

Eliciting Kemeny Rankings

Anne-Marie George¹, Christos Dimitrakakis^{1, 2}

¹University of Oslo, Norway

²University of Neuchatel, Switzerland
annemage@uio.no, christos.dimitrakakis@unine.ch

Abstract

We formulate the problem of eliciting agents' preferences with the goal of finding a Kemeny ranking as a Dueling Bandits problem. Here the bandits' arms correspond to alternatives that need to be ranked and the feedback corresponds to a pairwise comparison between alternatives by a randomly sampled agent. We consider both sampling with and without replacement, i.e., the possibility to ask the same agent about some comparison multiple times or not.

We find approximation bounds for Kemeny rankings dependant on confidence intervals over estimated winning probabilities of arms. Based on these we state algorithms to find Probably Approximately Correct (PAC) solutions and elaborate on their sample complexity for sampling with or without replacement. Furthermore, if all agents' preferences are strict rankings over the alternatives, we provide means to prune confidence intervals and thereby guide a more efficient elicitation. We formulate several adaptive sampling methods that use look-aheads to estimate how much confidence intervals (and thus approximation guarantees) might be tightened. All described methods are compared on synthetic data.

1 Introduction

Decision making based on the preferences of a group of individuals' is tackled in various settings by AI systems. To ensure a good societal outcome, the fair aggregation of the society members' preferences over alternatives plays a central role. Many such aggregation functions for agents' or voters' preferences have been proposed by politicians, mathematicians, economists and computer scientists. Research in the area of Social Choice and more specifically Voting Theory analyses and develops such aggregation functions algorithmically and axiomatically (Brandt et al. 2016). While impossibility results force all such functions to be flawed by computational complexity or violation of fairness guarantees, e.g. (Arrow 1950; Satterthwaite 1975), many are well researched and applied in practise. However, it is usually assumed that voters specify their preferences completely before the desired aggregation can be determined. This can be excessively time consuming or costly for voters, especially for large numbers of alternatives as in movie or

online-shopping platforms. Next we describe our contributions for efficiently eliciting aggregated rankings under Kemeny's rule, related work and outline of this paper.

Contributions and Related Work

We formulate the problem of eliciting voter preferences as a Dueling Bandit problem (Bengs et al. 2021) in which arms of bandits correspond to alternatives, and the bandit feedback for pulling two arms corresponds to a random voter's preference over the two arms. This feedback is then summarised in a matrix of winning probabilities over arms. Our goal is to elicit an aggregated ranking of voter preferences based on some voting rule. Here, the duelling bandit formulation is applicable for any voting rule that is a *C2 function* according to Fishburn's classification of preference aggregation functions (Fishburn 1977). That is, functions that only require the margin of voters that prefer one alternative over another as input, and not the exact combinations of pairwise preferences the voters have. For such functions, we can thus assume the input to be a *preference matrix* indicating the fraction of voters preferring an alternative over another for all pairs of alternatives. The results by (Boehmer et al. 2022) imply that it is NP-hard to decide whether a given matrix corresponds to a set of voters' preference orders (i.e., linear orders) over the alternatives. It follows that it is not possible to find a characterisation of such matrices that can be checked in polynomial time. Nevertheless, we identify properties of such matrices.

Contribution 1: In Section 2, we connect the vote elicitation and duelling bandit frameworks for C2 functions, and propose three properties of preference matrices, which will help devising adaptive elicitation strategies: Completeness, Triangle Inequality and Realisable Borda Scores.

Next, we focus our attention on the *voting rule* devised by John Kemeny in 1959 (Kemeny 1959), which is a preference aggregation function with compelling properties such as monotonicity and Condorcet winner consistency (Young and Levenglick 1978). The output of this rule is a set of rankings over the alternatives which minimise the *Kemeny score*. The Kemeny score measures voter disagreement in terms of pairwise disagreement over the ranked alternatives, i.e., Kendall's tau distance (Kendall 1938). In general, it is NP-hard to compute Kemeny rankings (Bartholdi, Tovey, and Trick 1989) but polynomial time solvable for specific

types of preferences, such as single-peaked ones (Brandt et al. 2015). A perk of Kemeny’s rule, as a C2 function, is that it can be readily applied to any matrix. That is, even when voters only specify parts of their preferences, i.e., an incomplete order over the alternatives, or when their preferences are non-transitive. The corresponding matrix in that case only contains fractions of voters that specified their preference over two alternatives. This enables us to compute an approximate Kemeny ranking even when not all preferences have been elicited yet. Our goal then is to sample pairs of arms (i.e., query voters to compare two alternatives) such that the approximate Kemeny ranking based on the sample means is close (for some distance measure) to an optimal Kemeny ranking. Note that while approximation algorithms exist for Kemeny rankings, e.g., (Betzler et al. 2008), the approximation bounds usually concern the approximation w.r.t. a given instance. Contrarily, we consider the approximation of Kemeny rankings w.r.t. a similar instance while applying an exact method.

Contribution 2: In Section 3, we formulate approximation bounds for (1) Kendall tau distance and (2) Kemeny Score difference, between the true Kemeny ranking of a matrix of winning probabilities and the Kemeny ranking of an approximate matrix. These bounds are based on confidence intervals around sample means for the winning probabilities.

Note that bandit problems with the goal of finding a ranking over the arms (as opposed to one or several winning arms) have been analysed for the C2 functions Borda and Copeland, see e.g. (Busa-Fekete, Szörényi, and Hüllermeier 2014; Falahatgar et al. 2017; Lin and Lu 2018), and other winner concepts such as von Neumann winners (Balsubramani et al. 2016). These works establish *probably approximately correct* (PAC) algorithms and analyse their sample complexity. To the best of our knowledge such considerations do not exist for Kemeny rankings. In fact, (Bengs et al. 2021) propose Kemeny rankings for future work. Also, no related work considers sampling from a fixed population without replacement (as would be reasonable when given a set of voters which shall not be queried twice about the same pair of alternatives). What is more, approximations of rankings w.r.t. Kemeny scores (i.e., agents’ disagreement) have not been considered under any aggregation function so far.

Contribution 3: In Section 4, we formulate PAC algorithms for Kemeny rankings w.r.t. Kemeny scores of a matrix of winning probabilities for two distinct cases of sampling methods: (1) sampling bandit feedback from voters u.a.r., i.e., sampling with replacement, (2) sampling bandit feedback from a hypergeometric distribution, i.e., sampling without replacement, meaning that voters cannot be asked multiple times about the same pair of alternatives. These PAC algorithms are based on our approximation results (see Contribution 2) and do not assume any properties for the matrix of winning probabilities. We give concrete sample complexities for desired approximation values and probabilities.

Contribution 4: In Section 5 we show how to prune confidence intervals for winning probabilities of arms based on the properties of preference matrices (see Contribution 1), i.e., for the case that the matrix originates from a set of (linear) preference orders. We discuss several adaptive sampling

methods that use look-aheads to anticipate possible effects of pruning. Finally, we analyse the sample complexity experimentally on synthetic data for all developed methods.

An extended version of this paper including more details, experimental results and all proofs of results marked with ★ can be found under <https://arxiv.org/abs/2312.11663>.

2 Voting and Bandits

In a voting setting, we assume to have a set N of voters and a set $C = [k]$ of k candidates or alternatives. Let $\mathcal{L}(C)$ be the set of all linear orders or rankings over C . We usually consider the preferences of the voters over the alternatives to be linear orders $\succ_v \in \mathcal{L}(C)$ for $v \in N$. A typical goal in this setting is to aggregate the preferences of voters into an outcome ranking or to determine a winner. A social preference function is a function that maps a *preference profile* $P = \{\succ_v \mid v \in N\}$ to one or several output rankings. One example of such a rule is Kemeny’s rule, which aims to minimise the disagreement of the voters with the output ranking.

In a duelling bandits setting, we have a set $C = [k]$ of k arms of bandits. In every time step, we pull two arms yielding feedback of which one is better. That is, when pulling arms i and j we get feedback $i \succ j$ with some (unknown) probability q_{ij} and feedback $j \succ i$ with inverse probability $q_{ji} = 1 - q_{ij}$. The winning probabilities can be summarised in a matrix $Q \in [0, 1]^{k \times k}$ where by convention $q_{ii} = 0.5$ for all $i \in [k]$. Let $\mathcal{Q}(k) = \{Q \in [0, 1]^{k \times k} \mid q_{ji} = 1 - q_{ij} \text{ and } q_{ii} = 0.5 \forall i, j \in [k]\}$ be the set of such matrices. We assume the true matrix of winning probabilities is unknown and must be approximated, with the ultimate goal of finding a good ranking of the arms. This goal can be expressed by a measure of regret and literature in this field usually analyses regret bounds in the number of samples.

To connect the two settings, we make the following observation: A preference profile $P \in \mathcal{L}(C)^{|N|}$ of voters N over alternatives $C = [k]$ can be translated to a preference matrix $Q \in \mathcal{Q}(k)$ by defining $q_{ij} = \frac{|\{v \in N \mid i \succ_v j\}|}{|N|}$ and $q_{ii} = 0.5$ for all $i, j \in [k]$. Thus, the alternatives are seen as arms of bandits and the winning probabilities q_{ij} are equal to the fraction of voters preferring alternative i over j . We denote the set of such *preference matrices* that are originating from preference profiles by $\mathcal{P}(k, n)$, where k is the number of alternatives and n the number of voters. Because voters’ preferences are linear orders, and transitive, $\mathcal{P}(k, n) \subsetneq \mathcal{Q}(k)$. We show further properties of such matrices.

Lemma 1 (★). *Let $P = \{\succ_v \in \mathcal{L}([k]) \mid v \in [n]\}$ be a profile of n voters and k alternatives. Then $Q \in \mathcal{P}(k, n)$ with $q_{ij} = \frac{1}{n} |\{v \in [n] \mid i \succ_v j\}|$ and $q_{ii} = 0.5$ satisfies:*

1. $q_{ij} \cdot n \in [n]$ and $q_{ij} + q_{ji} = 1$ for all distinct $i, j \in [k]$.
(Completeness)
2. $q_{lj} + q_{ji} \geq q_{li}$ for all distinct $i, j, l \in [k]$.
(Triangle Inequality)
3. $\sum_{i \in A} \sum_{j \in [k] \setminus \{i\}} q_{ij} \leq \frac{1}{n} \cdot |A| \cdot (k - \frac{|A|+1}{2})$
for all $A \subseteq [k]$.
(Realisable Borda Scores)

The next example shows that mapping preference profiles to preference matrices is not injective.

Example 1. Consider the two preference profiles P_1, P_2 over three alternatives below.

$$P_1 = \{1 \succ 2 \succ 3, 1 \succ 2 \succ 3, 3 \succ 2 \succ 1\}$$

$$P_2 = \{1 \succ 2 \succ 3, 2 \succ 1 \succ 3, 3 \succ 1 \succ 2\}.$$

The two profiles map to the same preference matrix Q :

$$Q = \begin{bmatrix} \frac{1}{2} & \frac{2}{3} & \frac{2}{3} \\ \frac{1}{3} & \frac{1}{2} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{2} \end{bmatrix}.$$

3 Kemeny’s Rule

Kemeny’s Rule outputs (after applying some tie-breaking) a strict ranking of alternatives / arms that minimises the disagreement of the voters. To measure disagreement of voters with a ranking we can apply Kendall’s tau distance (or “inversion distance”). The Kendall-tau distance between two preference orders $\succ_v, \succ_w \in \mathcal{L}(C)$ is defined as follows (where \mathbb{I} is the indicator function):

$$d_{\text{KT}}(\succ_v, \succ_w) = \sum_{\{i,j\} \subseteq C} \mathbb{I}(i \succ_v j) \cdot \mathbb{I}(j \succ_w i)$$

$$= \sum_{(i,j) \in \succ_v} \mathbb{I}(j \succ_w i).$$

The Kemeny score of a ranking $\succ \in \mathcal{L}(C)$ given the preference rankings of voters N is defined as $\text{KS}(\{\succ_v\}_{v \in N}, \succ) = \sum_{v \in N} d_{\text{KT}}(\succ, \succ_v)$ and the set of Kemeny rankings (before tie-breaking) is $\text{K}(\{\succ_v\}_{v \in N}) = \text{argmin}_{\succ \in \mathcal{L}(C)} \text{KS}(\{\succ_v\}_{v \in N}, \succ)$. In the same spirit of measuring disagreement, we can define the Kemeny rankings for a given matrix $Q \in [0, 1]^{k \times k}$, where q_{ij} is the probability that arm i beats arm j . The Kemeny score of a ranking $\succ \in \mathcal{L}(C)$ with $C = [k]$ is proportional to the probability that the outcome of any random sample of a pair of arms is not aligned with \succ , i.e., $\text{KS}(Q, \succ) = \sum_{(i,j) \in \succ} q_{ji}$ and the set of Kemeny rankings is $\text{K}(Q) = \text{argmin}_{\succ \in \mathcal{L}(C)} \text{KS}(Q, \succ)$.

In general, Kemeny’s rule is not resolute, i.e., the minimum Kemeny score may be assumed by several rankings. A tie-breaking rule allows to choose one among the rankings that minimise the Kendall-tau distance. In this paper, we assume to break ties deterministically whenever necessary.

Given only an approximation of a preference matrix Q , we can also only approximate its Kemeny ranking. To measure the quality of such an approximation, two natural distance measures for such rankings can be applied: the Kendall-tau distance between the true and the approximate ranking, or the difference in their Kemeny scores w.r.t. the true matrix Q . In the following, we analyse approximation guarantees based on both measures. While Kemeny’s rule is NP hard to compute, there exist polynomial time approximation algorithms (Betzler et al. 2008). For an approximate Kemeny ranking of an approximate matrix, we can simply combine the approximation bounds. Assume matrix Q has Kemeny ranking τ and its approximation matrix \hat{Q} has Kemeny ranking $\hat{\tau}$ that is a ρ -approximation w.r.t. Kemeny-scores, i.e., $\text{KS}(Q, \hat{\tau}) \leq \text{KS}(Q, \tau) + \rho$. Then an additive α -approximation $\hat{\tau}$ of $\hat{\tau}$, i.e., $\text{KS}(\hat{Q}, \hat{\tau}) \leq \text{KS}(\hat{Q}, \hat{\tau}) + \alpha$, yields $\text{KS}(Q, \hat{\tau}) < \text{KS}(Q, \tau) + \rho + \alpha$. Similarly, a multiplicative λ -approximation $\hat{\tau}$ with $\text{KS}(\hat{Q}, \hat{\tau}) \leq (1 + \lambda)\text{KS}(\hat{Q}, \hat{\tau})$ yields $\text{KS}(Q, \hat{\tau}) < (1 + \lambda)\text{KS}(Q, \tau) + (1 + \lambda)\rho$.

Kendall-tau Distance of Approx. Kemeny Rankings

Consider two rankings $\tau, \hat{\tau} \in \mathcal{L}(C)$ with $|C| = k$. If the two rankings are equal, their Kendall-tau distance is 0; if they are reversed, their Kendall-tau distance is maximal, i.e., $\frac{(k-1)k}{2}$. More generally, the Kendall-tau distance of any two rankings $\tau, \hat{\tau} \in \mathcal{L}(C)$ with $|C| = k$, is an integer in $\{0, \dots, \frac{(k-1)k}{2}\}$. The following lemma shows that the Kendall-tau distance of the Kemeny rankings $\text{K}(Q)$ and $\text{K}(\hat{Q})$ of a preference matrix Q and a close approximation \hat{Q} of Q can be maximal. This is proven by construction of an example with values in Q that are very close to 0.5. Thus, no matter the number of samples, when considering the Kemeny ranking based on the sample means, we cannot guarantee any improvement in terms of Kendall-tau distance.

Lemma 2 (★). For any small¹ $\epsilon > 0$, there exist matrices $Q, \hat{Q} \in \mathcal{Q}(k)$ with $\|Q - \hat{Q}\|_1 = \epsilon$ and $\tau = \text{K}(Q)$ and $\hat{\tau} = \text{K}(\hat{Q})$ such that $d_{\text{KT}}(\tau, \hat{\tau}) = \frac{(k-1)k}{2}$.

If the matrix of winning probabilities is indeed a preference matrix, i.e., reflecting fractions of n voters, then a reasonable restriction is to ask every voter at most once about their preference over a pair of arms. Disregarding inference reached through transitivity of voter preferences, a total of $n \cdot \frac{(k-1)k}{2}$ queries is necessary to recover the true preference matrix. However, if values in the preference matrix are close to 0.5, even asking $(n-1) \cdot \frac{(k-1)k}{2}$ queries can still lead to an approximate Kemeny ranking with maximal Kendall-tau distance to the true Kemeny ranking.

Lemma 3 (★). For any number of arms $k > 2$ and number of voters $n > 2$, there exist preference profiles with corresponding preference matrices $Q \in \mathcal{P}(k, n)$ and $\hat{Q} \in \mathcal{P}(k, n-1)$ with $\|Q - \hat{Q}\|_1 = \frac{(k-1)k}{2(n-1)}$ and $\tau = \text{K}(Q)$ and $\hat{\tau} = \text{K}(\hat{Q})$ such that $d_{\text{KT}}(\tau, \hat{\tau}) = \frac{(k-1)k}{2}$.

This shows that especially for large numbers of voters, the approximation error for approximating matrix Q could be very small, while the Kendall-tau distance is maximal. A similar result can be obtained when querying voters about all but one pair of alternatives. Consider a Condorcet cycle for which $q_{12} = q_{23} = q_{31} = 2/3$. Then the tie-breaking rule, e.g., $1 \succ 2 \succ 3$, coincides with the Kemeny ranking with a score of $4/3$. When eliciting all preferences over all pairs of arms but $(2, 3)$, and initialising $q_{23} = q_{32} = 0.5$, the approximated Kemeny ranking is $3 \succ 2 \succ 1$, independent of tie-breaking, yielding a maximal KT-distance.

The results in this section are based on examples for which Kemeny rankings have very high Kemeny scores. Thus, while yielding very different Kemeny rankings, the matrices in the proofs also yield similar voter satisfaction. In the next section, we concentrate on bounds of Kemeny score differences between Kemeny rankings of similar matrices.

Kemeny Score of Approx. Kemeny Rankings

We can bound the difference of Kemeny scores for two Kemeny rankings of similar matrices.

¹For the construction of our proof, we require $\epsilon < \frac{(k-1)k}{2}$.

Lemma 4 (★). Let $Q, \hat{Q} \in \mathcal{Q}(k)$ with $\|Q - \hat{Q}\|_1 \leq \epsilon$. Furthermore, let τ be a Kemeny ranking of Q and $\hat{\tau}$ a Kemeny ranking of \hat{Q} . Then $KS(Q, \hat{\tau}) - KS(Q, \tau) \leq \epsilon$.

Note that because τ is a Kemeny ranking w.r.t. Q , the difference of Kemeny scores of τ and $\hat{\tau}$ cannot be negative. Lemma 4 shows that with a close approximation of a matrix we can gain a Kemeny ranking that approximates the Kemeny score, i.e., minimal voter disagreement, well.

When sampling preferences, or more specifically feedback from dueling bandits, we can express the value of our approximation of the winning matrix by confidence intervals for matrix entries. For this type of uncertainty quantification, we can adjust the approximation bounds from Lemma 4.

Corollary 1. Let $Q, \hat{Q} \in \mathcal{Q}(k)$. Consider Kemeny rankings τ of Q and $\hat{\tau}$ of \hat{Q} . Suppose $q_{ij} \in [\hat{q}_{ij} - c_{ji}, \hat{q}_{ij} + c_{ij}]$ for some matrix $C \in \mathbb{R}_{\geq 0}^{k \times k}$ holds for all $i, j \in [k], i \neq j$ with probability $(1 - \delta)$. Then $KS(Q, \hat{\tau}) - KS(Q, \tau) \leq 2 \sum_{1 \leq i < j \leq k} \max(c_{ij}, c_{ji})$ with probability $(1 - \delta)$.

We can, in fact, refine this approximation bound for a Kemeny ranking $\hat{\tau}$ of $\hat{Q} + C$, with confidence bounds C .

Lemma 5 (★). Let $Q, \hat{Q} \in \mathcal{Q}(k)$ and $C \in \mathbb{R}^{k \times k}$ be matrices such that $q_{ij} \in [\hat{q}_{ij} - c_{ji}, \hat{q}_{ij} + c_{ij}]$ holds for all $i, j \in [k], i \neq j$ with probability $(1 - \delta)$. Consider Kemeny rankings τ of Q and $\hat{\tau}$ of $\hat{Q} + C$. Then $KS(Q, \hat{\tau}) - KS(Q, \tau) \leq \sum_{1 \leq i < j \leq k} (c_{ij} + c_{ji})$ with probability $(1 - \delta)$.

Thus, to construct a good ranking in terms of Kemeny score, we can choose a Kemeny ranking w.r.t. upper confidences of an approximated matrix of winning probabilities.

The above described approximation bounds for Kemeny scores are dependent on the quality of the approximation of winning probabilities for all pairs of arms. This is unsurprising, since the Kemeny score itself takes all winning probabilities into account with equal weight. In the next section, we discuss concrete confidence bounds and PAC algorithms for Kemeny rankings w.r.t. Kemeny scores.

4 PAC Algorithms for Kemeny Rankings w.r.t. Kemeny Scores

To approximate a Kemeny ranking, we determine a sequence of pairs of arms to pull. We first consider strategies for the setting where we draw samples with replacement from Bernoulli distributions with the winning probabilities as means. Thus, voters are drawn uniformly at random from an unknown population and we might ask the same voter the same query twice. We then investigate the setting in which the population of voters is known and every voter may only be asked the same query once (sampling from a finite population without replacement).

Sampling With Replacement

Consider a matrix $Q \in \mathcal{Q}(k)$ of winning probabilities between arms such that every sample (i, j) of two arms returns $i \succ j$ with probability q_{ij} . Because confidence interval sizes are usually monotonically decreasing with the number of samples, Lemma 5 suggests that for symmetric confidence bounds and without taking any further inference into

account a uniform sampling strategy minimises the Kemeny-score-approximation error. More concretely, it follows from the next proposition that Algorithm 1 is a PAC algorithm for approximating Kemeny rankings w.r.t. Kemeny scores.

Algorithm 1:

KemenyEl(Q, ρ, δ): Sampling with Replacement from $Q \in \mathcal{Q}(k)$

$\hat{Q} \leftarrow \{1/2\}^{k \times k}$	▷ Estimate of Q
$x := \frac{k \cdot (k-1)}{\rho}, y := \ln \frac{\delta}{k(k-1)}$	▷ Constants
$t := \lceil -\frac{1}{2}x^2y \rceil, c := \sqrt{-\frac{1}{2t}y}$	▷ Sample Size and Confidence Bound
for arms $i, j \in [k]$ do	
$s_1, \dots, s_t \sim \text{Bernoulli}(q_{ij})$	
$\hat{q}_{ij} \leftarrow 1/t \sum_{i=1, \dots, t} s_i$	
$\hat{q}_{ji} \leftarrow 1 - \hat{q}_{ij}$	
end for	
return $K(\hat{Q} + \mathbb{1} \cdot c)$	

Proposition 1 (★). Let $\hat{Q} \in \mathcal{Q}(k)$ be the matrix of sample averages approximating some $Q \in \mathcal{Q}(k)$ after t samples of each pair of arms, and let $c_\delta^t = \sqrt{\frac{1}{2t}y}$ with $y := -\ln \frac{\delta}{k(k-1)} > 1$. Then, with probability $(1 - \delta)$ we have $KS(Q, \hat{\tau}) - KS(Q, \tau) \leq k \cdot (k-1) \cdot c_\delta^t$ for Kemeny ranking τ of Q and Kemeny ranking $\hat{\tau}$ of $\hat{Q} + \mathbb{1} \cdot c_\delta^t$.

Given a desired approximation rate ρ , we can reformulate Proposition 1 and obtain the necessary number of samples.

Corollary 2. For $k \geq 2$ arms, $Q \in \mathcal{Q}(k)$ with Kemeny ranking τ , fixed approximation guarantee $0 < \rho \leq \frac{k(k-1)}{2}$ and probability δ , Algorithm 1 is a (δ, ρ) -PAC algorithm. That is, after $t < \frac{1}{2}x^2y$ samples of each pair of arms, Algorithm 1 returns ranking $\hat{\tau}$ such that with probability $(1 - \delta)$, we have $KS(Q, \hat{\tau}) - KS(Q, \tau) \leq \rho$. Here $x := \frac{k \cdot (k-1)}{\rho}$ and $y := -\ln \frac{\delta}{k(k-1)} > 1$.

Sampling Without Replacement

Suppose that, while sampling voters uniformly at random, none of the voters $v \in N$ can be asked twice about their preferences over a pair of arms. Then the number s out of t asked voters that prefer i to j follows a hypergeometric distribution: $s \sim \text{Hypergeometric}(n, n \cdot q_{ij}, t)$. In this case, we can refine the confidence bounds of our sampled estimation of the true fractions of voters q_{ij} preferring i to j . This leads to better approximation for a fixed number of samples, or reversely, a lower sample complexity for a fixed approximation goal, compared to sampling with replacement. More concretely, it follows from the next propositions that Algorithm 2 is a PAC algorithm for Kemeny rankings w.r.t. Kemeny scores when sampling without replacement.

Proposition 2 (★). Consider $n > 1$ voters with preference matrix $Q \in \mathcal{P}(k, n)$ over $k \geq 2$ alternatives. Let $\delta < 0.5$ be the approximation probability and $y := -\ln \frac{\delta}{k(k-1)} > 1$.

Algorithm 2:

 KemenyEl(Q, ρ, δ): Sampling w/o Replacement from $Q \in \mathcal{P}(k, n)$

```

 $\hat{Q} \leftarrow \{1/2\}^{k \times k} \quad \triangleright$  Estimate of  $Q$ 
 $x := \frac{k \cdot (k-1)}{\rho}, y := -\ln \frac{\delta}{k(k-1)} \quad \triangleright$  Constants
if  $n < \frac{x^2 y - 4}{2}$  then
     $t := \frac{-2n + x^2 y n}{2n + x^2 y}, c := \sqrt{\frac{(n-t)(t+1)}{2t^2 n}} y \quad \triangleright$  Sample Size and Confidence Bound (small  $n$ )
else
     $t := \frac{x^2 y(n+1)}{2n + x^2 y}, c := \sqrt{\frac{n-t+1}{2tn}} y \quad \triangleright$  Sample Size and Confidence Bound (large  $n$ )
end if
for arms  $i, j \in [k]$  do
     $s \sim \text{Hypergeometric}(n, n \cdot q_{ij}, t)$ 
     $\hat{q}_{ij} \leftarrow s/t$ 
     $\hat{q}_{ji} \leftarrow 1 - \hat{q}_{ij}$ 
end for
return  $K(\hat{Q} + \mathbb{1} \cdot c)$ 
    
```

Furthermore let $\hat{Q} \in \mathcal{Q}(k)$ be the matrix of sample averages after t uniformly random samples from the population for each pair of arms (without replacement).

Then $c_{\delta, n}^t = \sqrt{\frac{n-t+1}{2tn}} y$ is a $(1 - \delta)$ -confidence bound for all entries in \hat{Q} . Furthermore, with probability $(1 - \delta)$ we have $\text{KS}(Q, \hat{\tau}) - \text{KS}(Q, \tau) \leq k \cdot (k-1) \cdot c_{\delta, n}^t$ for Kemeny ranking τ of Q and Kemeny ranking $\hat{\tau}$ of $\hat{Q} + \mathbb{1} \cdot c_{\delta, n}^t$.

For given ρ , we can guarantee an approximation $\text{KS}(Q, \hat{\tau}) - \text{KS}(Q, \tau) \leq \rho$ with probability $(1 - \delta)$ after $t = \frac{x^2 y(n+1)}{2n + x^2 y}$ samples (without replacement) of each pair of arms, where $x := \frac{k \cdot (k-1)}{\rho}$.

For the case that we have more samples than half of the population size, i.e., $t > n/2$, and a limited population size, we can apply better bounds on confidences over the winning probabilities. This leads to the following result.

Proposition 3 (★). Assume a fixed voter population of size $n > 1$ with preference matrix $Q \in \mathcal{P}(k, n)$ over $k \geq 2$ alternatives. Let $\delta < 0.5$ be the approximation probability and $y := -\ln \frac{\delta}{k(k-1)} > 1$. Furthermore, let $\hat{Q} \in \mathcal{Q}(k)$ be the matrix of sample averages after $t > n/2$ uniformly random samples for each pair of arms (without replacement).

Then $c_{\delta, n}^t = \sqrt{\frac{(n-t)(t+1)}{2t^2 n}} y$ is a $(1 - \delta)$ -confidence bound for all entries in \hat{Q} . Furthermore, with probability $(1 - \delta)$ we have $\text{KS}(Q, \hat{\tau}) - \text{KS}(Q, \tau) \leq k \cdot (k-1) \cdot c_{\delta, n}^t$ for Kemeny ranking τ of Q and Kemeny ranking $\hat{\tau}$ of $\hat{Q} + \mathbb{1} \cdot c_{\delta, n}^t$.

For given ρ and if the population is of size $n < \frac{x^2 y - 4}{2}$, we can guarantee an approximation $\text{KS}(Q, \hat{\tau}) - \text{KS}(Q, \tau) \leq \rho$ with probability $(1 - \delta)$ after $t = \frac{x^2 y n + 2n}{2n + x^2 y}$ samples (without replacement) of each pair of arms, where $x := \frac{k \cdot (k-1)}{\rho}$.

Comparison

Note that Proposition 2 and 3 use improved confidence bounds for the estimated winning probabilities compared to Proposition 1. This is not surprising, since we are sampling from a fixed population without replacement, and expect the confidence bound to be dependent on the relation between the number of samples and the population size. Furthermore, unsurprisingly, this improves the sample complexity. We can quantify the advantage as follows.

Corollary 3. Let $x := \frac{k \cdot (k-1)}{\rho}$ and $y := -\ln \frac{\delta}{k(k-1)} > 1$.

Consider the confidence bounds after t samples from Proposition 1, Proposition 2 and Proposition 3:

$$c = \sqrt{\frac{1}{2t}} y, c' = \sqrt{\frac{n-t+1}{2tn}} y, c'' = \sqrt{\frac{(n-t)(t+1)}{2t^2 n}} y.$$

Then we have $c' = \sqrt{\frac{n-t+1}{n}} c$ and $c'' = \sqrt{\frac{(n-t)(t+1)}{tn}} c$, i.e., tighter confidence bounds when sampling without replacement. Furthermore, $c' > c''$ if and only if $t > n/2$.

For fixed approximation guarantee $\frac{k(k-1)}{n} \geq \rho > 0$, consider the sample complexities from Proposition 1, 2, 3:

$$t = \frac{1}{2} x^2 y, t' = \frac{x^2 y(n+1)}{2n + x^2 y}, t'' = \frac{x^2 y n + 2n}{2n + x^2 y}.$$

Then we have $t' = \frac{2(n+1)}{2n + x^2 y} t$ and $t'' = \frac{2n + 4n/(x^2 y)}{2n + x^2 y} t$. For the reasonable assumption that $2/n < \rho$, $k \geq 2$ and $n \geq 2$ this means that fewer samples are needed when sampling without replacement than with replacement. Furthermore, for $n < \frac{x^2 y - 4}{2}$ we have $t' > t''$.

This comparison also shows that the more samples are taken, the starker the difference between confidence bounds when sampling with and without replacement. Furthermore, in order to reach an approximation of ρ , the smaller ρ and δ , the fewer samples we need when sampling without replacement compared to sampling with replacement.

5 Pruning of Confidence Intervals

In Section 4 we estimated the confidence intervals of the winning probabilities q_{ij} by the use of some known concentration inequalities. This treatment delivered symmetric bounds that, given the same number of samples, are the same for all pairs of arms $i, j \in [k]$. However, the values q_{ij} of a preference matrix $Q \in \mathcal{P}(k, n)$ are not independent. In fact, Lemma 1 implies that $q_{ij} + q_{ji} = 1$, for all $i, j \in [k]$. We can restrict confidence intervals to comply with this constraint.

Lemma 6 (★). Assume that the sampled means \hat{q}_{ij} after some samples from a preference matrix $Q \in \mathcal{P}(k, n)$ have (possibly asymmetric) confidence intervals:

$$q_{ij} \in [\hat{q}_{ij} - c_{ij}, \hat{q}_{ij} + \bar{c}_{ij}]$$

for all arms $i, j \in [k]$ with probability $(1 - \delta)$. Define matrix $C \in [0, 1]^{k \times k}$ by $c_{ij} := \min(\bar{c}_{ij}, c_{ji})$. Then with probability $(1 - \delta)$ for all arms $i, j \in [k]$:

$$q_{ij} \in [\hat{q}_{ij} - c_{ji}, \hat{q}_{ij} + c_{ij}]. \quad (1)$$

Lemma 1 also implies that for a preference matrix $Q \in \mathcal{P}(k, n)$ a triangle inequality must hold for all triples of arms: $q_{lj} + q_{ji} \geq q_{li}$ for all $i, j, l \in [k]$. We can comply with this constraint by pruning confidence intervals further.

Lemma 7 (★). *Assume that the sampled means \hat{q}_{ij} after some samples from a preference matrix $Q \in \mathcal{P}(k, n)$ have (possibly asymmetric) confidence intervals:*

$$q_{ij} \in [\hat{q}_{ij} - c_{ji}, \hat{q}_{ij} + c_{ij}]$$

for all arms $i, j \in [k]$ with probability $(1 - \delta)$. Define matrix $\hat{C} \in [0, 1]^{k \times k}$ by

$$\hat{c}_{ij} := \min(c_{ij}, \min_{l \in [k]} \hat{q}_{il} + c_{il} + \hat{q}_{lj} + c_{lj} - \hat{q}_{ij}).$$

Then with probability $(1 - \delta)$ for all arms $i, j \in [k]$:

$$q_{ij} \in [\hat{q}_{ij} - \hat{c}_{ji}, \hat{q}_{ij} + \hat{c}_{ij}]. \quad (2)$$

An example of pruning confidence intervals can be found in Appendix D of (George and Dimitrakakis 2023).

After pruning intervals once, it might be feasible to prune the new confidence intervals even more. This procedure is described by Algorithm 3. Note that Lemma 7 allows \hat{c}_{ij} to be negative, indicating that the sample mean \hat{q}_{ij} overestimates the actual winning probability q_{ij} . This does not influence the approximation result given in Lemma 5. However, we require that $1 \geq \hat{c}_{ij} + \hat{q}_{ij} \geq 0$ which is expressed in Algorithm 3 in the 4th and 10th line. Algorithm 3 converges, because in every iteration of the repeat-loop at least one value \hat{c}_{ij} strictly decreases and values never increase while maintaining a lower bound of $-\hat{q}_{ij}$. Furthermore, the values of initial confidence bounds and sample averages build a discrete set of values such that there are countably many combinations that can build the minimum for the updates.

We can now refine the elicitation methods from Section 4.

Algorithm 3: PruningCI($\hat{Q}, \underline{C}, \overline{C}$)

Pruning Confidence Intervals $q_{ij} \in [\hat{q}_{ij} - \underline{c}_{ij}, \hat{q}_{ij} + \overline{c}_{ij}]$ for all $i, j \in [k]$ of preference matrix $Q \in \mathcal{P}(k, n)$

$c_{ij} \leftarrow \min(\overline{c}_{ij}, \underline{c}_{ji})$ for all $i, j \in [k]$

▷ Upper/Lower Confidence Bounds based on Lemma 6

$\hat{c}_{ij} \leftarrow \min(c_{ij}, 1 - \hat{q}_{ij})$ for all $i, j \in [k]$

▷ Upper/Lower Confidence Bounds based on $q_{ij} \in [0, 1]$

$\hat{C} \leftarrow \hat{C}$

repeat

$\hat{C} \leftarrow \hat{C}$

for arms $i, j \in [k]$ **do**

$\hat{c}_{ij} \leftarrow \min(c_{ij}, \min_{l \in [k], l \neq i, j} \hat{q}_{il} + c_{il} + \hat{q}_{lj} + c_{lj} - \hat{q}_{ij})$

$\hat{c}_{ij} \leftarrow \max(-\hat{q}_{ij}, \hat{c}_{ij})$

end for

until $\hat{C} == \hat{C}$

return \hat{C}

Adaptive Sampling Strategies

Recall that our approximation bounds for Kemeny rankings from Section 3 are determined by the confidence intervals'

sizes. Furthermore, the confidence bounds for sampling with or without replacement given in Section 4 can only be influenced by the number of pulls of the pairs of arms. Thus uniform sampling guarantees to always sample from a pair of arms with the largest confidence interval. However, when pruning confidence intervals after every new sample, the "most uncertain" pair of arms, i.e., one with the largest confidence interval, might not necessarily be one with the lowest number of pulls so far. We can also estimate the outcome and effect of the next ℓ samples. To simplify notation and computational complexity, we concentrate on such look-aheads with $l = 1$. This allows the following sample strategies:

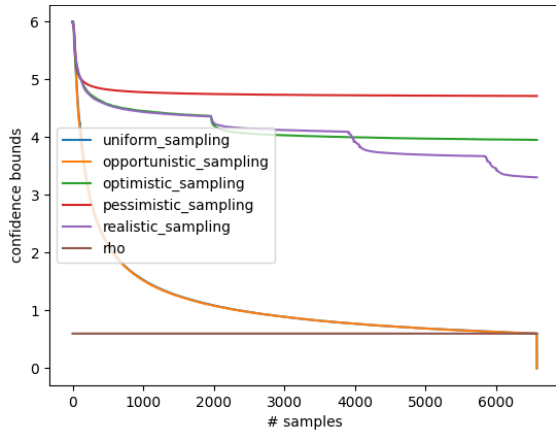
- **Uniform:** Sample a pair with minimal number of pulls.
- **Opportunistic:** Sample a pair with maximal confidence interval size.
- **Optimistic:** Sample a pair which, for the best case outcome, gives a maximal reduction in the sum of all confidence interval sizes.
- **Pessimistic:** Sample a pair which, for the worst case outcome, gives a maximal reduction in the sum of all confidence interval sizes.
- **Bayesian/Realistic:** Sample a pair which, in expectation w.r.t. the current estimation of winning probabilities, gives a maximal reduction in the sum of all confidence interval sizes.

Note that if in one round no pruning is possible for any sample outcome for any pair of arms, all sampling strategies collapse to opportunistic sampling. If additionally all pairs have been pulled equally often, this is the same as uniform sampling as any sample outcome will give same reduction.

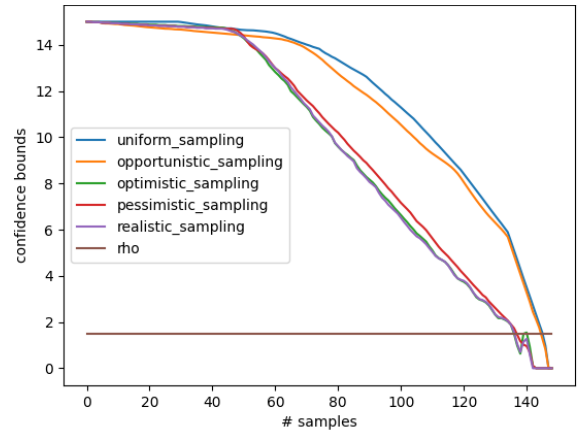
Experiments

Setup: We generate preference matrices $Q \in \mathcal{P}(k, n)$ for $n = 10$ voters uniformly at random for up to $k = 9$ arms. We set an approximation value of $\rho = 0.1 \frac{k(k-1)}{2}$, i.e., 10% of the worst case difference in approximated and optimal Kemeny score, and approximation probability $(1 - \delta) = 0.95$. For better comparison, we compare all sampling methods on the same preference matrices. Furthermore, we average all our results over runs on several instances. Note that the computational complexity for sampling with replacements only allows us to conduct experiments on 10 instances and for few arms, while for sampling without replacement we use 100 instances for experiments with up to 9 arms. To compare Kemeny scores at every sampling step, it is required to compute Kemeny rankings which is a known NP-hard problem. We use a simple ILP formulation for this. All code is written in Python and is publicly available under <https://github.com/annemage/Eliciting-Kemeny-Rankings>.

Results and Analysis: As all sampling methods correspond to uniform sampling when no pruning is taken into account, we restrict our attention here to the case where we apply Algorithm 3 to the confidence intervals indicated in Corollary 3 w.r.t. current numbers of pulls for every new sample. Unsurprisingly, our experiments show that the Kemeny score for a current approximation is closer to the true Kemeny score than the approximation bounds given by the



(a) Sampling With Replacement. Averages over 10 instances. $k = 4$, $\rho = 0.6$. Average true Kemeny score 2.08.



(b) Sampling Without Replacement. Averages over 100 instances. $k = 6$, $\rho = 1.5$. Average true Kemeny score 5.618.

Figure 1: Average Confidence Bounds for uniformly at random generated instances with $n = 10$ voters. $\delta = 0.05$.

total length of confidence intervals (see Lemma 5) would suggest, i.e., our approximation bounds are not tight. Because it is the approximation bounds, however, that influence the sample strategies, we show a comparison of these in Figure 1 for a sequence of samples. Note that here the length of the x-axis corresponds to the theoretical sample complexity given in Corollary 3. Obviously, the two plots confirm that sampling without replacement is much more efficient for all sample methods. Note that if for an instance sampling terminated because the approximation bound ρ was reached early, averages are taken over the remaining instances only.

It can be easily seen in Figure 1a that when sampling with replacement only uniform and opportunistic sampling reach the given approximation bound ρ within the theoretical sample complexity of 6576 samples. Note that in these experiments the look-ahead methods (optimistic, pessimistic and realistic sampling) often get stuck by sampling the same pairs of arms which apparently do not lead to enough pruning of confidence intervals to be worth it. This is surprising especially for pessimistic sampling, as this method should be conservative w.r.t. pruning possibilities. Furthermore, uniform and opportunistic sampling behave exactly the same and reach the desired approximation ρ in exactly the theoretical sample complexity. This might imply that pruning has no effect in these instances which could be an artifact of how the confidence intervals are computed in this case.

On the other hand, it can easily be seen from Figure 1b that for sampling without replacement even uniform sampling needs in average less than the theoretical number of samples indicated in Corollary 3 when pruning is applied, showing that pruning of confidence intervals is applicable in our instances. Here, the adaptive sampling strategies show a clear advantage over sampling uniformly. While opportunistic sampling seems the best choice when only sampling few pairs of arms, it is outperformed by optimistic, pessimistic and realistic (Bayesian) look-aheads for more samples. Naturally, the look-ahead methods can anticipate any possibility of pruning the confidence intervals, which seems to give

them an advantage in the long run. However, at this point we would like to point out that opportunistic sampling has much lower computational complexity than the look-ahead methods, which is an obvious advantage in practice. Optimistic and realistic sampling behave similarly. They need fewer samples and give a better approximation than pessimistic sampling. This might be because pessimistic sampling is too conservative in reacting to chances of pruning.

For different k experimental results show similar behaviour. Please refer to Appendix D in (George and Dimitrakakis 2023) for more results regarding different k and a comparison to preference profiles generated from a Mallows model (Mallows 1957) with $\phi = 0.2$ (i.e., preferences being fairly similar) or uniformly at random generated single-peaked preferences which always admit a Condorcet winner.

6 Conclusion

We phrased the problem of eliciting voters’ preferences for determining a Kemeny ranking in terms of a Dueling Bandit problem and found that while there are no efficient PAC algorithms for approximating the Kendall tau distance, we can formulate PAC algorithms for approximating the Kemeny score. Here, our approximation bounds for the Kemeny score are dependent on the sum of confidence bounds for the approximated winning probabilities of arms, rendering uniform sampling the best sampling strategy w.r.t. these bounds. We analyse the sample complexity for uniform sampling of arms for sampling with and without replacement. Further, we described ways of pruning confidence bounds in order to reduce the sample complexity, yielding non-uniform sample strategies, and demonstrate the effect of our methods experimentally.

In future work, one could transfer the methodology taken in this paper for finding confidence-bound-based approximation bounds for other C2 voting rules. Furthermore, the Kemeny score could be an interesting approximation measure for PAC algorithms under different voting rules.

Acknowledgements

This work was supported under the NRC Grant No 302203 “Algorithms and Models for Socially Beneficial Artificial Intelligence”.

References

- Arrow, K. J. 1950. A difficulty in the concept of social welfare. *Journal of political economy*, 58(4): 328–346.
- Balsubramani, A.; Karnin, Z.; Schapire, R. E.; and Zoghi, M. 2016. Instance-dependent regret bounds for dueling bandits. In *Conference on Learning Theory*, 336–360. PMLR.
- Bartholdi, J.; Tovey, C. A.; and Trick, M. A. 1989. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and welfare*, 6: 157–165.
- Bengs, V.; Busa-Fekete, R.; El Mesaoudi-Paul, A.; and Hüllermeier, E. 2021. Preference-based online learning with dueling bandits: A survey. *The Journal of Machine Learning Research*, 22(1): 278–385.
- Betzler, N.; Fellows, M. R.; Guo, J.; Niedermeier, R.; and Rosamond, F. A. 2008. Fixed-parameter algorithms for Kemeny scores. In *Algorithmic Aspects in Information and Management: 4th International Conference, AAIM 2008, Shanghai, China, June 23-25, 2008. Proceedings 4*, 60–71. Springer.
- Boehmer, N.; Faliszewski, P.; Niedermeier, R.; Szufa, S.; and Was, T. 2022. Understanding distance measures among elections. *arXiv preprint arXiv:2205.00492*.
- Brandt, F.; Brill, M.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2015. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research*, 53: 439–496.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D. 2016. *Handbook of computational social choice*. Cambridge University Press.
- Busa-Fekete, R.; Szörényi, B.; and Hüllermeier, E. 2014. PAC rank elicitation through adaptive sampling of stochastic pairwise preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, #1.
- Falahatgar, M.; Orlitsky, A.; Pichapati, V.; and Suresh, A. T. 2017. Maximum selection and ranking under noisy comparisons. In *International Conference on Machine Learning*, 1088–1096. PMLR.
- Fishburn, P. C. 1977. Condorcet social choice functions. *SIAM Journal on applied Mathematics*, 33(3): 469–489.
- George, A.-M.; and Dimitrakakis, C. 2023. Eliciting Kemeny Rankings. [arXiv:2312.11663](https://arxiv.org/abs/2312.11663).
- Kemeny, J. G. 1959. Mathematics without numbers. *Daedalus*, 88(4): 577–591.
- Kendall, M. G. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2): 81–93.
- Lin, C.-C.; and Lu, C.-J. 2018. Efficient mechanisms for peer grading and dueling bandits. In *Asian Conference on Machine Learning*, 740–755. PMLR.
- Mallows, C. L. 1957. Non-Null Ranking Models. I. *Biometrika*, 44(1/2): 114–130.
- Satterthwaite, M. A. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10(2): 187–217.
- Young, H. P.; and Levenglick, A. 1978. A consistent extension of Condorcet’s election principle. *SIAM Journal on applied Mathematics*, 35(2): 285–300.