Research Statement

Lirong Xia
lxia@cs.duke.edu

February 19, 2009

1 Overview of research

My research mainly focuses on the intersection of computer science and (micro)economics. 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. Technological advances in recent decades 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. As another example, we may want to find the "most important" webpage among a set of webpages, based on the fact that each webpage contains links to other webpages; this can be interpreted as that the webpage "votes" for those webpages.

However, these new applications also bring new challenges to voting theory, in the sense of computational complexity: there are often a large number of alternatives, or an overwhelming amount of available information. These challenges have led to a new and fast-growing area—*computational social choice*, which studies computational aspects of preference representation and aggregation in a multi-agent scenario [3]. Computational social choice is closely related to *mechanism design*, where the goal is to reach some decision (for example, in an auction, the allocation of resources), and agents will strategically report their preferences to obtain the best result for themselves. Mechanism design often uses payments to influence agents' incentives, but in social choice theory, usually there are no payments.

So far, most of my research has focused on computational social choice. My main contributions have been to the following topics: complexity of manipulation, and voting in multi-issue domains. The high-level goal of my research is to design computationally efficient ways for agents to present their preferences and make a joint decision, as well as to better understand the agents' (strategic) behavior in reporting preferences.

Complexity of manipulation

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 [11, 16]. 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 [1, 7, 12, 9].

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. I proved that for the *maximin* rule and *ranked pairs* rule, it is NP-complete to find a manipulation; for the *Bucklin* rule, I designed a polynomial-time algorithm to find manipulations [23]. This work leaves only the complexity of manipulation under the *positional scoring* rule unsolved.

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? I 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* [17]. I showed that for some voting rules this problem is NP-hard, while for the others, I designed efficient algorithms to compute the winner.

However, NP-hardness is a *worst-case* concept. Recent work suggests that a manipulation is easy to find for *most cases* [15, 6, 10, 26, 14]. Within this line of research, I derived a set of conditions on voting rules that admits a very simple randomized algorithm to find an instance of manipulation [19]. To study the *frequency of manipulability*, that is, the frequency with which a randomly drawn profile is manipulable, I introduced the concept of *generalized scoring rules*, and studied the frequency of manipulability for any such rule [18]. Generalized scoring rules include all the common voting rules. In a recent paper, I completely characterized the class of generalized scoring rules as the set of voting rules that satisfy three very natural properties [20].

Voting in multi-issue domains

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, whether temperature is displayed in Celsius or Fahrenheit, *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 sciences. The motivation is that the number of alternatives is exponentially large, which introduces many difficulties for preference representation and aggregation. On the one hand, it is very hard or even impossible for a voter to represent her preferences as a *linear order* on the alternatives, which is required for applying common voting rules. On the other hand, even if the voters are able to represent their preferences as linear orders, it is computationally extremely costly to apply common voting rules. 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* [2]—a compact language for representing preferences in multi-issue domains—to model the voters' preferences, and sequentially apply a voting rule to each issue separately [13]. I 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) [13], and whether or not such sequential voting processes satisfy some very natural properties, such as *neutrality* and *efficiency* [24].

One drawback of the sequential voting process is that the language used for the voters to represent their preferences is too restrictive [22]. To relax this constraint, I proposed two extensions of the sequential voting process, for two different representation languages. The first extension adopts a language that allows voters to use acyclic CP-nets that are consistent with a non-fixed order of issues (in contrast to a fixed order of issues in the sequential voting process) [25]. The second provides a general framework for the voters to use (possibly cyclic) CP-nets to represent their preferences [22]. For the second extension, I also designed an algorithm to compute the winner, which is very efficient when there are strong conditional independencies among issues.

Another downside of sequential elections is that the chair(person) can control the outcome of the election by changing the order in which issues are voted on. In a recent joint paper, we showed several hardness and easiness results for the chair's control problem [4].

As an alternative to using computational complexity as a barrier against manipulation, an approach that has been studied by economists to eliminate manipulation is to restrict the voters' preference domain. That is, the voters are only allowed to represent their preferences by choosing a linear order from a restricted set—for example, the set of all linear orders that can be represented by a CP-net with no edges. Although much work has been done on this problem in multi-issue domains, none of this work uses a language that captures conditional independencies among issues. In the context of sequential voting processes and their extensions, I characterized strategy-proof voting rules in three types of restricted preference domains, when the voters' preferences are represented by CP-nets consistent with a common order [21].

2 Future work

I am excited to work more on preference representation and aggregation in multi-issue domain. This is a promising new line of research, because on the one hand, many real-life instances are multi-issue settings, and on the other hand, almost all approaches developed previously in the social choice community are encumbered by an extremely high computational cost in this context. There are many important problems to work on. Recently, we wrote a joint paper on how voting rules correspond to *maximum likelihood estimators* [5], and we plan to extend this to multi-issue settings. Another important topic of study is the *communication complexity* of winner determination in multi-issue domains. Recently, my supervisor received a $429,212 NSF grant for research on voting in multi-issue domains.

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. Beyond computational social choice, I am also interested in game theory and mechanism design. Recently, I started some joint work with other researchers at Duke on auctions with budget constraints (based on earlier work [8]), and we are making good progress.

References

[1] John Bartholdi, III, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[2] Craig Boutilier, Ronen Brafman, Carmel Domshlak, Holger Hoos, and David Poole. CP-nets: a tool for representing and reasoning with conditional ceteris paribus statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[3] Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. A short introduction to computational social choice. In *SOFSEM'07*, 2007.

[4] Vincent Conitzer, Jérôme Lang, and Lirong Xia. How hard is it to control sequential elections via the agenda? In submission, 2009.

[5] Vincent Conitzer, Matthew Rognlie, and Lirong Xia. Social welfare functions that score rankings and maximum likelihood estimation. In *Proceedings of the Second International Workshop on Computational Social Choice (COMSOC)*, Liverpool, England, 2008.

[6] Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 627–634, Boston, MA, USA, 2006.

[7] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 1–33, 2007. Early versions in AAAI-02 and TARK-03.

[8] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *FOCS*, pages 260–269, 2008.

[9] Piotr Faliszewski, Edith Hemaspaandra, and Henning Schnoor. Copeland voting: ties matter. In *Proceedings of AAMAS '08*, pages 983–990, 2008.

[10] Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 243–249, 2008.

[11] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–602, 1973.

[12] Edith Hemaspaandra and Lane A. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.

[13] Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences. In Press*, 2009.

[14] Ariel D. Procaccia and Jeffrey S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Honolulu, HI, USA, 2007.

[15] Ariel D. Procaccia and Jeffrey S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, February 2007.

[16] Mark Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[17] Lirong Xia and Vincent Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 196–201, Chicago, IL, USA, 2008.

[18] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 109–118, Chicago, IL, USA, 2008.

[19] Lirong Xia and Vincent Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 99–108, Chicago, IL, USA, 2008.

[20] Lirong Xia and Vincent Conitzer. Finite local consistency characterizes generalized scoring rules. In submission, 2009.

[21] Lirong Xia and Vincent Conitzer. Strategy-proof voting rules over multi-issue domains with restricted preferences. In submission, 2009.

[22] Lirong Xia, Vincent Conitzer, and Jérôme Lang. Voting on multiattribute domains with cyclic preferential dependencies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 202–207, Chicago, IL, USA, 2008.

[23] Lirong Xia, Vincent Conitzer, Ariel Procaccia, and Jeffrey Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the Second International Workshop on Computational Social Choice (COMSOC)*, Liverpool, England, 2008.

[24] Lirong Xia and Jérôme Lang. On the neutrality and efficiency of seat-by-seat voting rules and correspondences. In *The 9th International Meeting of the Society of Social Choice and Welfare*, 2008.

[25] Lirong Xia, Jérôme Lang, and Mingsheng Ying. Strongly decomposable voting rules on multiattribute domains. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, BC, Canada, 2007.

[26] Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artif. Intell.*, 173(2):392–412, 2009.