick, M. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6(3):227–241.
- Black, D. 1948. On the rationale of group decision-making. *Journal of Political Economy* 56(1):23–34.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2011. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation* 209(2):89–107.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2010. Using complexity to protect elections. *Communications of the ACM* 53(11):74–82.
- Fitzsimmons, Z.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2013. Control in the presence of manipulators: Cooperative and competitive cases. In *Proc. of IJCAI-13*, 113–119.
- Fitzsimmons, Z.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2014. X THEN X: Manipulation of same-system runoff elections. Technical Report arXiv:1301.6118 [cs.GT], arXiv.org.
- Fitzsimmons, Z. 2014. Single-peaked consistency for weak orders is easy. Technical Report arXiv:1406.4829 [cs.GT], arXiv.org.
- Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41(4):587–601.
- Lackner, M. 2014. Incomplete preferences in single-peaked electorates. In *Proc. of AAAI-14*, 742–748.
- Mattei, N., and Walsh, T. 2013. PREFLIB: A library for preferences. In *Proc. of ADT-13*, 259–270.
- Satterthwaite, M. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10(2):187–217.
- Walsh, T. 2011. Where are the hard manipulation problems? *JAIR* 42(1):1–29.