

Imputation, Social Choice, and Partial Preferences

John A. Doucette

David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada

Within the field of Artificial Intelligence (AI) research, the subfield of Computational Social Choice considers the application of AI techniques to problems in Social Choice (a.k.a. Voting Theory) (Chevaletyre et al. 2007). Starting in the early 1990's, computer scientists began to take an interest in social choice. Initial work was concerned with circumventing the impossibility results implied by earlier work, using computational hardness. This work showed that, in many voting systems, a strategic voter would need to solve an NP-Hard problem to determine how best to cast its ballot (Bartholdi III and Orlin 1991). More recently, interest in this field has broadened with the growth of multi-agent systems in AI. Topics like multi-agent resource allocation, public goods games, and recommender systems can all be understood as forms of social choice problems. My research addresses an unresolved problem in social choice and proposes the use of machine learning techniques, to arrive at a principled solution. The problem I address is deciding the outcome of an election with partial preference ballots, and my proposed solution could lead to an important linkage between machine learning and social choice theory.

As is conventional in social choice, I model an election as a choice among a set of alternatives C . This model is broad enough to encapsulate a wealth of problems in artificial intelligence, where individuals' preferences must be aggregated to make a decision. A set of voters V constitutes the electorate. Each voter is assumed to have preferences over C which can be represented as a linear ordering. For example, a voter who prefers candidate X to candidate Y , and candidate Y to candidate Z would write their preferences $X \succ Y \succ Z$. Voters may each cast a ballot, which is also represented as a linear order over C , and a voting rule S is used to map sets of ballots to outcomes, deciding the election. For example, the ballots in Figure 1 show a situation where 8 voters have submitted the ballot $X \succ Y \succ Z$, 7 the ballot $Y \succ Z \succ X$, and 6 the ballot $Z \succ Y \succ X$. Different voting rules, like plurality or Single Transferable




Figure 1

Vote (STV) select winners from sets of ballot. For example, under plurality candidate X has won this election by virtue of having being at the top of the largest number of ballots. Under STV Y would win instead.

In this work, we model voters as having *partial* orderings over the candidates for their preferences, instead of linear orderings. For instance, a voter who is certain that they like X , but unsure which of Y or Z is better would have the preferences $X \succ Y$, $X \succ Z$, $Y \sim Z$. When voters cast their ballots, they may represent this ambiguity by including recording the partial ordering directly, taking the form of a directed acyclic graph where the vertices are candidates, and there is an edge from X to Y iff $X \succ Y$ in the voter's preferences. Some voting systems are unable to support this general format however. A simpler format assumes that voters have good knowledge about their top t most preferred candidates, and are unsure about the relative orderings of the remainder. Thus, a top- t ballot is a linear ordering over the voter's top- t most preferred alternatives, with the remainder of the ordering left blank, indicating the remaining $|C| - t$ candidates are below all those listed, in an uncertain order.

The problem with allowing for partial preferences lies in deciding the outcome of an election. Voting rules must address the missing information in ballots with a consistent policy, and must select the alternative that *would* have won if all the voters had specified their complete preferences.

Existing approaches to the problem of elections with partial preference ballots treat missing information in the ballots in a number of different ways. Many of the systems used in practice simply treat it as absent. For instance, in the 'single transferable vote' system ballots that run out of candidates are discarded, and the remainder of the election proceeds as though they had never been cast. Recent interest in avoiding this has led to more principled approaches.

Other approaches looked at different ways to interpret the missing information in the ballots, using optimization techniques or statistical inference. Approaches include assuming the missing pairwise preferences are uniformly distributed (Xia and Conitzer 2011), or selecting the winner based on a worst-case completion of the ballots (e.g. min-max regret) (Lu and Boutilier 2011). My approach is instead to assume that the missing preferences are distributed similarly to the preferences that *were* observed. Specifically, I assume that voters are similar to one another, given that

whatever information is contained on their ballots is similar. I then utilize machine learning techniques to impute the missing information, and decide the election on the basis of the completed ballots. In this way, my aim is to complete the missing information with the true preferences that users would have specified, had they been able.

More formally, my approach consists of re-phrasing the problem of deciding an election with partial preference ballots as a classification prediction task. Given the preferences that a particular voter has written on their ballots, we want to predict their true values for the missing information on the ballot. To do this, we break the problem into a number of sub-tasks. In the first task, we take all ballots with at least two preferences, and train a classifier to predict the second preference from the first. I then use this classifier to predict the missing second preferences of any ballot with just one preference recorded. Repeating this process, we can learn to predict third preferences from all ballots with at least three recorded values, and impute any missing third preferences on the remainder of the ballots. The process is iterated until all ballots are completed, and the complete ballots are then used to decide the final outcome of the election.

My initial work on this approach consisted of both a theoretical and experimental component. I first showed that under the top-t preference model, applying classifiers to predict the missing components of the ballots was equivalent to using the ballots themselves as votes under a specific election rule, to select from among the set of possible completions. That is, using machine learning to impute the missing preference is itself a form of social choice. This exciting result facilitates future theoretical work, making it possible to analyze my treatment of missing information using established tools from computational social choice, and hints at a possible bridging theory that connects the two fields tightly. It is this topic I hope to expand on in great detail. Beyond this exciting theoretical result, I was able to show that the proposed technique is useful in practice, with a detailed empirical evaluation which established the ability of machine learning models to impute missing information from real-world human preferences. In particular, I evaluated the performance of the new approach across ballots from ten real-world elections and found that it compares well with state-of-the-art techniques. Performance was also quite good in absolute terms. The correct winner was recovered in more than 90% of runs, and the overall ranking of the candidates matched that obtained by aggregating the ground-truth ballots to a very high degree. Error rates for the scores of individual candidates under the commonly used Borda Count voting rule were typically less than 1%. Overall, these results indicate very strong performance from my initial models, and suggest great promise for further refinement.

By July 2014, I anticipate the completion of a machine learning model tailored specifically to the problem of imputing missing preferences. I hope to benchmark that model, and to perform a more rigorous axiomatic analysis of it as a voting rule, including interpretations of Arrow's conditions in the context of learning missing preference information. I also hope to provide a performance guarantee under certain assumptions about the way in which missingness is gener-

ated in the data. These proofs would show that, when the assumptions underlying the model are satisfied, the preferences of specific users will be recovered with high probability. I also hope to characterize the assumptions underlying the model as analogous to specific, known, properties of social choice functions (e.g. Neutrality, anonymity, etc.).

Looking further ahead, I plan to consider the following possible research paths: a detailed axiomatic analysis of broader classes of machine learning methods when applied in this context; a characterization of the conditions under which specific machine learning techniques will converge to accurate models of voters' preferences; an examination the applicability of this technique to domains with very sparse preferences, like resource allocation (especially challenging for problems with combinatorial valuations); and a look at incorporating techniques like Bagging and Boosting that could improve the performance in real-world problems. The topic of manipulation (i.e. the strategic behaviors of both the voters and the electioneer) would also be valuable to explore. Ultimately I hope to make substantial contributions toward connecting machine learning models to social choice functions, providing benefits for both fields.

The potential value of this research is to offer a principled solution for selecting winners when voting with partial preference ballots, which should provide important progress in a number of central MAS research areas, including social choice, multiagent resource allocation, and areas of game theory like public goods games. Beyond the possible applications however, my topic is useful insofar as it draws an explicit link between machine learning and computational social choice in this domain. I am not merely applying machine learning methods to social choice problems, but showing that the application of machine learning models is really itself a form social choice, where participants are implicitly voting on how missing preference information should be treated. Fleshing out the details of this connection would be a major contribution to both fields, and offers tantalizing potential for deep insights. For example, showing that particular classification techniques are equivalent to particular classes of voting rules would tie the fields together strongly, and would allow many existing theoretical results in both to be merged.

References

- Bartholdi III, J. J., and Orlin, J. B. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8(4):341–354.
- Chevalleyre, Y.; Endriss, U.; Lang, J.; and Maudet, N. 2007. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science*, 51–69. Springer.
- Lu, T., and Boutilier, C. 2011. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, 287–293. AAAI Press.
- Xia, L., and Conitzer, V. 2011. A maximum likelihood approach towards aggregating partial orders. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence-Volume Volume One*, 446–451. AAAI Press.