
Electing the Most Probable Without Eliminating the Irrational: Voting Over Intransitive Domains

Edith Elkind

University of Oxford, UK
elkind@cs.ox.ac.uk

Nisarg Shah

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

Abstract

Picking the best alternative in a given set is a well-studied problem at the core of social choice theory. In some applications, one can assume that there is an objectively correct way to compare the alternatives, which, however, cannot be observed directly, and individuals' preferences over the alternatives (votes) are noisy estimates of this ground truth. The goal of voting in this case is to estimate the ground truth from the votes. In this paradigm, it is usually assumed that the ground truth is a ranking of the alternatives by their true quality. However, sometimes alternatives are compared using not one but multiple quality parameters, which may result in cycles in the ground truth as well as in the preferences of the individuals. Motivated by this, we provide a formal model of voting with possibly intransitive ground truth and preferences, and investigate the maximum likelihood approach for picking the best alternative in this case. We show that the resulting framework leads to polynomial-time algorithms, and also approximates the corresponding \mathcal{NP} -hard problems in the classic framework.

1 INTRODUCTION

Typically, voting rules are viewed as vehicles for aggregating subjective preferences of individuals into a consensus or societal preference. However, another paradigm of voting theory, which dates back to Marquis de Condorcet [12], became increasingly popular in recent years, motivated in part by its relevance to the design of crowdsourcing platforms and human computation systems [15]. Condorcet suggested that votes cast by the individuals should be viewed as noisy estimates of an underlying objective ground truth—a ranking of the available alternatives by their true quality, and the aim of voting should be to aggregate the votes in order to uncover the ground truth

and thereby pick the best alternative. He proposed a simple approach for modeling the noise present in individuals' votes, which is known today as Mallows' model [19]. In this model, every voter compares each pair of alternatives independently, and orders them correctly (as in the ground truth) with a fixed probability $p > 1/2$, and incorrectly with probability $1-p$. If the generated vote contains cycles, it is discarded and the process restarts, continuing until the pairwise comparisons form a total order over the alternatives. While this model is somewhat unrealistic [20], it is widely used, in part because it provides control of the level of noise in the votes through a single parameter.

However, in many applications, alternatives are compared using not one, but multiple quality parameters [22, 29]. Under multi-criteria decision making, the preference relation that arises from the pairwise comparisons may contain cycles. Such preferences are known as tournaments. Also, when the number of alternatives is large (e.g., in many human computation systems), it is hard for voters to submit a total order over the alternatives. Hence, most systems employ vote elicitation techniques where the individuals iteratively submit parts of their preference, such as pairwise comparisons or partial orders. Bounded rationality of voters may again lead to cyclic preferences in this case. There are also settings where a voter may in fact be a meta-voter, representing a group of individuals (e.g., a country or region). Synthesizing preferences of the people in a group may also lead to cyclic preference for the meta-voter.

Motivated by this, we introduce a variant of Mallows' model where both the ground truth and the noisy preferences generated may be tournaments rather than rankings of the alternatives. In this model, the vote generation process described above simplifies: there is no need to restart the process if the generated vote has cycles. Consequently, the pairwise comparisons are independent of each other, resulting in a more tractable model: indeed, it appears that Young [30, p. 1238] in his analysis of Condorcet's approach to choosing the most likely winner uses the tournament-based model in his calculations, even though his intention was to study the ranking-based model.

Recently, Procaccia et al. [23] have formalized, corrected, and extended Young’s analysis of the optimal rule to pick the best alternative. They focused on the limiting case of Mallows’ model where the noise is very high ($p \rightarrow 1/2$), as they were motivated by crowdsourcing settings, in which this is often the case. As a side result, they have also analyzed the other extreme case of very low noise ($p \rightarrow 1$).

Our Contribution: We introduce the tournament variant of Mallows’ model and show that the most likely winners in our model can be identified in polynomial time for a given value of the noise parameter, as well as in the limiting cases of extremely high and low noise; this is in sharp contrast with the ranking-based model.¹ We then focus on the limiting cases of the tournament-based model and show empirically that they provide a good approximation for the corresponding cases of the ranking-based model. As a side result we prove that for winner determination, Tideman’s rule [25, pp. 199-201] (closely related to the high-noise setting in the tournament-based model) is a 2-approximation of Kemeny’s rule (closely related to the high-noise setting in the ranking-based model), a result that may be of independent interest to the social choice community. Finally, we propose an agnostic voting rule that circumvents the problem of not knowing the noise parameter and returns the set of alternatives that are MLE at some value of the parameter. Using simulations, we show that this rule is quite decisive, i.e., returns very few alternatives.

2 RELATED WORK

The maximum likelihood estimation (MLE) approach to voting was proposed by Condorcet [12]. Young [30] formalized Condorcet’s ideas, and showed that Condorcet’s approach to choosing the best ranking results in a voting rule that is known as Kemeny’s rule. Young has also considered the problem of selecting the most likely winner, focusing on the limit cases where the noise is extremely high or extremely low. However, his analysis of this setting is presented by means of an example and appears to be flawed. Procaccia et al. [23] formalized Young’s analysis and extended it to objectives other than picking the best alternative. Recently, Caragiannis et al. [6] have further generalized this approach by focusing on the design of voting rules that demonstrate robustness to noise originating from a wide family of noise models. Such robustness is also a feature of our agnostic rule, in that it returns a set of alternatives that is guaranteed to contain the most likely alternative irrespective of the value of the underlying parameter of Mallows’ model from which votes are generated.

Other variants of the maximum likelihood approach have been considered in the computational social choice litera-

¹Simultaneously and independently, Azari Soufiani et al. [2] also introduced the tournament variant of Mallows’ model, but for a different goal of studying the optimal Bayesian estimators.

ture [10, 9, 28]. Perhaps the closest to our work is that of Xia et al. [29], who studied the MLE approach in multi-issue domains, where alternatives represent combinations of multiple issues. Xia et al. used CP-nets to represent the (possibly cyclic) preferences of the voters. However, they focused on dealing with the huge space of alternatives created by an exponential number of combinations.

Finally, we show that Tideman’s rule provides a very simple and elegant deterministic 2-approximation for the Kemeny winner. Note that the Kemeny winner and the Kemeny ranking are \mathcal{NP} -hard to compute [23, 3]. Approximations of the Kemeny ranking have been studied extensively in the literature, varying from deterministic approximations [11, 27] through randomized approximations [1] to a polynomial time approximation scheme (PTAS) [14]; further, any approximation of the Kemeny ranking implies an identical approximation of the Kemeny winner (see the discussion following Theorem 7).

3 PRELIMINARIES

Let $[k] = \{1, \dots, k\}$. We consider a set of *alternatives* A with $|A| = m$. Let $\mathcal{L}(A)$ denote the set of *votes*, where a vote is a ranking (linear order) over the alternatives, denoted $\sigma : \{1, \dots, m\} \rightarrow A$. Thus, alternative $\sigma(i)$ is the i -th most preferred alternative in σ ; $\sigma(1)$ and $\sigma(m)$ are, respectively, the most and the least preferred alternatives in σ . Note that $|\mathcal{L}(A)| = m!$. A *profile* $\pi \in \mathcal{L}(A)^n$ is a collection of n votes. For alternatives $a, b \in A$, let n_{ab} denote the number of votes in π that rank a above b . Hence, $n_{ab} + n_{ba} = n$ for all $a, b \in A$. Let $\Delta_{ab} = n_{ab} - n_{ba}$; this quantity can be thought of as the advantage of a over b .

Voting rules. A *voting rule* is a mapping from profiles to (a set of tied) *winning alternatives*. Formally, a voting rule is a mapping $f : \mathcal{L}(A)^n \rightarrow \mathcal{P}(A)$, where $\mathcal{P}(\cdot)$ denotes the power set.² We review three prominent voting rules that play a crucial role in this paper.

- *The Borda count.* Under the Borda count, each voter awards $m - i$ points to the alternative she ranks in position i , i.e., each alternative receives a number of points equal to the number of alternatives it defeats. The scores of the alternatives are tallied across the votes. That is, the Borda score of an alternative $a \in A$ in profile π is

$$\text{SC}^{BD}(a) = \sum_{\sigma \in \pi} \sum_{b \in A \setminus \{a\}} \mathbb{I}[a \succ_{\sigma} b] = \sum_{b \in A \setminus \{a\}} n_{ab},$$

where the second equality follows by switching the order of summation. The winner(s) are the alternative(s) with the highest score.

²Technically, such mappings are known as *social choice functions*. In contrast, *social welfare functions* map every profile to a ranking or a set of tied rankings over the alternatives.

- *Tideman’s rule.* More commonly known as *Tideman’s simplified Dodgson rule*,³ this rule was put forward by Tideman [25, pp. 199-201] as a polynomial-time computable approximation to Dodgson’s rule [13], which is \mathcal{NP} -hard to compute. Under Tideman’s rule, the score of an alternative a is given by

$$\text{SC}^{TD}(a) = \sum_{b \in A \setminus \{a\}} \max(0, \Delta_{ba}).$$

That is, the Tideman score of a is the cumulative advantage of all alternatives with a positive advantage over a . The winners are the alternatives with the *minimum* score.

- *Kemeny’s rule.* The Kendall tau distance between two rankings is the number of pairs of alternatives on which they disagree, i.e., for $\sigma_1, \sigma_2 \in \mathcal{L}(A)$, $d(\sigma_1, \sigma_2) = |\{(a, b) \mid a \succ_{\sigma_1} b, b \succ_{\sigma_2} a\}|$. With slight abuse of notation, for a profile π and a ranking σ , let $d(\pi, \sigma) = \sum_{\sigma' \in \pi} d(\sigma, \sigma')$. Under Kemeny’s rule, the score of an alternative $a \in A$ is the minimum distance from the input profile to any ranking that puts a first. Formally,

$$\text{SC}^{KM}(a) = \min_{\sigma \in \mathcal{L}(A): \sigma(1)=a} d(\pi, \sigma).$$

The winners are the alternatives with the *minimum* score. Equivalently, the rankings with the smallest distance from the profile are selected, and the winners are the alternatives appearing first in these rankings.

Refinement of voting rules. We say that voting rule \hat{f} is a *refinement* of voting rule f if $\hat{f}(\pi) \subseteq f(\pi)$ for every profile π . That is, \hat{f} can be seen as a combination of f with a (partial) tie-breaking rule.

4 MODEL

We begin by presenting the well-known ranking version of Mallows’ model, and then we introduce a more general tournament version of this model.

4.1 THE RANKING MODEL

Assume there is a hidden true ranking $\sigma^* \in \mathcal{L}(A)$ over the alternatives, which reflects the order of their true strengths. We also make the standard assumption that σ^* is selected using a uniform prior over $\mathcal{L}(A)$. Thus, $\sigma^*(1)$ denotes the true best alternative. A *noise model* describes how votes are generated given the true ranking. Votes in a profile are then assumed to be iid samples from the noise model.

Specifically, in Mallows’ model [19] (also known as the *Condorcet noise model* [12]), which was described informally in the introduction, the probability of generating a

³Tideman’s rule considered in this paper should not be confused with *the ranked pairs method*, also proposed by Tideman.

vote σ when the true ranking is σ^* is given by

$$\Pr[\sigma \mid \sigma^*] = \frac{\varphi^{d(\sigma, \sigma^*)}}{Z_\varphi^m}. \quad (1)$$

Here, $\varphi = \frac{1-p}{p} \in (0, 1)$ is the noise parameter of the model and $p \in (1/2, 1)$ is the probability of making the correct decision when comparing two alternatives. Thus, $\varphi \rightarrow 0$ represents a distribution concentrated around σ^* , whereas $\varphi \rightarrow 1$ converges to the uniform distribution, which has the greatest noise. Finally, d is the Kendall tau distance, and $Z_\varphi^m = \sum_{\sigma \in \mathcal{L}(A)} \Pr[\sigma \mid \sigma^*]$ is the normalization constant, which turns out to be independent of the true ranking σ^* (see, e.g., [17]).

Now, take a profile $\pi \in \mathcal{L}(A)^n$. Since individual votes are sampled iid, the probability of generating π is

$$\Pr[\pi \mid \sigma^*] = \prod_{\sigma \in \pi} \frac{\varphi^{d(\sigma, \sigma^*)}}{Z_\varphi^m} \propto \varphi^{d(\pi, \sigma^*)}.$$

Under the assumption of uniform prior over the true ranking σ^* , and for given φ , the probability of an alternative $a \in A$ being the true best alternative is proportional to

$$\sum_{\substack{\sigma^* \in \mathcal{L}(A): \\ \sigma^*(1)=a}} \Pr[\pi \mid \sigma^*] \propto \sum_{\substack{\sigma^* \in \mathcal{L}(A): \\ \sigma^*(1)=a}} \varphi^{d(\pi, \sigma^*)}. \quad (2)$$

Let $\Gamma_\varphi^R(a)$ be the “likelihood polynomial” of a , as given in the final expression of Equation (2). Then, the *maximum likelihood estimator* of the true best alternative is the set of alternatives having the highest probability of being the best alternative, given by $\text{MLE}_\varphi^R(\pi) = \arg \max_{a \in A} \Gamma_\varphi^R(a)$. Theorem 3.2 by Procaccia et al. [23] shows the following.

Theorem 1 (Procaccia et al. [23]). *Computing MLE_φ^R is \mathcal{NP} -hard.*

4.2 THE TOURNAMENT MODEL

We now introduce a variant of Mallows’ model where both the ground truth and the samples need not be total orders. Rather, they can be *tournaments*, i.e., sets of pairwise comparisons (one for each pair of alternatives). A tournament need not be transitive: it can be the case that a beats b , b beats c , and c beats a . As argued in the introduction, this is common when the alternatives are compared based on multiple quality parameters instead of a single parameter, and/or users are not required to submit total orders. Note that every ranking can be seen as a tournament.

Let $\mathcal{T}(A)$ denote the set of all tournaments over alternatives in A . We still use $a \succ_T b$ to denote that alternative a is preferred to alternative b in the tournament T . The Kendall tau distance extends to $\mathcal{T}(A)$ in a natural way: given two tournaments $T, T' \in \mathcal{T}(A)$, $d(T, T')$ is the number of pairs of alternatives on which T and T' disagree. Further, the

quantities $(n_{ab})_{a,b \in A}$ remain well-defined for a profile of tournaments $\pi \in \mathcal{T}(A)^n$. As the three voting rules introduced in Section 3 (the Borda count, Tideman’s rule, and Kemeny’s rule) can be defined in terms of $(n_{ab})_{a,b \in A}$, these rules are well-defined over profiles of tournaments as well.

Let $T^* \in \mathcal{T}(A)$ denote the hidden true tournament over the alternatives. We assume that T^* is selected using a uniform prior over $\mathcal{T}(A)$. That is, for each pair of alternatives $a, b \in A$ we independently decide whether $a \succ_{T^*} b$ or $b \succ_{T^*} a$, with both possibilities being equally likely. When generating a vote, each pairwise comparison in T^* is retained with a fixed probability $1/2 < p < 1$ and flipped with probability $1 - p$. Unlike in the ranking model, the pairwise comparisons in the samples are independent of each other. Accordingly, the probability of generating a tournament T when the true tournament is T^* is given by

$$\Pr[T | T^*] = p^{\binom{m}{2} - d(T, T^*)} (1 - p)^{d(T, T^*)} = p^{\binom{m}{2}} \varphi^{d(T, T^*)}.$$

We consider profiles consisting on n tournaments, which are sampled iid from the noise model. Let $d(\pi, T^*) = \sum_{T \in \pi} d(T, T^*)$. Then, the probability of generating a profile $\pi \in \mathcal{T}(A)^n$ is proportional to $\varphi^{d(\pi, T^*)}$, similarly to the ranking-based model.

Procaccia et al. [23] introduced the noisy choice model as the generalization of Mallows’ model where the ground truth was a ranking but the samples could be tournaments. In that sense, our model is a further generalization where even the ground truth may be a tournament.

However, this causes a potentially serious problem: The best alternative in a ranking σ^* is $\sigma^*(1)$. But the definition of the best alternative in a tournament T^* is unclear. Following Condorcet’s own definition of “Condorcet winners” for cyclic majority preferences, we say that an alternative is the winner in a tournament if it is preferred to every other alternative. Note that not every tournament has a winner. For a tournament T , define $\text{win}(T)$ to be the winner of T if it exists, and \emptyset otherwise.

Given a profile $\pi \in \mathcal{T}(A)^n$, we can now compute the likelihood of an alternative $a \in A$ being the best alternative in the unknown true tournament. Indeed, for every $T^* \in \mathcal{T}(A)$ with $\text{win}(T^*) = a$ we have

$$\begin{aligned} d(\pi, T^*) &= \sum_{b \in A \setminus \{a\}} n_{ba} \\ &+ \sum_{c, d \in A \setminus \{a\}} (n_{cd} \cdot \mathbb{I}[d \succ_{T^*} c] + n_{dc} \cdot \mathbb{I}[c \succ_{T^*} d]). \end{aligned}$$

Further, for each possible combination of pairwise comparisons of the alternatives in $A \setminus \{a\}$, the set $\{T^* \in \mathcal{T}(A) \mid \text{win}(T^*) = a\}$ contains exactly one tournament that realizes this combination. Hence, we have

$$\begin{aligned} \sum_{\substack{T^* \in \mathcal{T}(A): \\ \text{win}(T^*) = a}} \varphi^{d(\pi, T^*)} &= \varphi^{\sum_{b \in A \setminus \{a\}} n_{ba}} \cdot \prod_{c, d \in A \setminus \{a\}} (\varphi^{n_{cd}} + \varphi^{n_{dc}}) \\ &\propto \prod_{b \in A \setminus \{a\}} \frac{\varphi^{n_{ba}}}{\varphi^{n_{ba}} + \varphi^{n_{ab}}} = \prod_{b \in A \setminus \{a\}} \frac{1}{1 + \varphi^{n_{ab} - n_{ba}}}. \end{aligned}$$

Now, for an alternative $a \in A$, define its likelihood polynomial $\Gamma_\varphi^T(a) = \prod_{b \in A \setminus \{a\}} (1 + \varphi^{n_{ab} - n_{ba}})$. Technically, $\Gamma_\varphi^T(a)$ is a Laurent polynomial, i.e., some of the powers of φ may be negative. Therefore, we will sometimes work with the function $\hat{\Gamma}_\varphi^T(a) = \varphi^{nm} \Gamma_\varphi^T(a)$, which is a polynomial of degree at most $2nm$. Note that the likelihood polynomial of a is proportional to the inverse of the likelihood, and is therefore to be minimized. Thus, the maximum likelihood estimator for the best alternative is given by $\text{MLE}_\varphi^T(\pi) = \arg \min_{a \in A} \Gamma_\varphi^T(a)$, or, equivalently, $\text{MLE}_\varphi^T(\pi) = \arg \min_{a \in A} \hat{\Gamma}_\varphi^T(a)$. Since $\Gamma_\varphi^T(a)$ can be computed for every alternative $a \in A$ and every $\varphi \in \mathbb{Q}$ in polynomial time, the following is trivial.

Theorem 2. *Computing MLE_φ^T is in \mathcal{P} .*

5 LIMITING VOTING RULES

We will now study the extreme cases with very low and very high noise in input votes, i.e., $\varphi \rightarrow 0$ and $\varphi \rightarrow 1$, respectively. First, we observe that both for rankings and for tournaments and in both limiting cases, the limiting rule is well-defined, i.e., there exist α and β with $0 < \alpha < \beta < 1$ such that for $P \in \{R, T\}$ $\text{MLE}_\varphi^P = \text{MLE}_\alpha^P$ for all $0 < \varphi \leq \alpha$ and $\text{MLE}_\varphi^P = \text{MLE}_\beta^P$ for all $\beta \leq \varphi < 1$. Indeed, fix a profile π . For each $a \in A$ the degree of the likelihood polynomials $\Gamma_\varphi^R(a)$ and $\hat{\Gamma}_\varphi^T(a)$ is finite, and therefore we can pick $\alpha^\pi, \beta^\pi \in (0, 1)$ so that no two of these polynomials for π intersect in $(0, \alpha^\pi)$ or in $(\beta^\pi, 1)$. Since the number of profiles with a *fixed number of votes* is finite, taking the minimum of α^π and the maximum of β^π over all such profiles gives the desired values of α and β . Note, however, that this argument breaks down if the number of votes may vary. For example, if $\Gamma_\varphi^R(a)$ and $\Gamma_\varphi^R(b)$ for a profile π intersect at φ , then $\Gamma_\varphi^R(a)$ and $\Gamma_\varphi^R(b)$ for the profile $k\pi$ intersect at $\sqrt[k]{\varphi}$ (where $k\pi$ is the profile where each entry of π is repeated k times). As k is unbounded, we obtain $\beta = \sup_\pi \beta^\pi = 1$.

Notation. For the ranking model, let $\text{MLE}_{\text{Acc}}^R$ and $\text{MLE}_{\text{Inacc}}^R$ denote the limiting rules in the accurate case ($\varphi \rightarrow 0$) and in the inaccurate case ($\varphi \rightarrow 1$), respectively. Similarly, for the tournament model, let $\text{MLE}_{\text{Acc}}^T$ and $\text{MLE}_{\text{Inacc}}^T$ denote the limiting rules in the accurate case and in the inaccurate case, respectively.

5.1 THE RANKING MODEL

The accurate case ($\varphi \rightarrow 0$): Procaccia et al. [23] showed that when $\varphi \rightarrow 0$, every MLE best alternative is first in some Kemeny ranking. Further, they also showed that finding even a single Kemeny winner is \mathcal{NP} -hard.⁴ These results can be restated as follows.

⁴Both results can be found in the proof of Theorem A.1 in the appendix of the full version available at <http://www.cs.cmu.edu/~arielpro/papers/mle.full.pdf>.

Theorem 3 (Procaccia et al. [23]). $\text{MLE}_{\text{Acc}}^R$ is a refinement of Kemeny’s rule, and is \mathcal{NP} -hard to compute.

In fact, our tools enable us to describe $\text{MLE}_{\text{Acc}}^R$ in more detail; see Appendix A of the full version of the paper.⁵

The inaccurate case ($\varphi \rightarrow 1$): Procaccia et al. [23, Theorem 4.1] proved that every MLE best alternative in this case is also a Borda winner. However, they left open the question of computational complexity. Despite significant effort, we were unable to settle the computational complexity either. We conjecture that $\text{MLE}_{\text{Inacc}}^R$ is \mathcal{NP} -hard to compute.

Theorem 4 (Procaccia et al. [23]). $\text{MLE}_{\text{Inacc}}^R$ is a refinement of the Borda count.

Once again, the exact refinement is given in Appendix A of the full version. While the computational complexity of $\text{MLE}_{\text{Inacc}}^R$ is unknown, we remark that computing its output is easy whenever Borda’s rule produces a unique winner, which is often the case.

See Section 5.3 for an example showing the computation of $\text{MLE}_{\text{Acc}}^R$ and $\text{MLE}_{\text{Inacc}}^R$ for a given profile.

5.2 THE TOURNAMENT MODEL

The accurate case ($\varphi \rightarrow 0$): In this case we show the following.

Theorem 5. $\text{MLE}_{\text{Acc}}^T$ is a refinement of Tideman’s rule, and can be computed in polynomial time.

Proof. To determine the winner(s) under $\text{MLE}_{\text{Acc}}^T$ we need to compare the likelihood polynomials $\Gamma_\varphi^T(a) = \prod_{b \in A \setminus \{a\}} (1 + \varphi^{n_{ab} - n_{ba}})$ when $\varphi \rightarrow 0$. Pick $\alpha \in (0, 1)$ so that no two likelihood polynomials intersect in $(0, \alpha)$. As $\varphi \rightarrow 0$, the dominating term in $\Gamma_\varphi^T(a)$ is the term with the smallest power of φ . Denote the smallest power by $t_\varphi(a)$.

$$t_\varphi(a) = \sum_{\substack{b \in A \setminus \{a\}, \\ n_{ab} \leq n_{ba}}} n_{ab} - n_{ba} = - \sum_{b \in A \setminus \{a\}} \max\{0, \Delta_{ba}\}$$

(where we take the sum over an empty set to be 0). Hence, for $\varphi \in (0, \alpha)$ we have $\Gamma_\varphi^T(a) < \Gamma_\varphi^T(b)$ whenever $t_\varphi(a) > t_\varphi(b)$, or, equivalently, whenever $\text{SC}^{TD}(a) < \text{SC}^{TD}(b)$. Recall that we are interested in alternatives with the smallest value of the likelihood polynomial on $(0, \alpha)$; our calculation shows that every such alternative is a Tideman winner.

To show that $\text{MLE}_{\text{Acc}}^T$ is polynomial-time computable, it is not sufficient to observe that the functions $\Gamma_\varphi^T(a)$, $a \in A$, can be evaluated in polynomial time, as we also need to

find a small enough value of φ at which they should be compared. Nevertheless, comparing likelihood polynomials at $\varphi \rightarrow 0$ is not difficult. We first multiply the terms of $\hat{\Gamma}_\varphi^T(a)$ one-by-one, followed by expansion at each stage, to obtain the coefficients of this polynomial. Note that the degree of $\hat{\Gamma}_\varphi^T(a)$ is at most $2mn$, so this step can be implemented efficiently. To compare two polynomials at $\varphi \rightarrow 0$, it suffices to consider their coefficients lexicographically, starting with the lowest-order terms. The details are given in Appendix A of the full version of the paper. \square

The inaccurate case ($\varphi \rightarrow 1$): This case has striking similarity with the inaccurate case of the ranking model.

Theorem 6. $\text{MLE}_{\text{Inacc}}^T$ is a refinement of the Borda count, and can be computed in polynomial time.

Proof. Note that $\hat{\Gamma}_1^T(a) = \Gamma_1^T(a) = 1$ for all $a \in A$. Therefore, to compare the likelihood polynomials as $\varphi \rightarrow 1$, we will first compare their derivatives at $\varphi = 1$. We have

$$\begin{aligned} \left. \frac{d}{d\varphi} \Gamma_\varphi^T(a) \right|_{\varphi=1} &= 2^{m-2} \sum_{b \in A \setminus \{a\}} \left. \frac{d}{d\varphi} (1 + \varphi^{n_{ab} - n_{ba}}) \right|_{\varphi=1} \\ &= 2^{m-2} \sum_{b \in A \setminus \{a\}} (n_{ab} - n_{ba}). \end{aligned}$$

As φ approaches 1 from the left, we have $\Gamma_\varphi^T(a) < \Gamma_\varphi^T(b)$ whenever $\left. \frac{d}{d\varphi} \Gamma_\varphi^T(a) \right|_{\varphi=1} > \left. \frac{d}{d\varphi} \Gamma_\varphi^T(b) \right|_{\varphi=1}$. Using $n_{ba} = n - n_{ab}$, we observe that the latter condition is equivalent to $\text{SC}^{BD}(a) > \text{SC}^{BD}(b)$. Thus, the winners under $\text{MLE}_{\text{Inacc}}^T$ must have the highest Borda score, i.e., $\text{MLE}_{\text{Inacc}}^T$ is a refinement of the Borda rule.

In contrast with the ranking-based model, the rule $\text{MLE}_{\text{Inacc}}^T$ can be computed in polynomial time. Similarly to the accurate case of the tournament model (see the proof of Theorem 5), we multiply the terms of each $\hat{\Gamma}_\varphi^T(a)$, $a \in A$, in order to obtain the coefficients of these polynomials. Then, for each polynomial we compute its first $2nm$ derivatives at $\varphi = 1$. As the degree of each of these polynomials does not exceed $2nm$, comparing two such polynomials at $\varphi \rightarrow 1$ amounts to lexicographically comparing these two lists of values. \square

Alternatively, we can show that $\text{MLE}_{\text{Acc}}^T$ and $\text{MLE}_{\text{Inacc}}^T$ are polynomial-time computable by using results on *root separation* of polynomials. A classic paper by Mahler [18] proved the following.

Fact: Any two distinct roots of a polynomial are separated by at least H^{-k+1} , where H is the maximum absolute value of any coefficient, and k is the degree.

See [4] for further explanation and improved results for polynomials with integer coefficients. In our case, we can consider the polynomial $P = \prod_{a, b \in A} (\hat{\Gamma}_\varphi^T(a) - \hat{\Gamma}_\varphi^T(b))$;

⁵The full version can be found at <http://www.cs.cmu.edu/~nkshah/papers.html>

its degree does not exceed $2nm^3$ and its coefficients are at most exponential in $\text{poly}(n, m)$. Thus, $\Delta = H^{-k+1}$ has *polynomially many bits*. An exponential upper bound on H (and hence the respective lower bound on Δ) can be computed without expanding P . Since no two likelihood polynomials intersect on $(0, \Delta)$ or on $(1 - \Delta, 1)$, the outputs of $\text{MLE}_{\text{Acc}}^T$ and $\text{MLE}_{\text{Inacc}}^T$ can be computed by evaluating and comparing $\Gamma_\varphi^T(a)$, $a \in A$, in their product form at $\varphi = \Delta/2$ and $\varphi = 1 - \Delta/2$, respectively. We prefer the approach presented in the proofs above because it seems to work faster in practice, possibly due to the fact that it relies only on integer arithmetic.

5.3 THE TOURNAMENT MODEL VERSUS THE RANKING MODEL

The goal of this section is to compare the ranking-based model and the tournament-based model. We begin by presenting a profile on which the limiting voting rules for the two models differ.

Example 1. Consider a profile π consisting of the following 3 rankings over 4 alternatives.

$$a \succ b \succ c \succ d, \quad d \succ a \succ b \succ c, \quad c \succ d \succ b \succ a.$$

Recall that the likelihood polynomials $\Gamma_\varphi^R(a)$ (in the ranking model) and $\Gamma_\varphi^T(a)$ (in the tournament model) of an alternative $a \in A$ are given by

$$\Gamma_\varphi^R(a) = \sum_{\substack{\sigma^* \in \mathcal{L}(A): \\ \sigma^*(1)=a}} \varphi^{d(\pi, \sigma^*)},$$

$$\Gamma_\varphi^T(a) = \prod_{b \in A \setminus \{a\}} (1 + \varphi^{n_{ab} - n_{ba}}).$$

For profile π , the likelihood polynomials are given below.

$$\begin{aligned} \Gamma_\varphi^R(a) &= \varphi^8 + \varphi^8 + \varphi^8 + \varphi^9 + \varphi^9 + \varphi^9, \\ \Gamma_\varphi^R(b) &= \varphi^9 + \varphi^9 + \varphi^9 + \varphi^{10} + \varphi^{10} + \varphi^{10}, \\ \Gamma_\varphi^R(c) &= \varphi^8 + \varphi^9 + \varphi^9 + \varphi^{10} + \varphi^{10} + \varphi^{11}, \\ \Gamma_\varphi^R(d) &= \varphi^7 + \varphi^8 + \varphi^8 + \varphi^9 + \varphi^9 + \varphi^{10}, \\ \Gamma_\varphi^T(a) &= (1 + \varphi^1) (1 + \varphi^1) (1 + \varphi^{-1}), \\ \Gamma_\varphi^T(b) &= (1 + \varphi^{-1}) (1 + \varphi^1) (1 + \varphi^{-1}), \\ \Gamma_\varphi^T(c) &= (1 + \varphi^{-1}) (1 + \varphi^{-1}) (1 + \varphi^1), \\ \Gamma_\varphi^T(d) &= (1 + \varphi^1) (1 + \varphi^1) (1 + \varphi^{-1}). \end{aligned}$$

Now, we can compute the limiting voting rules using the likelihood polynomials as explained in Sections 5.1 and 5.2. The results of these rules along with those of the Borda count, Kemeny’s rule, and Tideman’s rule are given in Table 1.

While Example 1 shows that the accurate and the inaccurate cases of the tournament model differ from the respective cases of the ranking model, we show that the tournament model serves as a satisfactory polynomial-time approximation of the ranking model, where computing the

Ranking, Accurate	$\text{MLE}_{\text{Acc}}^R(\pi) = \{d\}$
Ranking, Inaccurate	$\text{MLE}_{\text{Inacc}}^R(\pi) = \{d\}$
Tournament, Accurate	$\text{MLE}_{\text{Acc}}^T(\pi) = \{a, d\}$
Tournament, Inaccurate	$\text{MLE}_{\text{Inacc}}^T(\pi) = \{a, d\}$
Borda count	$\text{Borda}(\pi) = \{a, d\}$
Kemeny’s rule	$\text{Kemeny}(\pi) = \{d\}$
Tideman’s rule	$\text{Tideman}(\pi) = \{a, d\}$

Table 1: Various voting rules applied on π .

limiting rules is non-trivial (and provably \mathcal{NP} -hard in the accurate case). While the tournament model—where both the ground truth and the estimates may be cyclic—has its intrinsic motivation (see Section 1), this offers an additional strong motivation for the model. The similarity of both models in the inaccurate case is evident: Both $\text{MLE}_{\text{Inacc}}^R$ and $\text{MLE}_{\text{Inacc}}^T$ are refinements of Borda’s rule. However, as seen in Example 1, these two rules are not identical. Moreover, we can show that neither of these rules is a refinement of the other: Appendix C in the full version presents an example with 8 rankings over 4 alternatives where both $\text{MLE}_{\text{Inacc}}^R$ and $\text{MLE}_{\text{Inacc}}^T$ have unique winners that are different. We remark, however, that these two rules return the same output most of the time (see Section 7), and always when the Borda winner is unique.

Motivated by the similarity in the inaccurate case, we compared Tideman’s rule and Kemeny’s rule, because the limiting rules in the accurate case of the ranking and the tournament models are refinements of Kemeny’s rule and Tideman’s rule, respectively. While the two rules were not previously thought to be connected, we show that Tideman’s rule is a 2-approximation of Kemeny’s rule for winner determination.

Theorem 7. *The Kemeny score of a Tideman winner is at most twice the Kemeny score of a Kemeny winner.*

Proof. First, we describe an alternative interpretation of Kemeny’s rule proposed by Conitzer et al. [8]. Given a profile π over A , define a *weighted pairwise majority (WPM) graph for a set of alternatives* $A' \subseteq A$ to be the directed graph $G_{A'}$ where the vertices are the alternatives in A' , and there is an edge between every pair of alternatives $a, b \in A'$ with weight $|\Delta_{ab}| = |n_{ab} - n_{ba}|$. The edge goes from a to b if $n_{ab} > n_{ba}$ and from b to a if $n_{ba} > n_{ab}$. When $n_{ab} = n_{ba}$, the edge with zero weight may be drawn in either direction.

The *feedback* of a ranking with respect to a WPM graph $G_{A'}$ is defined as the sum of the weights of edges of $G_{A'}$ going in direction opposite to the ranking. Conitzer et al. [8] showed that Kemeny’s rule is equivalent to first finding the rankings with the smallest feedback with respect to G_A , and then returning their top alternatives.

Equivalently, we can say that in G_A for every pair of alter-

natives $a, b \in A$ there is an edge from a to b with weight $\max(0, \Delta_{ab})$. Thus, the Tideman score of an alternative $a \in A$ is the sum of weights of its incoming edges, and the Tideman winners are the vertices that minimize this sum. For a subset of alternatives $S \subseteq A$, let $F(S)$ be the smallest feedback with respect to G_S , over all rankings of A . To compute the Kemeny score of an alternative $a \in A$, we consider the set of all rankings \mathcal{L}_a that put a first, and find a ranking in \mathcal{L}_a that has the minimum feedback with respect to G_A . Note that the feedback of any ranking in \mathcal{L}_a contains all incoming edges of a . Thus, to minimize feedback over \mathcal{L}_a , we order the alternatives in $A \setminus \{a\}$ so as to minimize the feedback over $G_{A \setminus \{a\}}$. Hence,

$$\text{SC}^{KM}(a) = \text{SC}^{TD}(a) + F(A \setminus \{a\}). \quad (3)$$

Further, for all $a, b \in A$ with $a \neq b$, we have

$$F(A \setminus \{a\}) \leq \text{SC}^{TD}(b) + F(A \setminus \{a, b\}). \quad (4)$$

Indeed, the left-hand side of (4) is the feedback of the best ranking with respect to $G_{A \setminus \{a\}}$, whereas the right-hand side of (4) is the feedback of the best ranking with respect to $G_{A \setminus \{a\}}$ among those that put b first, plus $\max(0, \Delta_{ab})$.

Now, consider a profile π . Let $a \in A$ be a Tideman winner and let $b \in A$ be a Kemeny winner. We want to show that $\text{SC}^{KM}(a) \leq 2\text{SC}^{KM}(b)$. If $a = b$, this is trivial. Thus, assume $a \neq b$. Combining (3) and (4), we obtain

$$\begin{aligned} \text{SC}^{KM}(a) &= \text{SC}^{TD}(a) + F(A \setminus \{a\}) \\ &\leq \text{SC}^{TD}(a) + \text{SC}^{TD}(b) + F(A \setminus \{a, b\}) \\ &\leq \text{SC}^{TD}(b) + \text{SC}^{TD}(b) + F(A \setminus \{b\}) \\ &\leq 2(\text{SC}^{TD}(b) + F(A \setminus \{b\})) = 2 \cdot \text{SC}^{KM}(b). \end{aligned}$$

Hence, the Kemeny score of any Tideman winner is a 2-approximation of the optimal Kemeny score. \square

Appendix D in the full version gives an example using 4 alternatives where the approximation factor is exactly 2. Hence, the result of Theorem 7 is tight. Caragiannis et al. [5] show that Tideman’s rule, which was originally proposed as an approximation to Dodgson’s rule, is actually an asymptotically optimal approximation of Dodgson’s rule. Theorem 7 shows that it is also a 2-approximation of Kemeny’s rule. Dodgson’s rule and Kemeny’s rule are deeply connected [24]: While Dodgson’s rule makes the smallest number of pairwise swaps to reach a profile with a Condorcet winner (an alternative preferred by a majority of voters to every other alternative), Kemeny’s rule makes the swaps until the majority opinion becomes acyclic and then returns the first alternative in the acyclic order. Tideman’s rule can now be seen as a hybrid that provides a good approximation to both Dodgson’s rule and Kemeny’s rule.

While approximations of the Kemeny ranking are studied extensively in the literature, we are not aware of any work

explicitly studying approximations of the Kemeny winner. However, it is easy to check that the first alternative of a ranking that is a c -approximation of the Kemeny ranking is also a c -approximation of the Kemeny winner. Hence, we can directly compare the result of Theorem 7 against approximations of the Kemeny ranking in the literature. The Kemeny ranking admits a polynomial time approximation scheme (PTAS) [14], which is, however, rather impractical. Constant approximations are therefore studied because they are fast and simple [11, 1, 27]. Tideman’s rule, which admits an elegant closed-form expression (see Section 3), is the simplest deterministic 2-approximation of the Kemeny winner that we are aware of. We believe that this result may be of independent interest to the social choice community.

6 THE AGNOSTIC RULE

While the limiting cases of $\varphi \rightarrow 0$ and $\varphi \rightarrow 1$ may be appropriate in some scenarios (and their analysis yields interesting connections, e.g., Theorem 7), in most practical settings the level of noise is unknown. We could include φ as one of the unknown parameters and infer the best possible values for the true ranking/tournament and φ (see, e.g., [17]). However, this approach gives one specific value (a “point estimate”) of φ . If the point estimate is wrong, the estimate for the best alternative is also sub-optimal.

Consider instead an agnostic approach that refrains from estimating the value of φ . Rather, given a profile, it returns the set of *all* alternatives that are the most likely winners for *some value of* φ . Let MLE_{Ag}^R and MLE_{Ag}^T denote the agnostic rules in the ranking model and in the tournament model, respectively. Then for a profile π ,

$$\begin{aligned} \text{MLE}_{\text{Ag}}^R(\pi) &= \bigcup_{\varphi \in (0,1)} \text{MLE}_{\varphi}^R(\pi), \\ \text{MLE}_{\text{Ag}}^T(\pi) &= \bigcup_{\varphi \in (0,1)} \text{MLE}_{\varphi}^T(\pi). \end{aligned}$$

There are three advantages of this approach over the inference approach: First, as we show below, MLE_{Ag}^T can be computed in polynomial time, and results presented in Section 7 demonstrate that MLE_{Ag}^T is a good approximation of MLE_{Ag}^R as well. Thus, it is easy to compute and use the agnostic rule. Most inference problems, on the other hand, are hard to solve [17]. Second, if the data is indeed generated from a Mallows’ (ranking or tournament) model, the MLE best alternative is guaranteed to be in the set returned by the agnostic rule, which is not the case for the inference approach. Third, while the set of winners under the agnostic rules contains the set of winners for any specific value of φ , our simulations in Section 7 show that on average only a few winning alternatives are returned.

Now, we show that the agnostic rule can be computed in polynomial time for the tournament model. Note that this is not obvious: While the limiting cases can be analyzed

by looking at the coefficients of the likelihood polynomials or using a root separation approach, these methods do not work for values of φ away from 0 and 1. We use another result regarding polynomials, which is known as root isolation (see, e.g., [7, 26]).

Fact: Given a polynomial, one can compute, in time polynomial in the input size, a set of disjoint intervals that isolate the roots of the polynomial, i.e., disjoint intervals that collectively contain all roots of the polynomial but each interval only contains a single distinct root.

Theorem 8. Computing MLE_{Ag}^T is in \mathcal{P} .

Proof. Once again, consider the polynomial $P = \prod_{a,b \in A} (\hat{\Gamma}_{\varphi}^T(a) - \hat{\Gamma}_{\varphi}^T(b))$. We have argued that the degree of P is at most $2nm^3$, and its coefficients can be computed in polynomial time. Next, we use root isolation to isolate the roots of P in polynomial time. Note that any value of $\varphi \in (0, 1)$ where some alternatives a, b , $a \neq b$, have equal likelihood is a root of P . Hence, in any region between two consecutive roots of P , the order of likelihoods of different alternatives is fixed.

Therefore, computing MLE_{Ag}^T amounts to taking one value of φ between each consecutive pair of isolating intervals, evaluating MLE_{φ}^T at every such φ , and returning the set of all MLE alternatives found. Note that the number of roots of P and therefore the number of evaluations of MLE_{φ}^T is polynomial in the input size, and we have already established that evaluating MLE_{φ}^T itself can be done in polynomial time (Theorem 2). Hence, the overall running time is polynomial in the input size. \square

7 EXPERIMENTS

In this section, we complement our theoretical results by two sets of experiments. The results of Section 5.3 establish that the tournament model can be thought of as a polynomial-time approximation to the ranking model. The first set of experiments analyzes how close the limiting rules (and the agnostic rules) in the two models are to each other on average. The second set of experiments aims to check whether the agnostic rules return reasonably small sets of winning alternatives. To this end, we compute the average number of winning alternatives returned by the agnostic rules in the two models.

For both sets of experiments, we generate profiles using iid samples from Mallows' (ranking) model with the noise parameter φ taking 10 different values from 0.1 to 1.⁶ In each case, we average our results over 5000 sampled profiles. It is easy to check that the Borda winner, the Tideman winner, and the Kemeny winner always coincide in case of 3

⁶We use the ranking model to generate profiles so that the rules for both models can be applied on the generated profiles which contain rankings, which are also tournaments.

alternatives. Hence, in our experiments we set the number of alternatives m to 5 or 7; the number of votes n is also either 5 or 7. Thus, each of the graphs presented has four lines; one for each of $(n, m) = (5, 5), (5, 7), (7, 5),$ and $(7, 7)$. In all the graphs, the x -axis shows the noise parameter φ used to generate profiles. Importantly, while the tournament model admits polynomial-time algorithms, and can therefore be used with a large number of alternatives, we use a small number of alternatives to be able to compare the associated voting rules with the exponential-time rules of the ranking model.

For the first set of experiments, we measure the dissimilarity between three pairs of rules in terms of the dissimilarity between the sets of winning alternatives they return. As the measure of dissimilarity between two sets A and B , we use the *Jaccard distance*, which is defined as follows.

$$d_J(A, B) = \frac{|A \cup B| - |A \cap B|}{|A \cup B|}.$$

Figures 1(a), 1(b), and 1(c) respectively show the dissimilarity between the limiting rules for the accurate case, the limiting rules for the inaccurate case, and the agnostic rules of the two models—the ranking model and the tournament model—as a function of their noise parameter φ (note that, while the limiting rules were derived for a specific range of the noise parameter φ , we compare them at *all* values of φ). An interesting observation is that while the dissimilarity is quite low in both the accurate and the agnostic cases, it is surprisingly low in the inaccurate case. This observation holds true for all combinations of (n, m) . This indicates that the MLE rules of the tournament model are in general good approximations of the MLE rules of the ranking model, and the approximation becomes very good for the rules derived under the assumption of very high noise.

Recall that Theorem 7 establishes that Tideman's rule is a 2-approximation of Kemeny's rule in the worst case (in terms of the Kemeny score of the winner). This is a significant improvement over the Borda count, which is known to give a 4-approximation in the worst case [11]. Thus, Tideman's rule improves over Borda's rule by a factor of 2 in the worst case. It is interesting to check if this relationship holds even in the average case.

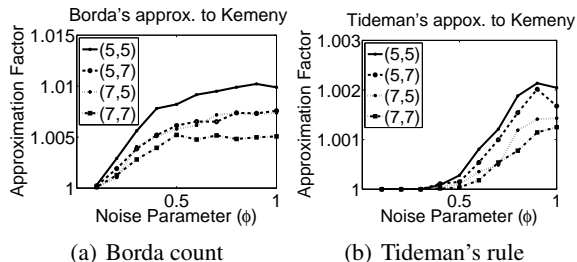


Figure 2: Approximations of Kemeny's rule.

Figures 2(a) and 2(b) show the average-case approxima-




Figure 1: The dissimilarity between the ranking model and the tournament model.

tion factors of the Borda count and Tideman’s rule, respectively, as a function of the noise parameter φ . It is evident that for both rules their average-case approximation factors are much better than their worst-case approximation factors. However, in the average case, the Borda count quickly reaches an approximation ratio of 1.01, while the approximation ratio of Tideman’s rule stays well below 1.003. That is, the improvement of Tideman’s rule over the Borda count is at least as good—in fact, slightly better—in the average case as in the worst case.




Figure 3: The average number of winning alternatives returned by the agnostic rules.

In our second set of experiments, we analyze the average number of winning alternatives returned by the agnostic rules in the ranking and the tournament models, again as a function of the noise parameter φ . Figures 3(a) and 3(b) show the results for the ranking and the tournament models, respectively. It can be seen that the agnostic rule—despite returning a set of alternatives that is guaranteed to contain the MLE best alternative for all values of $\varphi \in (0, 1)$ —outputs an average of less than 1.3 and 1.2 alternatives in the ranking model and the tournament model, respectively. In fact, in our simulations, the agnostic rules in both models return a single alternative, which is guaranteed to be the MLE best alternative for all values of φ , more than 80% of the time, for every $\varphi \in (0, 1)$.

8 DISCUSSION

We have studied methods for picking the best alternative given noisy estimates of an objective true comparison be-

tween the alternatives. Besides studying the standard Mallows’ model where both the ground truth and the estimates are acyclic total orders, we introduced and studied the setting where both may contain cycles. Procaccia et al. [23] studied the case where the ground truth is acyclic, but the estimates may or may not be cyclic. The only case not studied in the literature is that of possibly cyclic ground truth and acyclic estimates. However, this setting does not appear natural, and is also technically challenging: the denominator Z_φ^m in the probability expression of Mallows’ model (Equation 1) would not be independent of the ground truth, rendering the analysis extremely difficult.

Generalizations of Mallows’ model have been proposed in the literature [16, 21]. Some of these use critical information regarding positions of the alternatives in the ground truth ranking. Future work may also involve adapting such models to the case of tournaments; for example, one can use the number of alternatives defeated by a given alternative as a proxy for its rank, or one can develop distance metrics over tournaments to replace the Kendall tau distance in Equation (1). It would be interesting to see if such adaptations provide tractable approximations of the original ranking model.

The maximum likelihood approach to voting focuses solely on maximizing the likelihood of selecting the best alternative. This results in voting rules that can be difficult to understand, but have performance guarantees nonetheless. While simplicity is usually an important goal in the design of voting rules, it is less of an issue in human computation contexts, where the workers are paid for their input and do not need to know or understand how their estimates would be aggregated. Yet, in some practical applications, one may wish to use rules with additional desirable properties, either motivated by the application itself or as a safeguard in case the assumptions about the nature of the noise fail. A very exciting direction is to use Mallows’ model in order to inform the design of voting rules while trading off some of the likelihood for axiomatic properties such as Condorcet consistency or monotonicity.

References

- [1] N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: Ranking and clustering. In *Proc. of 37th STOC*, pages 684–693, 2005.
- [2] H. Azari Soufiani, D. C. Parkes, and L. Xia. A statistical decision-theoretic framework for social choice. In *Proc. of 28th NIPS*, 2014. Forthcoming.
- [3] J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6:157–165, 1989.
- [4] Y. Bugeaud and M. Mignotte. Polynomial root separation. *International Journal of Number Theory*, 6(3):587–602, 2010.
- [5] I. Caragiannis, C. Kaklamanis, N. Karanikolas, and A. D. Procaccia. Socially desirable approximations for Dodgson’s voting rule. *ACM Transactions on Algorithms*, 10(2):1–28, 2014.
- [6] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? In *Proc. of 14th EC*, pages 143–160, 2013.
- [7] G. E. Collins and A. G. Akritas. Polynomial real root isolation using Descartes’s rule of signs. In *Proc. of 3rd ISSAC*, pages 272–275, 1976.
- [8] V. Conitzer, A. Davenport, and H. Kalagnanam. Improved bounds for computing Kemeny rankings. In *Proc. of 21st AAAI*, pages 620–626, 2006.
- [9] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proc. of 21st IJCAI*, pages 109–115, 2009.
- [10] V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *Proc. of 21st UAI*, pages 145–152, 2005.
- [11] D. Coppersmith, L. Fleischer, and A. Rudra. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *Proc. of 17th SODA*, pages 776–782, 2006.
- [12] M. de Condorcet. *Essai sur l’application de l’analyse à la probabilité de décisions rendues à la pluralité de voix*. Imprimerie Royal, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.
- [13] C. L. Dodgson. *A Method for Taking Votes on More than Two Issues*. Clarendon Press, 1876.
- [14] C. Kenyon-Mathieu and W. Schudy. How to rank with few errors. In *Proc. of 39th STOC*, pages 95–103, 2007.
- [15] E. Law and L. von Ahn. *Human Computation*. Morgan & Claypool, 2011.
- [16] G. Lebanon and J. Lafferty. Cranking: Combining rankings using conditional probability models on permutations. In *Proc. of 9th ICML*, pages 363–370, 2002.
- [17] T. Lu and C. Boutilier. Learning Mallows models with pairwise preferences. In *Proc. of 28th ICML*, pages 145–152, 2011.
- [18] K. Mahler. An inequality for the discriminant of a polynomial. *The Michigan Mathematical Journal*, 11(3):257–262, 1964.
- [19] C. L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957.
- [20] A. Mao, A. D. Procaccia, and Y. Chen. Better human computation through principled voting. In *Proc. of 27th AAAI*, pages 1142–1148, 2013.
- [21] M. Meila, K. Phadnis, A. Patterson, and J. A. Bilmes. Consensus ranking under the exponential model. In *Proc. of 23rd UAI*, pages 285–294, 2007.
- [22] S. Merrill and B. Grofman. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press, 1999.
- [23] A. D. Procaccia, S. J. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proc. of 28th UAI*, pages 695–704, 2012.
- [24] T. C. Ratliff. A comparison of Dodgson’s method and Kemeny’s rule. *Social Choice and Welfare*, 18(1):79–89, 2001.
- [25] N. Tideman. *Collective Decisions and Voting*. Ashgate, 2006.
- [26] J. V. Uspensky. *Theory of equations*. McGraw-Hill New York, 1948.
- [27] A. v. Zuylen and D. P. Williamson. Deterministic algorithms for rank aggregation and other ranking and clustering problems. In *Proc. of 5th WAOA*, pages 260–273, 2007.
- [28] L. Xia and V. Conitzer. A maximum likelihood approach towards aggregating partial orders. In *Proc. of 22nd IJCAI*, pages 446–451, 2011.
- [29] L. Xia, V. Conitzer, and J. Lang. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proc. of 9th AAMAS*, pages 399–408, 2010.
- [30] H. P. Young. Condorcet’s theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.