Subset Selection Via Implicit Utilitarian Voting

Ioannis Caragiannis University of Patras, Greece caragian@ceid.upatras.gr

Ariel D. Procaccia Carnegie Mellon University, USA arielpro@cs.cmu.edu

Abstract

How should one aggregate ordinal preferences expressed by voters into a measurably superior social choice? A well-established approach - which we refer to as *implicit utilitarian voting* — assumes that voters have latent utility functions that induce the reported rankings, and seeks voting rules that approximately maximize utilitarian social welfare. We extend this approach to the design of rules that select a subset of alternatives. We derive analytical bounds on the performance of optimal (deterministic as well as randomized) rules in terms of two measures, distortion and regret. Empirical results show that regret-based rules are more compelling than distortion-based rules, leading us to focus on developing a scalable implementation for the optimal (deterministic) regret-based rule. Our methods underlie the design and implementation of an upcoming social choice website.

1 Introduction

We are interested in the classic social choice problem of aggregating the preferences of a set of voters — represented as *rankings* over a set of alternatives — into a collective decision. Traditional social choice theory typically takes a *normative* approach, by specifying desirable axioms that the aggregation method (also known as a *voting rule*) should satisfy [Arrow, 1951]. In contrast, researchers in *computational social choice* [Brandt *et al.*, 2016] often advocate *quantitative* approaches to the same problem. The high-level idea is to identify a compelling objective function, and design voting rules that optimize this function.

Here we focus on a specific objective function: *utilitarian social welfare*. Specifically, we assume that each voter assigns a utility to each possible outcome, and the socially optimal outcome maximizes the sum of utilities. This sounds simple enough at first glance, but there is a major obstacle we must overcome: voters' preferences are expressed as *ordinal* preferences (rankings), rather than *cardinal* preferences (utilities). While this reduces the cognitive load on voters, and Swaprava Nath Carnegie Mellon University, USA swapravn@cs.cmu.edu

Nisarg Shah Carnegie Mellon University, USA nkshah@cs.cmu.edu

makes preference elicitation much easier, it does seem to be at odds with the utilitarian viewpoint.

Procaccia and Rosenschein [2006] reconcile these differences via an approach that we refer to as implicit utilitarian voting.¹ They propose that voters have latent utility functions, and report rankings that are consistent with these utilities, that is, the voters rank the alternatives by their utility. The performance of a voting rule — which can only access the submitted rankings, not the implicit utility functions can then be quantified via a measure called distortion: the worst-case (over utility functions consistent with the reported profile of rankings) ratio between the social welfare of the optimal (welfare-maximizing) alternative, and the social welfare of the alternative selected by the voting rule. While Procaccia and Rosenschein focus on analyzing the distortion of existing voting rules, Boutilier et al. [2015] design voting rules that minimize distortion. In particular, they bound the worst-case distortion, and show that the distortion-minimizing (randomized) voting rule can be implemented in polynomial time.

The work of Boutilier *et al.* [2015] provides a good understanding of optimized aggregation of rankings from the utilitarian viewpoint — but only when a single alternative is selected by the voting rule. Indeed, this understanding does not extend to common applications that require selection of a subset of alternatives, such as choosing a committee, or selecting restaurants for the next four group lunches. Our goal is therefore to

... build on the utilitarian approach to design optimal voting rules for selecting a subset of alternatives, and understand the guarantees they provide, as well as their performance in practice.

We make four main contributions. First, on a conceptual level, we introduce the additive notion of *regret* into the implicit utilitarian voting setting, as an alternative to the multiplicative notion of distortion. Second, in Section 3, we derive worst-case bounds on the distortion and regret of optimal deterministic and randomized voting rules. Third, in Section 4, we compare the *worst-case-optimal* deterministic voting rules with respect to distortion and regret — denoted

¹Cf. *utilitarian voting*, which has sporadically been used to refer to both approval voting and range voting.

 f_{dist}^* and f_{reg}^* , respectively — with a slew of well-known voting rules, in terms of *average-case* distortion and regret, using experiments on synthetic and real data. We find that f_{reg}^* outperforms all other rules on average, even when measuring distortion! Fourth, in Section 5, we develop a scalable implementation for f_{reg}^* (which, we show, is \mathcal{NP} -hard to compute).

1.1 Direct Real-World Implications

Research in computational social choice has frequently been justified by potential applications in multiagent systems. But recently researchers have begun to realize that, arguably, the most exciting products of this research are computer programs that help *humans* make decisions via AI-driven algorithms. One example is *Spliddit* (www.spliddit.org), a fair division website [Goldman and Procaccia, 2014]. In the voting space, existing examples include *Whale* (whale3.noiraudes.net/whale3/) and *Pnyx* (pnyx.dss.in.tum.de) — but these websites generally adopt the axiomatic viewpoint.

Since May 2015, some of us have been working on the design and implementation of a new not-for-profit social choice website, RoboVote (www.robovote.org), which is scheduled to launch in May 2016. The novelty of RoboVote is that it relies on optimization-based approaches. For the case of objective votes - when a ground truth ranking of the alternatives exists (e.g., the order of different stocks by the relative change in their prices tomorrow) - RoboVote implements voting rules that pinpoint the most likely best alternative [Young, 1988], or the set most likely to contain it [Procaccia et al., 2012]. For the case of subjective votes — the classic setting which is the focus of this paper, with applications to everyday scenarios such as a group of friends selecting a movie to watch or a restaurant to go to - we use the results of Boutilier et al. [2015] to select a single alternative. But, previously, the extension to subset selection was unavailable — this is precisely the motivation for the work described herein. Based on the results of Sections 4 and 5, we have implemented the deterministic regret minimization rule on RoboVote.

1.2 Related Work

In addition to the aforementioned papers [Procaccia and Rosenschein, 2006; Boutilier *et al.*, 2015], several other papers employ the notion of distortion to quantify how close one can get to maximizing utilitarian social welfare when only ordinal preferences are available [Caragiannis and Procaccia, 2011; Anshelevich *et al.*, 2015; Anshelevich and Sekar, 2016]. In particular, Anshelevich *et al.* [2015] study the same setting as Boutilier *et al.* [2015], but in addition assume the preferences of voters are consistent with distances in a metric space. We refer the reader to the paper by Boutilier *et al.* [2015, Section 1.2] for a thorough discussion of work (in philosophy, economics, and social choice theory) related to implicit utilitarian voting more broadly.

There is quite a bit of work in computational social choice on voting rules that select subsets of alternatives. Typically it is assumed that ordinal preferences are translated into a position-based score for each alternative (in contrast to our work). Just to give a few examples, under the Chamberlin-Courant method, each voter assigns a score to a set equal to the highest score of any alternative in the set, and the (computationally hard) objective is to choose a subset of size k that maximizes the sum of scores [Chamberlin and Courant, 1983; Procaccia *et al.*, 2008]. Skowron *et al.* [2015] generalize the way in which the score of a voter for a subset of alternatives is computed. The *budgeted social choice* framework of Lu and Boutilier [2011a] is more general in that the number of alternatives to be selected is not fixed; rather, each alternative has a cost that must be paid to add it to the selection.

2 The Model

Let $[t] = \{1, \ldots, t\}$. Let A be the set of *alternatives*, and denote m = |A|. Let N = [n] be the set of *voters*. Let $\mathcal{L} = \mathcal{L}(A)$ denote the set of rankings over the alternatives. Each voter $i \in [n]$ submits a ranking $\sigma_i \in \mathcal{L}$ over the alternatives, and which can alternatively be seen as a permutation of A. Therefore, $\sigma_i(a)$ is the position in which voter i ranks alternative a (1 is best, m is worst). Moreover, $a \succ_{\sigma_i} b$ denotes that voter i prefers alternative a to alternative b. The collection of voters' (submitted) rankings is called the *preference profile*, and denoted by $\vec{\sigma} \in \mathcal{L}^n$.

We assume the rankings are induced by comparisons between the voters' underlying utilities. For $i \in N$ and $a \in A$, let $u_i(a) \in [0, 1]$ be the utility of voter *i* for alternative *a*. As in previous papers [Boutilier *et al.*, 2015; Caragiannis and Procaccia, 2011], we assume that the utilities are normalized such that $\sum_{a \in A} u_i(a) = 1$ for all $i \in N$. The collection of voter utilities, denoted \vec{u} , is called the *utility profile*. We say that utility profile \vec{u} is consistent with preference profile $\vec{\sigma}$ — denoted $\vec{u} \triangleright \vec{\sigma}$ — if for all $a, b \in A$ and $i \in N, a \succ_{\sigma_i} b$ implies $u_i(a) \ge u_i(b)$.

Next we need to define the utility of a voter for a set of alternatives. For $S \subseteq A$, we define $u_i(S) = \max_{a \in S} u_i(a)$, that is, each voter derives utility for his favorite alternative in the set; this is in the same spirit as previous papers on set selection [Chamberlin and Courant, 1983; Monroe, 1995; Procaccia *et al.*, 2008; Lu and Boutilier, 2011a]. Then, the (utilitarian) *social welfare* of S given the utility profile \vec{u} is $sw(S, \vec{u}) = \sum_{i=1}^{n} u_i(S)$.

We are interested in voting rules that, given a preference profile, select a subset of given cardinality k.² Therefore, it will be useful to denote $\mathcal{A}_k = \{S \subseteq A : |S| = k\}$. In order to unify notation, we directly define a *randomized* voting rule as a function $f : \mathcal{L}^n \to \Delta(\mathcal{A}_k)$, that is, the rule is allowed to select alternatives randomly, and formally $f(\vec{\sigma})$ is a *probability distribution* over \mathcal{A}_k . A *deterministic* voting rule simply gives probability 1 to a specific subset.

A voting rule can only access the preference profile $\vec{\sigma}$, yet the goal is to maximize social welfare with respect to the latent utility function $\vec{u} \triangleright \vec{\sigma}$. We study two notions that quantify how well a rule achieves this goal: distortion and regret.

The distortion [Procaccia and Rosenschein, 2006] of a

²Formally, this is a special case of *social choice correspondences* with fixed output cardinality [Campbell and Kelly, 1996].

(randomized) voting rule f on a preference profile $\vec{\sigma}$ is

$$\operatorname{dist}(f, \vec{\sigma}) = \sup_{\vec{u} \triangleright \vec{\sigma}} \frac{\max_{S \in \mathcal{A}_k} \operatorname{sw}(S, \vec{u})}{\mathbb{E}[\operatorname{sw}(f(\vec{\sigma}), \vec{u})]}$$

In words, it is the worst-case — over utility profiles consistent with the given preference profile — ratio between the social welfare of the best subset, and the expected social welfare of the subset selected by the voting rule. We define the distortion of a voting rule f by taking the worst case over preference profiles: $dist(f) = \max_{\vec{\sigma} \in \mathcal{L}^n} dist(f, \vec{\sigma})$.

The second measure is *regret*. While it has not been studied as part of the agenda of implicit utilitarian voting, it has been explored in other social choice settings, especially partial preferences [Lu and Boutilier, 2011b]; similar measures have been extensively studied in decision theory and machine learning [Blum and Mansour, 2007; Bubeck and Cesa-Bianchi, 2012]. The *regret* of a (randomized) voting rule f on a preference profile $\vec{\sigma}$ is given by

$$\operatorname{reg}(f,\vec{\sigma}) = \frac{1}{n} \cdot \sup_{\vec{u} \triangleright \vec{\sigma}} \left(\max_{S \in \mathcal{A}_k} \operatorname{sw}(S,\vec{u}) - \mathbb{E}[\operatorname{sw}(f(\vec{\sigma}),\vec{u})] \right).$$

As before, define the regret of a rule f to be $\operatorname{reg}(f) = \max_{\vec{\sigma} \in \mathcal{L}^n} \operatorname{reg}(f, \vec{\sigma})$. We divide by n because the total (worst-case) regret of any voting rule f is provably linear in n (so this is *per vote* regret). Note that distortion is a multiplicative measure of loss, whereas regret is its additive version.

3 Worst-Case Bounds

In this section we provide bounds on worst-case distortion and regret, for both deterministic and randomized voting rules. Boutilier *et al.* [2015] show that for selecting a single winner (k = 1), we can achieve $O(\sqrt{m} \cdot \log^* m)$ distortion using a randomized rule, where $\log^* m$ is the iterated logarithm of m (the number of alternatives). This bound is asymptotically almost tight: they also show that the worstcase distortion is always $\Omega(\sqrt{m})$.

For a large k, though, one can hope for a better bound. Clearly, when k = m there is only one voting rule (which selects every alternative), and its distortion is 1. More generally, it is easy to show that the voting rule f that selects a subset from \mathcal{A}_k uniformly at random has $dist(f) \leq m/k$. However, since we can already achieve $O(\sqrt{m} \cdot \log^* m)$ distortion for k = 1, a bound of m/k provides an improvement only for $k = \Omega(\sqrt{m}/\log^* m)$. Can we achieve better distortion for smaller values of k as well? It is not even clear whether the optimal worst-case distortion should monotonically decrease in k, because as our flexibility grows with k, so does the flexibility of the welfare-maximizing solution. In fact, a part of our main result shows that the worst-case distortion remains $\Omega(\sqrt{m})$ for all values of k up to $\Theta(\sqrt{m})$.

Theorem 1. Let m = |A|, and let k be the number of alternatives to be selected.

1. Distortion, deterministic rules: There exists a deterministic voting rule f^* with dist $(f^*) \le 1 + m (m-k)/k$. Moreover, for every deterministic voting rule f,

$$dist(f) \geq \begin{cases} 1 + \frac{m(m-3k)}{6k} & \text{if } k \leq \frac{m}{9}, \\ 1 + m & \text{if } \frac{m}{9} < k \leq \frac{m}{2}, \\ 1 + \frac{m(m-k)}{k} & \text{otherwise.} \end{cases}$$

These bounds are tight up to a constant factor of 7.3.

2. Distortion, randomized rules: There exists a randomized voting rule f* such that

$$\operatorname{dist}(f^*) \leq \begin{cases} 2\sqrt{m \cdot H_m} & \text{if } k \leq \frac{2 \cdot m \cdot H_m}{m + H_m}, \\ 4\sqrt{m \cdot k} & \text{if } \frac{2 \cdot m \cdot H_m}{m + H_m} < k \leq \left(\frac{m}{4}\right)^{\frac{1}{3}}, \\ \frac{m}{k} & \text{otherwise}, \end{cases}$$

where $H_m = \Theta(\log m)$ is the m^{th} harmonic number. Moreover, for every randomized voting rule f,

$$dist(f) \geq \begin{cases} \frac{\sqrt{m}}{2} & \text{if } k \leq \frac{m \cdot (\sqrt{m} - 1)}{m - 1} \approx \sqrt{m}, \\ \frac{m}{k + m/k} & \text{otherwise.} \end{cases}$$

These bounds are tight up to a factor of $9.4 \cdot m^{1/6}$.

3. **Regret, deterministic rules:** There exists a deterministic voting rule f^* such that

$$\operatorname{reg}(f^*) \leq \begin{cases} \frac{1}{2} & \text{if } k \leq \frac{m}{2}, \\ 1 - \frac{k}{m} & \text{otherwise}, \end{cases}$$

and this upper bound is completely tight.

4. Regret, randomized rules: There exists a randomized voting rule f^* such that $\operatorname{reg}(f^*) \leq \frac{1}{2} \cdot (1 - \frac{k^2}{m^2})$. Moreover, for every randomized voting rule f,

$$\operatorname{reg}(f) \geq \begin{cases} \frac{1}{4} & \text{if } k \leq m/2 \\ \frac{1}{2} \cdot \frac{k}{m} \left(1 - \frac{k}{m}\right) & \text{otherwise.} \end{cases},$$

All the upper bounds above can be achieved via polynomialtime algorithms.



Figure 1: The upper and lower bounds on worst-case distortion and regret for m = 100.

The bounds presented above are simplified forms of the exact bounds that we derive. Figure 1 shows our exact bounds for $m = 100.^3$

The intricate proof of Theorem 1 appears in the full version of the paper.⁴ Below, we just sketch one part of the proof that we find especially interesting.

³The second upper bound in part 2 of Theorem 1 (which increases with k) does not play a role unless m is very large.

⁴Available at: www.cs.cmu.edu/~arielpro/papers.html

Proof sketch of the upper bound in part 2 of Theorem 1. Our construction builds on the one used by Boutilier *et al.* [2015] for k = 1, but uses additional tools and introduces novel techniques. As mentioned at the beginning of this section, choosing a set uniformly at random from \mathcal{A}_k (under which the marginal probability of every alternative being chosen is k/m) has distortion at most m/k. However, this approach does not work well if some alternatives are significantly better than others.

In that case, one may wish to choose the alternatives with probabilities proportional to their "quality". For $a \in A$, let us define its quality by its harmonic score $har(a) = \sum_{i \in [n]} 1/\sigma_i(a)$. Then, we wish to choose alternative a with marginal "probability" $k \cdot har(a) / \sum_{b \in A} har(b)$. Note that this quantity may be greater than 1. Moreover, this approach fails when all sets are almost equally good. Hence, we employ a combination of the two approaches.

Fix $0 \le \alpha \le 1$, and for an alternative $a \in A$ define

$$p_a = \alpha \cdot \frac{k}{m} + (1 - \alpha) \cdot \frac{k \cdot \operatorname{har}(a)}{\sum_{b \in A} \operatorname{har}(b)}.$$
 (1)

Using the bihierarchy extension [Budish *et al.*, 2013] of the Birkhoff-von Neumann theorem [Birkhoff, 1946; von Neumann, 1953], we can show that there exists a distribution over A_k under which the marginal probabilities of selected alternatives are consistent with Equation (1) if and only if

$$\forall a \in A, 0 \leq p_a \leq 1 \quad \text{and} \quad \sum_{a \in A} p_a = k.$$

Suppose such a distribution D exists. Consider a preference profile $\vec{\sigma}$ and a utility profile $\vec{u} \triangleright \vec{\sigma}$. Let $S^* \in \arg \max_{S \in \mathcal{A}_k} \mathrm{sw}(S, \vec{u})$. Define

$$X = \sqrt{\frac{H_m}{m} \cdot \frac{\alpha}{1 - \alpha}},$$

where $H_m = \sum_{t=1}^m 1/t$ is the m^{th} harmonic number. Note that $\sum_{a \in A} har(a) = n \cdot H_m$. Now, consider two cases.

Case 1: Suppose $sw(S^*, \vec{u}) \leq n \cdot X$. Then,

$$\mathbb{E}_{S\sim D}[\mathbf{sw}(S, \vec{u})] = \sum_{S\in\mathcal{A}_k} \Pr_D[S] \cdot \left(\sum_{i=1}^n \max_{a\in S} u_i(a)\right)$$
$$\geq \sum_{i=1}^n \left(\sum_{S\in\mathcal{A}_k} \Pr_D[S] \cdot \frac{\sum_{a\in S} u_i(a)}{k}\right)$$
$$= \frac{1}{k} \sum_{i=1}^n \sum_{a\in A} u_i(a) \cdot \Pr_{S\sim D}[a\in S]$$
$$\geq \frac{1}{k} \sum_{i=1}^n \sum_{a\in A} u_i(a) \cdot \alpha \cdot \frac{k}{m} = \alpha \cdot \frac{n}{m}.$$

Hence, the distortion is

$$\frac{\mathsf{sw}(S^*, \vec{u})}{\mathbb{E}_{S \sim D}[\mathsf{sw}(S, \vec{u})]} \le \frac{n \cdot X}{\alpha \cdot n/m} = \frac{X \cdot m}{\alpha} = \sqrt{\frac{m \cdot H_m}{\alpha \cdot (1 - \alpha)}}.$$

Case 2: Suppose $\mathfrak{sw}(S^*, \vec{u}) > n \cdot X$. Then, for each alternative $a \in S^*$, let N_a denote the subset of voters who rank a above any other alternative of S^* , i.e.,

$$N_a = \{ i \in [n] : \forall b \in S^* \setminus \{a\}, \ a \succ_{\sigma_i} b, \}.$$

Let $\operatorname{sw}_{N_a}(S, \vec{u})$ denote the welfare of the voters in N_a for the set of alternatives S under the utility profile \vec{u} . Let T_a denote the total utility that agents in N_a have for alternative a, i.e., $T_a = \sum_{i \in N_a} u_i(a)$. It can be shown (although it is nontrivial) that $\operatorname{har}(a) \geq T_a$ for all $a \in A$. Because $\{N_a\}_{a \in S^*}$ is a partition of the set of voters, we have

$$\begin{split} \mathbb{E}_{S\sim D}[\mathbf{sw}(S, \vec{u})] &= \mathbb{E}_{S\sim D} \left[\sum_{a \in S^*} \mathbf{sw}_{N_a}(S, \vec{u}) \right] \\ &\geq \sum_{a \in S^*} T_a \cdot \Pr_{S\sim D}[a \in S] \\ &\geq \sum_{a \in S^*} T_a \cdot (1 - \alpha) \cdot \frac{k \cdot \operatorname{har}(a)}{\sum_{b \in A} \operatorname{har}(b)} \\ &\geq \frac{(1 - \alpha) \cdot k}{n \cdot H_m} \cdot \sum_{a \in S^*} (T_a)^2 \geq \frac{1 - \alpha}{n \cdot H_m} \cdot \left(\sum_{a \in S^*} T_a \right)^2 \\ &= \frac{1 - \alpha}{n \cdot H_m} \cdot (\mathbf{sw}(S^*, \vec{u}))^2 \,. \end{split}$$

Here, the fourth transition uses $har(a) \ge T_a$, the fifth transition uses the power-mean inequality, and the final transition uses $sw(S^*, \vec{u}) = \sum_{a \in S^*} T_a$. Now, the distortion is

$$\frac{\mathrm{sw}(S^*,\vec{u})}{\mathbb{E}_{S\sim D}[\mathrm{sw}(S,\vec{u})]} \leq \frac{n \cdot H_m}{(1-\alpha) \cdot \mathrm{sw}(S^*,\vec{u})} < \sqrt{\frac{m \cdot H_m}{\alpha \cdot (1-\alpha)}},$$

where the final transition uses our assumption $sw(S^*, \vec{u}) > n \cdot X$ along with the definition of X.

Combined analysis: In both cases, the distortion is at most $\sqrt{mH_m/(\alpha(1-\alpha))}$. The final step involves choosing the optimal value of α by minimizing this quantity subject to our constraints: $p_a \leq 1$ for all $a \in A$. Simplifying the bound obtained along with our universal distortion bound of m/k yields the required upper bound.

4 Empirical Comparisons

In Section 3 we provided analytical results for both deterministic and randomized rules. In our view, randomized rules are especially practicable when the output distribution is sampled multiple times, or when the voters are well-informed, or when the voters are indifferent about the outcome (e.g., they are software agents). Moreover, we believe that the results for randomized rules are of substantial theoretical interest. But our work is partly driven by its direct applications in RoboVote (see Section 1.1), which does not satisfy the above conditions. This leads us to use deterministic voting rules, which is what we focus on hereinafter.

Let f_{dist}^* and f_{reg}^* be the *deterministic* rules that minimize the worst-case distortion and regret, respectively, on every given preference profile. The deterministic results of Section 3 establish upper and lower bounds on their *worst-case*





Figure 4: Preference profiles from the Sushi and the T-Shirt datasets, uniformly random consistent utility profiles.

distortion/regret. In this section, we evaluate their *averagecase* performance on simulated as well as real data, and compare them against nine well-known voting rules: plurality, approval voting, Borda count, STV, Kemeny's rule, the maximin rule, Copeland's rule, Bucklin's rule, and Tideman's rule.⁵

We perform three experiments: (i) choosing a utility profile uniformly at random from the simplex of all utility profiles, (ii) drawing a real-world utility profile from the Jester datasets [Goldberg *et al.*, 2001], and (iii) drawing a realworld preference profile from the PrefLib datasets [Mattei and Walsh, 2013], and choosing a consistent utility profile uniformly at random. For each experiment, we have 8 voters and 10 alternatives, and test for $k \in [4]$.⁶ For each setting, we perform 10 000 random simulations, and measure both distortion and regret for the *actual utility profile*, as opposed to the worst-case utility profile. The figures show the average performance with 95% confidence intervals.

In all of our simulations, we observed that three of the classical voting rules stand out: Borda count performs well for choosing a single alternative (but not for choosing larger subsets) whereas plurality and STV perform well for choosing larger subsets (but not for choosing a single alternative). Hence, all of our graphs specifically distinguish these three rules in addition to f_{dist}^* and f_{reg}^* .

Figure 2 shows the results for the first experiment where we choose the utility profile uniformly at random. Figure 3 shows the results for the second experiment where real-world utility profiles are drawn from one of the Jester datasets, in which more than 50 000 voters rated 150 jokes on a realvalued scale; the results from the other Jester dataset are almost identical. Finally, Figure 4 shows the results for the third experiment where real-world preference profiles are drawn from the Sushi dataset (5 000 voters ranking 100 different kinds of sushi) and the T-Shirt dataset (30 voters ranking 11 T-shirt designs) from PrefLib. Experiments on other datasets from PrefLib (AGH Course Selection, Netflix, Skate, and Web Search) yielded similar results.

Right off the bat, one can observe that the average-case distortion and regret values are much lower than their worst-case counterparts. For example, average regret is generally lower than 0.1 — compare with the tight worst-case deterministic bound of 1/2 for $k \le m/2$.

Much to our surprise, in all of our experiments, f_{reg}^* outperforms f_{dist}^* in terms of both average-case distortion (multiplicative loss) *and* regret (additive loss). While both measures of loss have been studied extensively in the literature,

⁵For the score-based rules, the *k*-subset is selected by picking the top k alternatives based on their scores.

⁶In RoboVote, we expect typical instances to have few voters and alternatives. But we chose m > 2k because otherwise the problem would be trivial: for $k \ge m/2$, picking the top k alternatives based on plurality scores is optimal for both distortion and regret.

we are not aware of any previous work that compares the two approaches. At least in our social choice domain, the regretbased approach is clearly better on average.

Moreover, in all cases but one (k = 1 in the Jester experiment), f_{reg}^* also outperforms all the classical voting rules under consideration. We therefore conclude that, on random as well as on real-world instances, f_{reg}^* provides superior performance in terms of social welfare maximization.

5 Computation and Implementation

In this section, we analyze and compare the two deterministic optimal rules — f_{dist}^* and f_{reg}^* — from a computational viewpoint. Selecting optimal subsets turns out to be challenging, as both rules are \mathcal{NP} -hard to compute; the proof of this nontrivial result appears in the full version.

Theorem 2. Given a preference profile $\vec{\sigma}$ and an integer k, computing a k-subset of alternatives that has the minimum distortion or the minimum regret on $\vec{\sigma}$ is \mathcal{NP} -hard.

Given that f_{reg}^* outperforms f_{dist}^* in the experiments of Section 4, and that both rules are computationally hard, f_{reg}^* stands out as the clear choice for implementation in our website RoboVote. We therefore devoted our efforts to developing a scalable implementation for f_{reg}^* .

The first step is to simplify the description of f_{reg}^* . Given a ranking σ and an alternative $a \in A$, recall that $\sigma(a)$ denotes the position of a in σ . For a set $S \subseteq A$, let $\sigma(S) = \min_{a \in S} \sigma(a)$. For sets $S, T \subseteq A$, we say $T \succ_{\sigma} S$ if $\sigma(T) < \sigma(S)$, i.e., if there exists an alternative in T that is preferred in σ to every alternative in S. Using these notations, it is relatively straightforward to prove that

$$f_{\mathsf{reg},k}^{*}(\vec{\sigma}) = \underset{T \in \mathcal{A}_{k}}{\operatorname{arg\,min}} \max_{S \in \mathcal{A}_{k}} \sum_{i \in N: S \succ_{\sigma_{i}} T} \frac{1}{\sigma_{i}(S)}.$$
 (2)

To better understand this equation, we consider the special case of k = 1. In this case,

$$f^*_{\mathsf{reg}}(\vec{\sigma}) \in \operatorname*{arg\,min}_{a \in A} \max_{b \in A} \sum_{i \in [n]: b \succ_{\sigma_i} a} \frac{1}{\sigma_i(b)}.$$

Note that this voting rule is very similar to the classical maximin rule: replacing $1/\sigma_i(b)$ with 1 would yield the maximin rule. Thus, in some sense, this is a smooth version of the maximin rule, where the "victory" of *b* over *a* in voter *i*'s vote is weighted by the strength of *b* in this vote (measured by $1/\sigma_i(b)$). In our view, this intuitive structure makes f_{reg}^* even more compelling.

We now briefly describe six approaches we have developed for computing f_{reg}^* :

- 1. Naïve: This uses Equation (2), and requires $\Omega(n \cdot {\binom{m}{k}}^2)$ operations, which is prohibitive even for small m.
- Submodular: The regret for set S in choosing set T, i.e., ∑_{i∈[n]:S≻σ_iT} 1/σ_i(S), is submodular in S. Hence, for each T ∈ A_k we can optimize over S ∈ A_k using any algorithm for the submodular maximization subject to cardinality constraint (SMCC) problem. We use the SFO toolbox for Matlab [Krause, 2010].

- 3. Submodular+Greedy: This improves the previous approach by first computing a 1 1/e greedy approximation to the SMCC instance for set *T*, and pruning *T* if this is already greater than the best regret found so far.
- 4. MultiILP: Instead of using SMCC, for each $T \in \mathcal{A}_k$ we optimize over $S \in \mathcal{A}_k$ by solving an integer linear program (ILP) with roughly $n \cdot m$ variables and $n \cdot m^2$ constraints. Note that $\binom{m}{k}$ such ILPs need to be solved.
- 5. MultiILP+Greedy: This improves the MultiILP approach by using a greedy pruning procedure as before.
- SingleILP: This approach solves a single but huge ILP with ^(m)/_k additional constraints.

Figure 5 shows the average running times of these approaches (and 95% confidence intervals) over 10000 instances with n = 15, k = 3, and m varying from 10 to 50.⁷ The experiments were performed on a single machine with quad-core 2.9 GHz CPU and 32 GB RAM. A time limit of 2 minutes was set because a running time greater than this would not be helpful for our website, where the results need to be delivered quickly to the users. While the greedy pruning procedure does help reduce the running time of both the Submodular and MultiILP approaches, SingleILP still computes f_{reg}^* much faster than any other approach, solving instances with 50 alternatives in less than 10 seconds. We have therefore implemented SingleILP on RoboVote.



Figure 5: Running times of six approaches to computing f^*_{rea} .

6 Discussion

We find it exciting that new theoretical questions in computational social choice are driven by concrete real-world applications. And while research in the field is often motivated by potential applications to multiagent systems, we focus on helping *people* — not software agents — make joint decisions.

We also remark that we consider the empirical dominance of $f_{\rm reg}^*$, in terms of both regret and (surprisingly) distortion, to be especially significant. It would be interesting to understand, on a theoretical level, why this happens. A promising starting point is to derive analytical bounds on the averagecase distortion of $f_{\rm dist}^*$ and $f_{\rm reg}^*$ under uniformly random utility profiles.

⁷The running time scales linearly in n, and increases with $\binom{m}{k}$.

Acknowledgments

This work was supported in part by NSF grants CCF-1215883, CCF-1525932, and IIS-1350598; and by a Sloan Research Fellowship.

References

- [Anshelevich and Sekar, 2016] E. Anshelevich and S. Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence* (AAAI), 2016. Forthcoming.
- [Anshelevich et al., 2015] E. Anshelevich, O. Bhardwaj, and J. Postl. Approximating optimal social choice under metric preferences. In Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), pages 777–783, 2015.
- [Arrow, 1951] K. Arrow. Social Choice and Individual Values. Wiley, 1951.
- [Birkhoff, 1946] G. Birkhoff. Three observations on linear algebra. Universidad Nacional de Tucumán, Revista A, 5:147–151, 1946.
- [Blum and Mansour, 2007] A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 4. Cambridge University Press, 2007.
- [Boutilier *et al.*, 2015] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [Brandt et al., 2016] F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors. *Hanbook of Compu*tational Social Choice. Cambridge University Press, 2016.
- [Bubeck and Cesa-Bianchi, 2012] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [Budish *et al.*, 2013] E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- [Campbell and Kelly, 1996] D. E. Campbell and J. S. Kelly. Arrovian social choice correspondences. *International Economic Review*, 37(4):803–823, 1996.
- [Caragiannis and Procaccia, 2011] I. Caragiannis and A. D. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9– 10):1655–1671, 2011.
- [Chamberlin and Courant, 1983] J. R. Chamberlin and P. N. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983.
- [Goldberg *et al.*, 2001] K. Y. Goldberg, T. Roeder, D. Gupta, and C. Perkins. Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval*, 4(2):133–151, 2001.

- [Goldman and Procaccia, 2014] J. Goldman and A. D. Procaccia. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges*, 13(2):41–46, 2014.
- [Krause, 2010] A. Krause. SFO: A toolbox for submodular function optimization. *Journal of Machine Learning Research*, 11:1141–1144, 2010.
- [Lu and Boutilier, 2011a] T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280– 286, 2011.
- [Lu and Boutilier, 2011b] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 287–293, 2011.
- [Mattei and Walsh, 2013] N. Mattei and T. Walsh. Preflib: A library of preference data. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory* (*ADT*), pages 259–270, 2013.
- [Monroe, 1995] B. L. Monroe. Fully proportional representation. American Political Science Review, 89(4):925–940, 1995.
- [Procaccia and Rosenschein, 2006] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop* on Cooperative Information Agents (CIA), pages 317–331, 2006.
- [Procaccia *et al.*, 2008] A. D. Procaccia, J. S. Rosenschein, and A. Zohar. On the complexity of achieveing proportional representation. *Social Choice and Welfare*, 30(3):353–362, 2008.
- [Procaccia et al., 2012] A. D. Procaccia, S. J. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proceedings of the 28th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 695–704, 2012.
- [Skowron *et al.*, 2015] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2131–2137, 2015.
- [von Neumann, 1953] J. von Neumann. A certain zero-sum two-person game equivalent to the optimal assignment problem. In W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 2, pages 5–12. 1953.
- [Young, 1988] H. P. Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.