

Improving Group Decision-Making by Artificial Intelligence

Lirong Xia

Rensselaer Polytechnic Institute
xial@cs.rpi.edu

Abstract

We summarize some of our recent work on using AI to improve group decision-making by taking a unified approach from statistics, economics, and computation. We then discuss a few ongoing and future directions.

1 Introduction

The One Hundred Year Study on Artificial Intelligence envisioned that “the field of AI is shifting toward building intelligent systems that can collaborate effectively with people, and that are more generally human-aware” [Stone *et al.*, 2016]. My research has been well-positioned within this trend: the goal is *developing and leveraging AI techniques to help human beings and software agents make better decisions, by bridging theory, practice, and education.*

For example, suppose a university is hiring a new faculty member. After the interviews, the committee members rank the candidates and vote to decide the top choice. How can the committee members reduce the uncertainty in the quality of the candidates in regard to their potential to do influential research, ability to teach, and their fit with the department? How can they make a fair decision when people have conflicting preferences? Might the committee worry about its members strategically misreporting their preferences and information?

This example illustrates the problem of group decision-making. Similar challenges appear in many applications such as political elections, meta search engines, recommender systems, crowdsourcing, etc. Evidently, addressing this fundamentally multi-disciplinary challenge requires considering three types of criteria:

- *Statistical criteria* evaluate the quality of decisions in the statistical sense. A typical research topic is multinomial logistic regression in statistics.
- *Socio-economic criteria* include various desirable normative properties in social choice theory and sociology, such as fairness, strategy-proofness, and ethics.
- *Computational criteria* are critical for big data. A typical research topic is rank aggregation.

Although there has been interdisciplinary work, much of it has overlooked at least one aspect. For example, research in

machine learning (Statistics+Computation) often overlooks fairness. Research in computational social choice (Economics+Computation) often lacks considerations from statistics.

2 Overview of My Research

My research tackles the multi-disciplinary challenge of group decision-making by taking a unified approach from statistics, economics, and computation. My work in the first direction (Statistics+Computation) improved the state of the art in learning from rank data; my work in the second direction (Economics+Computation) belong to the burgeoning field of computational social choice; and my work in the third direction (Statistics+Economics+Computation) is conceptually new.

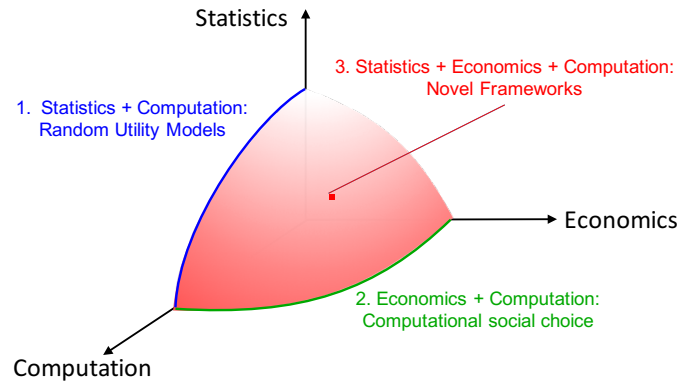


Figure 1: My research on group decision-making.

3 AI’s Role

Recently there have been many discussions and concerns on moral aspects of AI [Rossi, 2016]. In the context of group decision-making, should we trust AI to make right decisions for us, such as the next president, national defense strategies, and economic policies? How can we quantitatively measure the moral aspects of AI? How can we design moral AI algorithms?

The answer certain depends on the application. For low-stakes applications, such as a group of friends deciding where

to go for dinner, giving AI more control may improve the efficiency of decision-making without introducing too much risk. For high-stakes applications such as presidential elections, it seems a good idea to be more careful and use AI only as a supporter rather than the decision-maker.

4 Research Contributions

In light of ethical considerations of AI, the general theme of my research is *using AI to support group decision-making*, rather than letting AI make group decisions. Below I will briefly discuss three directions I have been pursuing, illustrated in Figure 1.

4.1 Direction 1 (Statistics+Computation): Learning Random Utility Models

Random utility models (RUMs) [Thurstone, 1927] are one of the most widely-justified and widely-applied models for decision-making from rank data. For example, McFadden was awarded the 2000 Nobel Prize in Economics for his contributions in the theory and practice of *discrete choice models*, which are special cases of RUMs. Other notable applications of RUMs include to elections, crowdsourcing, recommender systems, marketing, health care, transportation, and security.

In an RUM, each alternative a_i is parameterized by a utility distribution μ_i , which often belongs to a parameterized family of distributions, e.g. Gaussian distributions parameterized by means and variances. Agent’s rankings are generated in two steps. In the first step, a *latent utility* U_i for each alternative a_i is generated from μ_i . In the second step, the alternatives are ranked w.r.t. their utilities U_i in the decreasing order. An RUM for three alternatives and the process of generating $a_2 \succ a_1 \succ a_3$ is illustrated in Figure 2.




Figure 2: Generating $a_2 \succ a_1 \succ a_3$ in a RUM.

However, designing efficient algorithms for learning general RUMs is a well-known open question due to the lack of closed-form formulas for the likelihood function. Most previous research and applications were limited to the *Plackett-Luce model*, which is a computationally tractable subcase.

Algorithms for learning general RUMs.

With Hossein Azari Soufiani and David Parkes, I proposed an MC-EM algorithm [Azari Soufiani *et al.*, 2012], which is the first algorithm for computing the MLE of general RUMs. This helps to improve the quality of decisions in many applications due to better fitness of general RUMs. We further proposed a flexible *rank-breaking* framework to explore tradeoffs between computational efficiency and statistical efficiency [Azari Soufiani *et al.*, 2013a].

Using this framework, we obtained a much faster algorithm that achieves competitive statistical efficiency compared to the state-of-the-art algorithm [Hunter, 2004]. We characterized all *consistent* rank-breaking algorithms for the Plackett-Luce model, and extended them to general RUMs [Azari Soufiani *et al.*, 2014a] and mixtures of Plackett-Luce models [Zhao *et al.*, 2016]. We also designed preference elicitation algorithms under general RUMs according to various information-maximization principles [Azari Soufiani *et al.*, 2013b].

Identifiability of mixtures of Plackett-Luce models

The *identifiability* of mixtures of Plackett-Luce models has been a long-standing open question. Identifiability requires that different parameters of the model correspond to different distributions over data. Therefore, if a model is not identifiable, one must be very careful when interpreting the learned parameter, e.g. in clustering, because the ground truth can be completely different and the difference cannot be detected by any statistical method. However, identifiability was often overlooked by previous work, e.g. by Gormley and Murphy [2008]. For example, Gormley and Murphy [2008] used the mixture of four Plackett-Luce models to fit an Irish election dataset with five alternatives, and interpreted the learned parameters as *voting blocs*.

We proved the first theorems on the identifiability of finite mixtures of Plackett-Luce models [Zhao *et al.*, 2016] using algebraic geometry techniques especially tensor decomposition and analysis of Kruskal’s rank. Our theorems state that when $m \leq 2k - 1$, where m is the number of alternatives, the mixture of k Plackett-Luce models is not identifiable, which means that the results by Gormley and Murphy [2008] are potentially flawed (where $k = 4$ and $m = 5$). Our positive results are that the mixture model is identifiable for $k = 2$ and $m \geq 4$, and is *generically identifiable* under a much milder condition, i.e. $k \leq \lfloor \frac{m-2}{2} \rfloor!$.

4.2 Direction 2 (Economics+Computation): Computational Social Choice

Algorithmic game theory [Nisan *et al.*, 2007] and computational social choice [Brandt *et al.*, 2016] are recognized as one of the eleven “*fundamental methods and application areas*” of AI, according to The One Hundred Year Study on Artificial Intelligence [Stone *et al.*, 2016]. Some of my work focused on the following two key research topics.

Topic 1. Using high complexity to prevent agents’ strategic behavior

A recurring concern in the design of systems for group decision-making is that of manipulation, where a participant might be able to benefit by misreporting her true preferences, beliefs, or information. A well-known impossibility result from economic theory establishes that this is inevitable for reasonable voting rules. Moreover, manipulation is just one example of strategic behavior—other concerns, especially with the advent of Internet systems, include bribery and false-name manipulation (sometimes called a sybil attack).

For more than two decades, researchers have been interested in understanding whether computational intractability

can provide a barrier against manipulative behavior, and understanding *for which voting system computing a manipulation is hard*. This problem is particularly critical given that nowadays the participants can easily access powerful computational tools.

My research has established the NP-hardness of manipulation in many common voting schemes, which suggests that computational complexity does provide some protection against manipulation. Still, NP-hardness is a worst-case concept and does not imply that manipulation is hard to compute in practice. A series of work of mine suggests that despite worst-case results, computational complexity by itself may not provide an effective barrier against strategic behavior. Most of my results are summarized in a recent book chapter [Conitzer and Walsh, 2016] and a text book in preparation by Parkes and Seuken [2016].

Topic 2. Combinatorial voting

In many applications the number of alternatives is exponentially large in a natural description of the problem. A prominent example is *combinatorial voting*, where there are multiple *issues* and each alternative can be uniquely characterized by assigning a value to each issue. This is an important voting model in public choice and also for Internet applications. For instance, residents in Florida voted in the 2012 US election to decide 11 issues, 5 out of which are interrelated tax policies. For each issue a resident can vote for “pass” or “deny”. When using voting for meeting scheduling, users may need to determine at least (1) the location, and (2) the time.

Most previous work focused on *simultaneous voting*, where participants vote over issues *separately* and *at the same time*. However, this approach is not suitable when participants’ preferences over one issue depend on what is decided in regard to other issues. I have used compact knowledge representation languages such as *CP-nets* [Boutilier *et al.*, 2004] to represent agents’ preferences, and designed *sequential* mechanisms with high computational and economic efficiency. Most of my results are summarized a recent book chapter written with Jerome Lang [Lang and Xia, 2016].

4.3 Direction 3 (Statistics + Economics + Computation): Novel Frameworks

There are many more group decision-making scenarios today than were envisioned in classical social choice theory. Because of this, I believe that designing *application-specific* mechanisms is a promising direction where AI will play an important role. Therefore, we proposed the following two frameworks under the unified consideration of statistics, economics, and computation.

Framework 1: Statistical decision-theoretic framework

We proposed a principled *statistical decision-theoretic framework for social choice* [Azari Soufiani *et al.*, 2014b], denoted by $\mathcal{F} = (\mathcal{M}, \mathcal{D}, L)$, which has three parts: (1) a statistical model $\mathcal{M} = (\Theta, \mathcal{S}, \text{Pr})$; (2) a decision space \mathcal{D} , and (3) a loss function $L(\theta, d)$ that evaluates the loss of decision $d \in \mathcal{D}$ against parameter $\theta \in \Theta$. This framework allows us to design new mechanisms following either the Bayesian or the frequentist approach.

Bayesian approach. Given a framework \mathcal{F} and a prior distribution π , the *Bayesian estimator* can be used for making group decisions. We designed multiple new mechanisms following the Bayesian approach, analyzed their social choice normative properties [Azari Soufiani *et al.*, 2014b; Xia, 2016], and designed MCMC sampling algorithms to compute them [Hughes *et al.*, 2015]. To the best of my knowledge, we were the first to study fairness for statistical estimators [Azari Soufiani *et al.*, 2014b]. I also proved the first impossibility theorem on social choice normative properties of Bayesian estimators, which states that no Bayesian estimator can satisfy the *strict Condorcet criterion* [Xia, 2016]. Our MCMC algorithms are the first sampling-based algorithms for ranking models with theoretical guarantees [Hughes *et al.*, 2015].

Frequentist approach. Frequentists often evaluate a decision-making mechanism $f_{\mathbb{F}}$ w.r.t. a fixed parameter θ by the *frequentist loss*, denoted by $\text{FL}(\theta, f_{\mathbb{F}})$, which is the expected loss of the decision made by $f_{\mathbb{F}}$ when the parameter is fixed to be θ and the data is generated given θ . That is,

$$\text{FL}(\theta, f_{\mathbb{F}}) = E_{P \sim \text{Pr}_{\theta}} L(\theta, f_{\mathbb{F}}(P)),$$

where P represents the randomly-generated data. *Minimaxity* is a commonly-used optimality criterion that measures the robustness of the mechanism, which is the worst-case frequentist’s loss for the mechanism, where the worst-case is taken over all parameters, namely $\min_{\theta \in \Theta} \text{FL}(\theta, f_{\mathbb{F}})$. A mechanism $f_{\mathbb{F}}$ is *minimax*, if it has the lowest worst-case frequentist loss among *all* mechanisms. Recently, I proved that for many statistical decision-theoretic frameworks, the MLE is *minimax* [Xia, 2016]. This implies that for many frameworks, MLEs have the minimum sample complexity, which was only previously known for Mallows’ model [Caragiannis *et al.*, 2013].

Framework 2: Automated mechanism design

In a position paper [Xia, 2013], I proposed to use machine learning to automatically learn a mechanism that satisfies a user-specified set of desirable social choice normative properties. The main idea is that many desirable normative properties can be viewed as logical rules for generating new data. For example, *monotonicity* states that for any agent, raising the position of an alternative in her ranking does not hurt the alternative. This can be viewed as a data generation rule such that for any positive example (P, c) , where P is a profile and c is the winner-to-be, and for any P' obtained from P by raising the position of c in one ranking, we have that (P', c) must also be a positive example. After new data are generated according to the normative properties, we learn a mechanism within structured frameworks [Xia, 2015].

5 Ongoing and Future Research

It is important that theoretical results and frameworks can be applied to improve group decision-making in practice. Below I will briefly discuss three ongoing and future directions for research and development.

The screenshot shows the OPRA interface for submitting a CP-net. At the top, there is a navigation bar with 'Welcome Lirong' and icons for 'Regular Polls', 'Multi-Polls', and 'Groups'. Below this is a header for the current poll: '-< Date (Created by lirong)'. The main content area is divided into three sections:

- Area 1:** 'Conditional Preferences' section. It shows a dependency graph where 'Date' depends on 'Topic'. The graph is represented as:


```

            graph TD
            Topic --> Date
            
```
- Area 2:** 'Topic' selection section. It shows three topics: T1 (green), T2 (blue, selected), and T3 (red).
- Area 3:** 'Your preferences' and 'Remaining Options' sections. The 'Your preferences' section shows a list of items: 1 D1, 2 D2. The 'Remaining Options' section shows 3 D3. There are 'Clear' and 'Submit' buttons for preferences, and a 'Move All' button for remaining options.

Figure 3: Submitting CP-nets at OPRA.

Direction 1. Group decision support systems

Previous research in *group decision support systems* focused on removing communication barriers and building hardware systems [Desanctis and Gallupe, 1987], but often overlooked the role of decision-making mechanisms. My group has been working on establishing principled theoretical and algorithmic frameworks that unify group decision support systems and AI, especially computational social choice and machine learning. There are many places AI and economics can significantly improve the state of the art, such as in learning users' behavior for better UI design, measuring consensus in the group, and using machine learning to guide heuristic search for computing hard voting mechanisms [Jiang *et al.*, 2017].

Direction 2. Multi-type resource allocation

In *multi-type resource allocation* problems, items are categorized into multiple types and each agent must get at least one item per type. For example, the problem arises in allocating courses to students, computational resources to users in cloud computing, medical resources to patients, etc. Previously only negative results are known.

In two recent papers [Mackin and Xia, 2016; Sikdar *et al.*, 2017], we were able to obtain surprisingly positive results by using AI techniques. For example, when agents' preferences are *lexicographic* and are represented by CP-nets [Boutilier *et al.*, 2004], we designed an extension of the classical top-trading-cycles mechanism and proved that it satisfies many desirable properties. I believe that establishing theoretical and algorithmic foundations of mechanism design for multi-type resource allocation with the help of AI is a promising direction for future research.

Direction 3. Online Preference Reporting and Aggregation (OPRA) system

My group has built an open-source system for online group decision-making. The service is open to public at <http://opra.io>, and all source code can be found at Github (<https://github.com/PrefPy/opra>). OPRA has been used in classes at RPI for students to make various decisions, for CS department to decide best poster awards, and for running polls for RPI's Grand Marshall Week.

In my view, OPRA serves as a framework for bridging social choice theory and group decision-making in practice by collecting data, testing new mechanisms, verifying theoretical models, and providing new insights to theoretical problems. In fact, many outcomes of my research have already been integrated to OPRA. For example, OPRA measures and visualizes consensus in agents' preferences by computing the *margin of victory* [Xia, 2012] and mixtures of Plackett-Luce models [Zhao *et al.*, 2016].

As another example, Figure 3 shows OPRA's current support for combinatorial voting and multi-type resource allocation. A user is submitting a CP-net to represent her preferences over two types: Topic and Date, each of which has three items. Area 1 shows the configurable dependency graph of the CP-net. Area 2 provides indices to all topics, and currently T_2 is chosen. Area 3 offers an optimized UI for the user to submit a ranking over the three dates conditioned on the chosen topic (T_2).

Acknowledgements

This work is supported by National Science Foundation under grant IIS-1453542. We thank Vincent Conitzer, David Parkes, and Ariel Procaccia for helpful comments.

References

- [Azari Soufiani *et al.*, 2012] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Random utility theory for social choice. In *Proceedings of NIPS*, pages 126–134, Lake Tahoe, NV, USA, 2012.
- [Azari Soufiani *et al.*, 2013a] Hossein Azari Soufiani, William Chen, David C. Parkes, and Lirong Xia. Generalized method-of-moments for rank aggregation. In *Proceedings of NIPS*, Lake Tahoe, NV, USA, 2013.
- [Azari Soufiani *et al.*, 2013b] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Preference Elicitation For General Random Utility Models. In *Proceedings of UAI*, Bellevue, Washington, USA, 2013.
- [Azari Soufiani *et al.*, 2014a] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Computing Parametric Ranking Models via Rank-Breaking. In *Proceedings of ICML*, Beijing, China, 2014.
- [Azari Soufiani *et al.*, 2014b] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Statistical decision theory approaches to social choice. In *Proceedings of NIPS*, Montreal, Quebec, Canada, 2014.
- [Boutilier *et al.*, 2004] 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.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jerome Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Caragiannis *et al.*, 2013] Ioannis Caragiannis, Ariel Procaccia, and Nisarg Shah. When do noisy votes reveal the truth? In *Proceedings of ACM EC*, Philadelphia, PA, 2013.
- [Conitzer and Walsh, 2016] Vincent Conitzer and Toby Walsh. Barriers to Manipulation in Voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia, editors, *Handbook of Computational Social Choice*, chapter 6. Cambridge University Press, 2016.
- [Desanctis and Gallupe, 1987] Gerardine Desanctis and R. Brent Gallupe. A foundation for the study of group decision support systems. *Management Science*, 33(5):589–609, 1987.
- [Gormley and Murphy, 2008] Isobel Claire Gormley and Thomas Brendan Murphy. Exploring voting blocs within the irish exploring voting blocs within the irish electorate: A mixture modeling approach. *Journal of the American Statistical Association*, 103(483):1014–1027, 2008.
- [Hughes *et al.*, 2015] David Hughes, Kevin Hwang, and Lirong Xia. Computing Optimal Bayesian Decisions for Rank Aggregation via MCMC Sampling. In *Proceedings of UAI*, pages 385–394, 2015.
- [Hunter, 2004] David R. Hunter. MM algorithms for generalized Bradley-Terry models. In *The Annals of Statistics*, volume 32, pages 384–406, 2004.
- [Jiang *et al.*, 2017] Chunheng Jiang, Sujoy Sikdar, Jun Wang, Lirong Xia, and Zhibing Zhao. Practical Algorithms for Computing STV and Other Multi-Round Voting Rules. In *EXPLORE-2017: The 4th Workshop on Exploring Beyond the Worst Case in Computational Social Choice*, 2017.
- [Lang and Xia, 2016] Jérôme Lang and Lirong Xia. Voting in Combinatorial Domains. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia, editors, *Handbook of Computational Social Choice*, chapter 9. Cambridge University Press, 2016.
- [Mackin and Xia, 2016] Erika Mackin and Lirong Xia. Allocating Indivisible Items in Categorized Domains. In *Proceedings of IJCAI*, pages 359–365, 2016.
- [Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [Parkes and Seuken, 2016] David C. Parkes and Sven Seuken. *Economics and Computation*. Book in preparation, 2016.
- [Rossi, 2016] Francesca Rossi. Moral preferences. In *The 10th Workshop on Advances in Preference Handling (MPREF)*, 2016.
- [Sikdar *et al.*, 2017] Sujoy Sikdar, Sibel Adali, and Lirong Xia. Mechanism Design for Multi-Type Housing Markets. In *Proceedings of AAI*, 2017.
- [Stone *et al.*, 2016] Peter Stone, Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyan Krishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller. Artificial intelligence and life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA, September 2016.
- [Thurstone, 1927] Louis Leon Thurstone. A law of comparative judgement. *Psychological Review*, 34(4):273–286, 1927.
- [Xia, 2012] Lirong Xia. Computing the margin of victory for various voting rules. In *Proceedings of ACM EC*, pages 982–999, Valencia, Spain, 2012.
- [Xia, 2013] Lirong Xia. Designing social choice mechanisms using machine learning. In *Proceedings of AAMAS*, pages 471–474, St. Paul, MN, USA, 2013.
- [Xia, 2015] Lirong Xia. Generalized Decision Scoring Rules: Statistical, Computational, and Axiomatic Properties. In *Proceedings of ACM EC*, pages 661–678, Portland, Oregon, USA, 2015.
- [Xia, 2016] Lirong Xia. Bayesian estimators as voting rules. In *Proceedings of UAI*, pages 785–794, 2016.
- [Zhao *et al.*, 2016] Zhibing Zhao, Peter Piech, and Lirong Xia. Learning Mixtures of Plackett-Luce Models. In *Proceedings of ICML*, 2016.