Probabilistic and Utility-theoretic Models in Social Choice: Challenges for Learning, Elicitation, and Manipulation

Craig Boutilier
University of Toronto
Department of Computer Science
cebly@cs.toronto.edu

Tyler Lu
University of Toronto
Department of Computer Science
tl@cs.toronto.edu

1 Introduction

The abundance of inexpensive preference data facilitated by online commerce, search, recommender systems, and social networks has the potential to stretch the boundaries of social choice. Specifically, concepts and models usually applied to high stakes domains such as political elections, public or corporate policy decisions, and the like, will increasingly find themselves used in the lower stakes, high-frequency domains addressed by online systems.

Here are just two of many examples that are not typically interpreted as social choice problems, but which in fact, can profitably viewed as such. First, consider the (first page of) results returned by your favorite search engine to a specific query. While pure personalization, taking into account your specific preferences for results, would be ideal, this is generally not possible because of data scarcity. Hence the small amount of information known about your preferences is aggregated with the (equally scarce) data about users similar to you to determine the best results. This is a consensus decision making problem, since a single set of results is constructed for a collection of users, each of whom may have somewhat different preferences. Indeed, within the subfield of rank learning within machine learning, the *label ranking* paradigm [4] makes this assumption explicit. As a second example, consider the problem of an online retailer determining which *subset of size* k of potential products to offer to its target market. Ideally, the retailer would *segment* its audience into k groups such that a single product would be desirable to each member of the group [6]. Again, since a single choice is being proposed for all members of the group, this is a social choice problem [7].

Several factors make these and related problems both interesting and rather novel from a social choice perspective. First, the expression of complete preferences is wildly impractical: users will simply not tolerate much in the way of elicitation; and typically preferences will be *estimated* from choice behavior, partial ratings data, etc. Second, massive amounts of such data will in fact make it feasible to learn quite compelling probabilistic models of user preferences. Third, approximation will be an absolute necessity for several reasons: the need for "nearly instantaneous" recommendations will demand computational approximation; the incompleteness of preference data will demand informational approximation; and finally, very clear (usually economic) tradeoffs can be made that greatly facilitate the design of approximation methods (unlike, say, in political elections, where an "approximate winner" is unlikely to be viewed as satisfactory).

Issues of computational approximation have been studied extensively in social choice; informational approximation (dealing with incomplete preferences) has been too (though to a lesser extent); and probabilistic models have been used in analysis.[1] However, we feel the new demands of online systems call for a different style of analysis of social choice models and algorithms. Two key components lie at the heart of our proposal for such analyses: (a) utility-theoretic approximation, be it informational or computational; and (b) learning and exploiting probabilistic models of user preferences. We outline four broad categories of research challenges based on these components.

In what follows, we use A to denote a set of alternatives; U, a set of users or voters; v, a ranking, permutation, or *vote* over A; V, the set of permutations; $\mathbf{v}$, a *profile* with one (ranked) vote per voter; and r a voting rule, with $r(\mathbf{v})$ denoting the selected alternative given $\mathbf{v}$.

2 Learning Preferences

By a *probabilistic model*, we simply mean some distribution P over the set of rankings (or preferences) V. We'll discuss below various ways to exploit probabilistic models of user preferences when tackling various problems in social choice. However, one first needs *realistic* models of user preferences that support tractable inference and can be effectively learned from readily available data. Analysis of voting schemes in social choice tends to focus on models such as impartial culture which have little connection to reality in the settings mentioned above (or even in electoral data [11]).[2]

A number of models have been developed in econometrics, statistics and psychometrics that explicitly try to reflect the processes by which human comparison judgements are made, and are used to model population preferences. It is impossible to do justice to this literature here [10], but several of these models—especially the Mallows and Plackett-Luce models—have been appropriated by the machine learn-

[1] In this short position paper, we unfortunately must exclude references, even representative ones, on these topics.

[2] And even then, the questions addressed using such models tend to be very different than those we outline below.

ing community under the guise of "learning to rank" (LeToR). This has precipitated the development of many interesting methods for tractable learning and probabilistic inference with such models. This work is vitally important for the application of computational social choice, and we believe a rapprochement between the two disciplines is in order.

Of course, the flow works in both directions: the problems that arise in social choice must influence the development of new models and algorithms for learning and probabilistic inference. As one example, most work in LeToR assumes that observed rankings are noisy estimates of some underlying objective ranking (rather than representing genuinely distinct preferences). Because of the types of data sets considered, several important problems have gone unaddressed. For example, learning Mallows models is widely considered to be intractable with choice data consisting of pairwise comparisons of form $a_i \succ a_j$, obviously an important form of evidence in any social choice problem. We've developed a new model that allows Mallows models (and mixtures thereof) to be effectively learned from such data [8]. At its heart is the *generalized repeated insertion model (GRIM)*, that that allows approximate sampling of rankings conditioned on pairwise evidence.[3] With several real-world data sets, we've learned interesting population models with this technique.

Of course, this is just a start. More general models that support effective inference and tractable learning are needed, especially models that are tuned to the types of preference distributions we expect to find in consensus decision making domains. For example, realistic, tractable models for distributions over single-peaked preferences seem largely to have been unaddressed (and the "riffle independence" concept developed in ML may prove useful [3]).

3 Optimization

A second key issue is critical in the design of social choice methods for online settings, centered on the notion of utility-theoretic approximation of recommendations or "winners," especially when we have partial information about user preferences. While incomplete preferences are studied in a variety of guises, little attention is paid to the question of how to *select* a winner in such a situation.[4] In recent work, we've proposed using the notion of *minimax regret (MMR)* for just this purpose [9].

Most voting rules can be defined using a natural *scoring function* $s(a, \mathbf{v})$ that measures the quality or utility of alternative a given profile $\mathbf{v}$, i.e., $r(\mathbf{v}) \in \text{argmax}_{a \in A}\, s(a, \mathbf{v})$. Now suppose we have access only to partial votes of some of the voters; i.e., replace each vote v with a (possibly empty) partial order p, or a collection of pairwise comparisons. Let $\mathbf{p}$ denote this *partial profile*. How should one select a winner? Intuitively, we measure the quality of a given $\mathbf{p}$ by considering how far from optimal a could be in the worst case (i.e., given any *completion* or extension $\mathbf{v} \in C(\mathbf{p})$ of $\mathbf{p}$). The minimax optimal solution is any alternative that is nearest to optimal in the worst case. More formally:

$$\begin{aligned} \mathit{Regret}(a, \mathbf{v}) &= \mathit{max}_{a' \in A} s(a', \mathbf{v}) - s(a, \mathbf{v}) \\ \mathit{MR}(a, \mathbf{p}) &= \mathit{max}_{\mathbf{v} \in C(\mathbf{p})} \mathit{Regret}(a, \mathbf{v}) \\ \mathit{MMR}(\mathbf{p}) &= \mathit{min}_{a \in A} \mathit{MR}(a, \mathbf{p}) \\ a^*_{\mathbf{p}} &\in \operatorname*{argmin}_{a \in A} \mathit{MR}(a, \mathbf{p}) \end{aligned}$$

This is a natural robustness criterion: the minimax winner $a^*_{\mathbf{p}}$ provides us with the tightest possible bound on loss of "societal utility." MMR can be computed in polytime for a variety of voting rules, and can offer quite distinct recommendations compared to selecting among possible winners [9].

One might consider minimax regret to be too pessimistic, though we argue below that it is, in fact, a very effective driver of vote elicitation/active learning. MMR also fails to exploit distributional information P about voter preferences. With such a probabilistic model, one can instead select a winner by maximizing expected utility (MEU): $a^*_{\mathbf{p}} = \text{argmax} \sum_{\mathbf{v}} P(\mathbf{v}|\mathbf{p}) s(a, \mathbf{v})$. The investigation of algorithms for solving this computationally challenging problem for various combinations of voting rules and preference distributions is, in our opinion, a vital direction.

Notice that MEU ensures (Bayesian) optimality in the presence of a partial profile, but provides no guidance w.r.t. potential loss relative to choosing a winner with a complete profile $\mathbf{v}$. This stands in contrast to MMR, which tells us the potential value of adding new evidence to complete the vote profile. In the probabilistic case, *expected regret* is the most natural measure of loss regarding a proposed alternative a: $ER(a, \mathbf{p}) = \sum_{\mathbf{v}} P(\mathbf{v}|\mathbf{p}) \mathit{Regret}(a, \mathbf{v})$.[5] Notice, of course, that the same alternative $a^*_{\mathbf{p}}$ maximizes expected utility and minimizes expected regret; but ER is much more informative and useful for elicitation purposes.

4 Elicitation

Preference/vote elicitation is another critical process that has received insufficient attention in social choice. By explicitly articulating a notion of "societal" utility, and developing suitable probabilistic models, natural approaches to elicitation emerge that exploit the optimization criteria discussed above. Connections to active learning also become much clearer when adopting this perspective.

Without a probabilistic model P, MMR is probably the most natural criterion for robust selection of alternatives. But if MMR is too great, the potential error associated with *any* winner will be unacceptable. MMR can be reduced by asking some voter(s) some query(ies) about their preferences. In [9] we developed elicitation schemes that exploit the current solution to the minimax problem to determine appropriate voter-query pairs: on both synthetic and real-world voting and preference data, these methods performed extremely well, asking only a fraction of the queries that would be require to fully elicit voter rankings.[6] This is true despite the rather pessimistic worst-case results on the communication complexity of many voting rules. MMR also provides strong, distribution-free quality guarantees.

[3] This generalizes the *repeated insertion model* [2] for unconditional Mallows sampling.

[4] Necessary and possible winners don't actually prescribe general methods for selection.

[5] See Smith [12] who uses score-based regret.

[6] See Kalech et al. [5] for an alternative approach to elicitation.

In the probabilistic case, expected regret is the appropriate measure of loss, and optimal queries are those with maximum *expected value of information (EVOI)*. EVOI can be very difficult to compute in general, so again, as with MEU and ER computation, interesting challenges lay ahead in the effective (possibly approximate) computation of EVOI for various families of distributions and voting rules.

Interestingly, there are very useful ways of combining the probabilistic and regret-based perspectives. One difficulty with vote elicitation is that it is unrealistic to expect a fully interactive approach: no user u will want to answer a query, then wait for other users to answer their queries before the system returns with the next query for u. There is a fundamental tradeoff between amount of information elicited and the number of "query rounds" [5]. Probabilistic models can be used to help *batch* queries to assess this tradeoff. For instance, given a voting rule and a distribution, we may ask about the impact of asking m random users a small set of queries, e.g., "what are your top t alternatives?" For any t we can assess the posterior distribution over either MMR or ER to determine the depth t that makes the right tradeoff. That is, for given voting rules and families of distributions, we'd like effective techniques to compute, say, $E_P[MMR(\mathbf{p})|m, t]$, where expectation is taken over possible responses to the top-t queries from m users. Alternatively, one might favor a PAC-style analysis, deriving appropriate values for m and t such that $P(MMR(\mathbf{p}) < \varepsilon) \geq 1 - \delta$: in other words, for the selected m and t, with high probability $1-\delta$, MMR will be less than some small value ε if we ask m voters for their top-t candidates. Analysis of this type (for various classes of queries) can be used to drastically limit the number of rounds while keeping the total amount of elicited information small.[7]

5 Manipulation

Finally, we close by suggesting that the utility-theoretic and probabilistic perspectives can provide a much more nuanced analysis of manipulation. Most manipulation analysis addresses the question of whether a small coalition of voters can change the outcome of an election by misreporting their preferences under *some distribution* of the preferences of the electorate. Typically, this distribution is a point distribution in which the coalition knows the exact preferences of other voters. Probabilistic information is sometimes used, but usually only to analyze the odds that a manipulation exists *assuming complete knowledge* on the part of the manipulators.

We suggest that two different styles of analysis would be much more useful when considering the application of social choice in the domains described above. First, assuming that manipulators know the full preference profile is unrealistic. Of course, it would be equally unrealistic to assume no knowledge: instead we suggest that analyses should *restrict* the manipulators' knowledge in reasonable ways. For example, we may insist that the distribution over preferences known to the manipulators has some minimum entropy; or we could restrict knowledge of preferences to that obtainable using a small number of samples from the underlying distribution. Such analysis of the potential for manipulation should also be undertaken using realistic distributions of preferences as opposed to impartial culture and related models.

The second change in analysis is suggested by the use of societal utility measures. Intuitively, if a *small* coalition can change the outcome from the *true* winner a to an alternative b, then it is highly likely b had a reasonably high societal utility to begin with. So rather than asking whether specific voting rules are manipulable, we can instead ask how much "damage" can a small coalition do: in other words, what is the maximum regret $MR(b, \mathbf{p})$ or expected regret $ER(b, \mathbf{p})$ given partial knowledge $\mathbf{p}$ obtained by the manipulators. The susceptibility of a voting rule to manipulation can then be characterized by placing limits on the form of $\mathbf{p}$, maximizing these damage metrics over possible manipulations b, and maximizing or taking expectation w.r.t. $\mathbf{p}$ of some limited form. Here is just one concrete question of this form: given distribution P, what is $E_{\mathbf{p}[m]\sim P} \max_b ER(b, \mathbf{p})$, where $\mathbf{p}[m] \sim P$ refers to random sample of m votes from P. This type of analysis may provide a very different view of the manipulability of various voting rules.

Acknowledgements: Thanks to Yann Chevaleyre, J erôme Lang, and Nicolas Maudet for very engaging discussions on several of these broad topics (and some of the specific problems mentioned here).

References

[1] M. Braverman and E. Mossel. Sorting from noisy information. Manuscript, arXiv:0910.1191, 2009.

[2] J. P. Doignon, A. Pekec, and M. Regenwetter. The repeated insertion model for rankings: Missing link between two subset choice models. *Psychometrika*, 69(1):33–54, 2004.

[3] J. Huang and C. Guestrin. Riffled independence for ranked data. *NIPS 21*, pp.799–807, Vancouver, 2009.

[4] E. Hüllermeier, J. Fürnkranz, W. Cheng, and K. Brinker. Label ranking by learning pairwise preferences. *Artificial Intelligence*, 172(16-17):1897–1916, 2008.

[5] M. Kalech, S. Kraus, G. Kaminka and C. Goldman Practical voting rules with partial information. *Journal of Autonomous Agents and Multi-Agent Systems*, 22:151–182, 2011.

[6] J. Kleinberg, C. Papadimitriou, and P. Raghavan. Segmentation problems. *Journal of the ACM*, 51:263–280, 2004.

[7] T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. *IJCAI-11*, Barcelona, 2011. To appear.

[8] T. Lu and C. Boutilier. Learning mallows models with pairwise preferences. *ICML-11*, Bellevue, WA, 2011. To appear.

[9] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. *IJCAI-11*, Barcelona, 2011. To appear.

[10] J. Marden. *Analyzing and modeling rank data*. Chapman and Hall, 1995.

[11] M. Regenwetter, B. Grofman, A. A. J. Marley, and I. Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006.

[12] W. Smith. Range voting. http://www.math.temple.edu/~wds/homepage/rangevote.pdf, 2000.

[7] Preliminary results suggest that reasonable bounds can be derived for Borda scoring with Mallows models. Some relevant results on sorting complexity for Mallows models are developed in [1].