

Computational Social Choice: Strategic and Combinatorial Aspects

Lirong Xia

Department of Computer Science
Duke University
Durham, NC 27708, USA
lxia@cs.duke.edu

When agents have conflicting preferences over a set of alternatives and they want to make a joint decision, a natural way to do so is by *voting*. How to design and analyze desirable voting rules has been studied by economists for centuries. In recent decades, technological advances, especially those in internet economy, have introduced many new applications for voting theory. For example, we can rate movies based on people's preferences, as done on many movie recommendation sites. However, in such new applications, we always encounter a large number of alternatives or an overwhelming amount of information, which makes computation in voting process a big challenge. Such challenges have led to a burgeoning area—*computational social choice*, aiming to address problems in computational aspects of preference representation and aggregation in a multi-agent scenario.

The high-level goal of my research is to better understand and prevent the agents' (strategic) behavior in voting systems, as well as to design computationally efficient ways for agents to present their preferences and make a joint decision.

Strategic Behavior Analysis

Manipulation is the phenomenon that an agent reports her preferences falsely to make herself better off. It is one of the central problems in mechanism design, e.g., how can we design an auction protocol in which no agent has incentive to lie? In social choice theory, we have the same problem of how to design a voting rule to discourage manipulation. A voting rule is *strategy-proof* if under it, manipulation is never beneficial for an agent. A very negative result states that there is no strategy-proof voting rule that satisfies several very natural properties (Gibbard 1973; Satterthwaite 1975). However, even if possibilities for manipulation exist, it might be computationally hard for the manipulator to find the manipulation. This means that computational complexity might be used as a barrier against manipulation. Recently, there has been a lot of work on the complexity of finding a manipulation in various settings for common voting rules (Bartholdi, Tovey, and Trick 1989; Conitzer, Sandholm, and Lang 2007; Hemaspaandra and Hemaspaandra 2007; Faliszewski, Hemaspaandra, and Schnoor 2008).

However, for several common voting rules the complexity of manipulation remained open for many years, in the

most basic setting where the voters are equally weighted. We proved that for the *maximin* rule and *ranked pairs* rule, it is NP-complete to find a manipulation; for the *Bucklin* rule, we designed a polynomial-time algorithm to find manipulations (Xia et al. 2009); in another paper (Xia, Conitzer, and Procaccia 2010), we proved that it is strongly NP-complete for a specific positional scoring rule.

Another problem in voting system that is closely related to manipulation and preference elicitation is the following: when the voters' preferences have not yet been completely revealed, can the winner already be determined? We conducted an exhaustive study on the complexity of determining the winner under common voting rules, when the voters' revealed preferences are modeled by *partial orders* (Xia and Conitzer 2008a).

However, NP-hardness is a *worst-case* concept. Recent work suggests that a manipulation is easy to find for *most cases* (Conitzer and Sandholm 2006; Procaccia and Rosenschein 2007; Friedgut, Kalai, and Nisan 2008; Zuckerman, Procaccia, and Rosenschein 2009). Within this line of research, we derived a set of conditions on voting rules that admits a very simple randomized algorithm to find an instance of manipulation (Xia and Conitzer 2008c). To study the frequency with which a randomly drawn profile is manipulable, we introduced the concept of *generalized scoring rules*, and studied such frequency for any generalized scoring rules (Xia and Conitzer 2008b), which include all the common voting rules. In a recent paper, we completely characterized the class of generalized scoring rules as the set of voting rules that satisfy two very natural axioms (Xia and Conitzer 2009).

Combinatorial Voting

In real life, often the set of alternatives has a multi-issue structure. That is, there is a set of issues (or attributes), and an alternative is uniquely identified by the values that these issues take. Such domains are called *multi-issue domains*. For example, imagine a website that would like its users to vote on the default settings of the website. There are multiple issues that need a decision (background theme, news stories displayed, etc.).

Some aspects of voting in multi-issue domains have been extensively studied by economists; recently, voting in multi-issue domains has attracted the interest of researchers in computer science. The motivation for the study is that the

number of alternatives is exponentially large, hence introduces many difficulties for preference representation and aggregation. Therefore, new languages for voters to express their preferences, as well as novel voting rules to aggregate these preferences, are desired.

One feasible solution is to use acyclic *CP-nets* (Boutilier et al. 2004)—a compact language for representing preferences in multi-issue domains—to model the voters’ preferences, and sequentially apply a voting rule to each issue sequentially (Lang and Xia 2009). We have extensively studied the properties of such sequential voting processes, including how properties of issue-wise rules transfer to the sequential voting process (and vice versa) (Lang and Xia 2009), and whether or not such sequential voting processes satisfy some very natural properties, such as *neutrality* and *efficiency* (Xia and Lang 2009).

To relax the restrictive constraint on the language used in sequential voting processes, we proposed two extensions of the sequential voting process. The first extension adopts a language that allows voters to use acyclic CP-nets that are consistent with a non-fixed order of issues (Xia, Lang, and Ying 2007). The second provides a general framework for the voters to use (possibly cyclic) CP-nets to represent their preferences (Xia, Conitzer, and Lang 2008). For the second extension, we designed an algorithm to compute the winner, which is very efficient when there are strong conditional independencies among issues. We also extend the *maximum likelihood approach* to voting rules (Conitzer and Sandholm 2005; Conitzer, Rognlie, and Xia 2009) to multi-issue domains (Xia, Conitzer, and Lang 2010).

Another downside of sequential elections is that the chair can control the outcome of the election by changing the order in which issues are voted on. In a recent paper, we showed several hardness and easiness results for the chair’s control problem (Conitzer, Lang, and Xia 2009).

Future work

Many real-life preference aggregation instances are multi-issue settings, yet almost all approaches developed previously in the social choice community are encumbered by an extremely high computational cost in this context. Therefore, preference representation and aggregation in multi-issue domain is a promising new line of research, which is of great academic interest and wide potential application.

Besides voting in multi-issue domains, I plan to develop techniques/tools in the generalized voting rules framework, for example, designing algorithms for finding a manipulation, and studying the complexity of storing and communicating part of the electorate’s results (Chevalleyre et al. 2009; Xia and Conitzer 2010a). Another interesting direction is to analyze voters’ strategic behavior in a game-theoretic perspective (e.g., in (Xia and Conitzer 2010b)).

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.
- Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. CP-nets: a tool for representing and reasoning with conditional ceteris paribus statements. *JAIR* 21:135–191.
- Chevalleyre, Y.; Lang, J.; Maudet, N.; and Ravilly-Abadie, G. 2009. Compiling the votes of a subelectorate. In *IJCAI*, 97–102.
- Conitzer, V., and Sandholm, T. 2005. Common voting rules as maximum likelihood estimators. In *UAI*, 145–152.
- Conitzer, V., and Sandholm, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *AAAI*, 627–634.
- Conitzer, V.; Lang, J.; and Xia, L. 2009. How hard is it to control sequential elections via the agenda? In *IJCAI*, 103–108.
- Conitzer, V.; Rognlie, M.; and Xia, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *IJCAI*, 109–115.
- Conitzer, V.; Sandholm, T.; and Lang, J. 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM* 54(3):Article 14, 1–33.
- Faliszewski, P.; Hemaspaandra, E.; and Schnoor, H. 2008. Copeland voting: ties matter. In *AAMAS*, 983–990.
- Friedgut, E.; Kalai, G.; and Nisan, N. 2008. Elections can be manipulated often. In *FOCS*, 243–249.
- Gibbard, A. 1973. Manipulation of voting schemes: a general result. *Econometrica* 41:587–602.
- Hemaspaandra, E., and Hemaspaandra, L. A. 2007. Dichotomy for voting systems. *Journal of Computer and System Sciences* 73(1):73–83.
- Lang, J., and Xia, L. 2009. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences* 57(3):304–324.
- Procaccia, A. D., and Rosenschein, J. S. 2007. Junta distributions and the average-case complexity of manipulating elections. *JAIR* 28:157–181.
- Satterthwaite, M. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187–217.
- Xia, L., and Conitzer, V. 2008a. Determining possible and necessary winners under common voting rules given partial orders. In *AAAI*, 196–201.
- Xia, L., and Conitzer, V. 2008b. Generalized scoring rules and the frequency of coalitional manipulability. In *EC*, 109–118.
- Xia, L., and Conitzer, V. 2008c. A sufficient condition for voting rules to be frequently manipulable. In *EC*, 99–108.
- Xia, L., and Conitzer, V. 2009. Finite local consistency characterizes generalized scoring rules. In *IJCAI*, 336–341.
- Xia, L., and Conitzer, V. 2010a. Compilation complexity of common voting rules. In *AAAI*.
- Xia, L., and Conitzer, V. 2010b. Stackelberg voting games: Computational aspects and paradoxes. In *AAAI*.
- Xia, L., and Lang, J. 2009. A dichotomy theorem on the existence of efficient or neutral sequential voting correspondences. In *IJCAI*, 342–347.
- Xia, L.; Zuckerman, M.; Procaccia, A. D.; Conitzer, V.; and Rosenschein, J. 2009. Complexity of unweighted coalitional manipulation under some common voting rules. In *IJCAI*, 348–353.
- Xia, L.; Conitzer, V.; and Lang, J. 2008. Voting on multiattribute domains with cyclic preferential dependencies. In *AAAI*, 202–207.
- Xia, L.; Conitzer, V.; and Lang, J. 2010. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *AAMAS*.
- Xia, L.; Conitzer, V.; and Procaccia, A. 2010. A scheduling approach to coalitional manipulation. In *EC*.
- Xia, L.; Lang, J.; and Ying, M. 2007. Strongly decomposable voting rules on multiattribute domains. In *AAAI*, 776–781.
- Zuckerman, M.; Procaccia, A. D.; and Rosenschein, J. S. 2009. Algorithms for the coalitional manipulation problem. *Artif. Intell.* 173(2):392–412.