
Maximum Selection and Ranking under Noisy Comparisons

Moein Falahatgar¹ Alon Orlitsky¹ Venkatadheeraj Pichapati¹ Ananda Theertha Suresh²

Abstract

We consider (ϵ, δ) -PAC maximum-selection and ranking using pairwise comparisons for general probabilistic models whose comparison probabilities satisfy strong stochastic transitivity and stochastic triangle inequality. Modifying the popular knockout tournament, we propose a simple maximum-selection algorithm that uses $\mathcal{O}\left(\frac{n}{\epsilon^2} \left(1 + \log \frac{1}{\delta}\right)\right)$ comparisons, optimal up to a constant factor. We then derive a general framework that uses noisy binary search to speed up many ranking algorithms, and combine it with merge sort to obtain a ranking algorithm that uses $\mathcal{O}\left(\frac{n}{\epsilon^2} \log n (\log \log n)^3\right)$ comparisons for $\delta = \frac{1}{n}$, optimal up to a $(\log \log n)^3$ factor.

1. Introduction

1.1. Background

Maximum selection and sorting using pairwise comparisons are computer-science staples taught in most introductory classes and used in many applications. In fact, sorting, also known as *ranking*, was once claimed to utilize 25% of all computer cycles, e.g., (Mukherjee, 2011).

In many applications, the pairwise comparisons produce only random outcomes. In sports, tournaments rank teams based on pairwise matches whose outcomes are probabilistic in nature. For example, Microsoft’s *TrueSkill* (Herbrich et al., 2006) software matches and ranks thousands of Xbox gamers based on individual game results. And in online advertising, out of a myriad of possible ads, each web page may display only a few, and a user will typically select at most one. Based on these random comparisons, ad companies such as Google, Microsoft, or Yahoo, rank the ads’ appeal (Radlinski & Joachims, 2007; Radlinski et al., 2008).

These and related applications have brought about a resur-

gence of interest in maximum selection and ranking using noisy comparisons. Several probabilistic models were considered, including the popular Bradley-Terry-Luce (Bradley & Terry, 1952) and its Plackett-Luce (PL) generalization (Plackett, 1975; Luce, 2005). Yet even for such specific models, the number of pairwise comparisons needed, or *sample complexity*, of maximum selection and ranking was known only to within a $\log n$ factor. We consider a significantly broader class of models and yet propose algorithms that are optimal up to a constant factor for maximum selection and up to $(\log \log n)^3$ for ranking.

1.2. Notation

Noiseless comparison assumes an unknown underlying ranking $r(1), \dots, r(n)$ of the elements in $\{1, \dots, n\}$ such that if two elements are compared, the higher-ranked one is selected. Similarly for noisy comparisons, we assume an unknown ranking of the elements, but now if two elements i and j are compared, i is chosen with some unknown probability $p(i, j)$ and j is chosen with probability $p(j, i) = 1 - p(i, j)$, where if i is higher-ranked, then $p(i, j) \geq \frac{1}{2}$. Repeated comparisons are independent of each other.

Let $\tilde{p}(i, j) = p(i, j) - \frac{1}{2}$ reflect the *additional probability* by which i is preferable to j . Note that $\tilde{p}(j, i) = -\tilde{p}(i, j)$ and $\tilde{p}(i, j) \geq 0$ if $r(i) > r(j)$. $|\tilde{p}(i, j)|$ can also be seen as a measure of dissimilarity between i and j . Following (Yue & Joachims, 2011), we assume that two natural properties, satisfied for example by the PL model, hold whenever $r(i) > r(j) > r(k)$: *Strong Stochastic Transitivity (SST)*, $\tilde{p}(i, k) \geq \max(\tilde{p}(i, j), \tilde{p}(j, k))$, and *Stochastic Triangle Inequality (STI)*, $\tilde{p}(i, k) \leq \tilde{p}(i, j) + \tilde{p}(j, k)$.

Two types of algorithms have been proposed for maximum selection and ranking under noisy comparisons: non-adaptive or offline (Rajkumar & Agarwal, 2014; Negahban et al., 2012; 2016; Jang et al., 2016) where the comparison pairs are chosen in advance, and *adaptive* or *online* where the comparison pairs are selected sequentially based on previous comparison results. We focus on the latter.

We specify the desired output via the (ϵ, δ) -PAC paradigm (Yue & Joachims, 2011; Szörényi et al., 2015) that requires the output to likely closely approximate the intended outcome. Specifically, given $\epsilon, \delta > 0$, with prob-

¹University of California, San Diego ²Google Research. Correspondence to: Venkatadheeraj Pichapati <dheera-jpv7@ucsd.edu>.

ability $\geq 1 - \delta$, maximum selection must output an ϵ -maximum element i such that for all j , $p(i, j) \geq \frac{1}{2} - \epsilon$. Similarly, with probability $\geq 1 - \delta$, the ranking algorithm must output an ϵ -ranking $r'(1), \dots, r'(n)$ such that whenever $r'(i) > r'(j)$, $p(i, j) \geq \frac{1}{2} - \epsilon$.

1.3. Outline

In Section 2 we review past work and summarize our contributions. In Section 3 we describe and analyze our maximum-selection algorithm. In Section 4 we propose and evaluate the ranking algorithm. In Section 5 we experimentally compare our algorithms with existing ones. In Section 6 we mention some future directions.

2. Old and new results

2.1. Related work

Several researchers studied algorithms that with probability $1 - \delta$ find the exact maximum and ranking. (Feige et al., 1994) considered a simple model where the elements are ranked, and $\tilde{p}(i, j) = \epsilon$ whenever $r(i) > r(j)$. (Busa-Fekete et al., 2014a) considered comparison probabilities $p(i, j)$ satisfying the Mallows model (Mallows, 1957). And (Urvoy et al., 2013; Busa-Fekete et al., 2014b; Heckel et al., 2016) considered general comparison probabilities, without an underlying ranking assumption, and derived rankings based on Copeland- and Borda-counts, and random-walk procedures. As expected, when the comparison probabilities approach half, the above algorithms require arbitrarily many comparisons.

To achieve finite complexity even with near-half comparison probabilities, researchers adopted the PAC paradigm. For the PAC model with SST and STI constraints, (Yue & Joachims, 2011) derived a maximum-selection algorithm with sample complexity $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{n}{\epsilon\delta})$ and used it to bound the regret of the problem's dueling-bandits variant. Related results appeared in (Syrngkanis et al., 2016). For the PL model, (Szörényi et al., 2015) derived a PAC ranking algorithm with sample complexity $\mathcal{O}(\frac{n}{\epsilon^2} \log n \log \frac{n}{\epsilon\delta})$.

Deterministic adversarial versions of the problem were considered by (Ajtai et al., 2015), and by (Acharya et al., 2014a; 2016) who were motivated by density estimation (Acharya et al., 2014b).

2.2. New results

We consider (ϵ, δ) -PAC adaptive maximum selection and ranking using pairwise comparisons under SST and STI constraints. Note that when $\epsilon \geq \frac{1}{2}$ or $\delta \geq 1 - 1/n$ for maximum selection and $\delta \geq 1 - 1/n^2$ for ranking, any output is correct. We show for $\epsilon < 1/4$, $\delta < \frac{1}{2}$ and any n :

- Maximum-selection algorithm with sample complexity $\mathcal{O}(\frac{n}{\epsilon^2} (1 + \log \frac{1}{\delta}))$, optimal up to a constant factor.
- Ranking algorithm with $\mathcal{O}(\frac{n}{\epsilon^2} (\log n)^3 \log \frac{n}{\delta})$ sample complexity.
- General framework that converts any ranking algorithm with sample complexity $\mathcal{O}(\frac{n}{\epsilon^2} (\log n)^x \log \frac{n}{\delta})$ into a ranking algorithm that for $\delta \geq \frac{1}{n}$ has sample complexity $\mathcal{O}(\frac{n}{\epsilon^2} \log n (\log \log n)^x)$.
- Using the above framework, a ranking algorithm with sample complexity $\mathcal{O}(\frac{n}{\epsilon^2} \log n (\log \log n)^3)$ for $\delta = \frac{1}{n}$.
- An $\Omega(\frac{n}{\epsilon^2} \log \frac{n}{\delta})$ lower bound on the sample complexity of any PAC ranking algorithm, matching our algorithm's sample complexity up to a $(\log \log n)^3$ factor.

3. Maximum selection

3.1. Algorithm outline

We propose a simple maximum-selection algorithm based on Knockout tournaments. Knockout tournaments are used to find a maximum element under non-noisy comparisons. Knockout tournament of n elements runs in $\lceil \log n \rceil$ rounds where in each round it randomly pairs the remaining elements and proceeds the winners to next round.

Our algorithm, given in `KNOCKOUT` uses $\mathcal{O}(\frac{n}{\epsilon^2} (1 + \log \frac{1}{\delta}))$ comparisons and $\mathcal{O}(n)$ memory to find an ϵ -maximum. (Yue & Joachims, 2011) uses $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{n}{\epsilon\delta})$ comparisons and $\mathcal{O}(n^2)$ memory to find an ϵ -maximum. Hence we get $\log n$ -factor improvement in the number of comparisons and also we use linear memory compared to quadratic memory. From (Zhou & Chen, 2014) it can be inferred that the best PAC maximum selection algorithm requires $\Omega(\frac{n}{\epsilon^2} (1 + \log \frac{1}{\delta}))$ comparisons, hence up to constant factor, `KNOCKOUT` is optimal.

(Yue & Joachims, 2011; Szörényi et al., 2015) eliminate elements one by one until only ϵ -maximums are remaining. Since they potentially need $n - 1$ eliminations, in order to apply union bound they had to ensure that each eliminated element is not an ϵ -maximum w.p. $1 - \delta/n$, requiring $\mathcal{O}(\log(n/\delta))$ comparisons for each eliminated element and hence a superlinear sample complexity $\mathcal{O}(n \log(n/\delta))$.

In contrast, `KNOCKOUT` eliminates elements in $\log n$ rounds. Since in Knockout tournaments, number of elements decrease exponentially with each round, we afford to endure more error in the initial rounds and less error in the latter rounds by repeating comparison between each pair more times in latter rounds. Specifically, let b_i be the highest-ranked element (according to the unobserved underlying ranking) at the beginning of round i . `KNOCKOUT` makes sure that w.p. $\geq 1 - \frac{\delta}{2^i}$, $\tilde{p}(b_i, b_{i+1}) \leq \epsilon_i$ by repeating

comparison between each pair in round i for $\mathcal{O}\left(\frac{1}{\epsilon_i^2} \log \frac{2^i}{\delta}\right)$ times. Choosing $\epsilon_i = \frac{c\epsilon}{2^{i/3}}$ with $c = 2^{1/3} - 1$, we make sure that comparison complexity is $\mathcal{O}\left(\frac{n}{\epsilon^2} (1 + \log \frac{1}{\delta})\right)$ and by union bound and STI, w.p. $\geq 1 - \delta$, $\tilde{p}(b_1, b_{\lceil \log n \rceil + 1}) \leq \sum_{i=1}^{\lceil \log n \rceil + 1} \frac{c\epsilon}{2^{i/3}} \leq \epsilon$.

For $\gamma \geq 1$, a relaxed notion of SST, called γ -stochastic transitivity (Yue & Joachims, 2011), requires that if $r(i) > r(j) > r(k)$, then $\max(\tilde{p}(i, j), \tilde{p}(j, k)) \leq \gamma \cdot \tilde{p}(i, k)$. Our results apply to this general notion of γ -stochastic transitivity and the analysis of KNOCKOUT is presented under this model. KNOCKOUT uses $\mathcal{O}\left(\frac{n\gamma^4}{\epsilon^2} (1 + \log \frac{1}{\delta})\right)$ comparisons.

Remark 1. (Yue & Joachims, 2011) considered a different definition of ϵ -maximum as an element i that is at most ϵ dissimilar to true maximum i.e., for j with $r(j) = n$, $\tilde{p}(j, i) \leq \epsilon$. Note that this definition is less restrictive than ours, hence requires fewer comparisons. Under this definition, (Yue & Joachims, 2011) used $\mathcal{O}\left(\frac{n\gamma^6}{\epsilon^2} \log \frac{n}{\epsilon\delta}\right)$ comparisons to find an ϵ -maximum whereas a simple modification of KNOCKOUT shows that $\mathcal{O}\left(\frac{n\gamma^2}{\epsilon^2} (1 + \log \frac{1}{\delta})\right)$ comparisons suffice. Hence we also get a significant improvement in the exponent of γ .

To simplify the analysis, we assume that n is a power of 2, otherwise we can add $2^{\lceil \log n \rceil} - n$ dummy elements that lose to every original element with probability 1. Note that all ϵ -maximums will still be from the original set.

3.2. Algorithm

We start with a subroutine COMPARE that compares two elements. It compares two elements i, j and maintains empirical probability \hat{p}_i , a proxy for $p(i, j)$. It also maintains a confidence value \hat{c} s.t., w.h.p., $\hat{p}_i \in (p(i, j) - \hat{c}, p(i, j) + \hat{c})$. COMPARE stops if it is confident about the winner or if it reaches its comparison budget m . It outputs the element with more wins breaking ties randomly.

Algorithm 1 COMPARE

Input: element i , element j , bias ϵ , confidence δ .

Initialize: $\hat{p}_i = \frac{1}{2}$, $\hat{c} = \frac{1}{2}$, $m = \frac{1}{2\epsilon^2} \log \frac{2}{\delta}$, $r = 0$, $w_i = 0$.

1. **while** $(|\hat{p}_i - \frac{1}{2}| \leq \hat{c} - \epsilon$ and $r \leq m)$
 - (a) Compare i and j . **if** i wins $w_i = w_i + 1$.
 - (b) $r = r + 1$, $\hat{p}_i = \frac{w_i}{r}$, $\hat{c} = \sqrt{\frac{1}{2r} \log \frac{4r^2}{\delta}}$.

if $\hat{p}_i \leq \frac{1}{2}$ **Output:** j . **else Output:** i .

We show that COMPARE w.h.p., outputs the correct winner if the elements are well separated.

Lemma 2. If $\tilde{p}(i, j) \geq \epsilon$, then

$$Pr(\text{COMPARE}(i, j, \epsilon, \delta) \neq i) \leq \delta.$$

Note that instead of using fixed number of comparisons, COMPARE stops the comparisons adaptively if it is confident about the winner. If $|\tilde{p}(i, j)| \gg \epsilon$, COMPARE stops much before comparison budget $\frac{1}{2\epsilon^2} \log \frac{2}{\delta}$ and hence works better in practice.

Now we present the subroutine KNOCKOUT-ROUND that we use in main algorithm KNOCKOUT.

3.2.1. KNOCKOUT-ROUND

KNOCKOUT-ROUND takes a set S and outputs a set of size $|S|/2$. It randomly pairs elements, compares each pair using COMPARE, and returns the set of winners. We will later show that maximum element in the output set will be comparable to maximum element in the input set.

Algorithm 2 KNOCKOUT-ROUND

Input: Set S , bias ϵ , confidence δ .

Initialize: Set $O = \emptyset$.

1. Pair elements in S randomly.
2. **for** every pair (i, j) :
Add COMPARE(i, j, ϵ, δ) to O .

Output: O

Note that comparisons between each pair can be handled by a different processor and hence this algorithm can be easily parallelized.

S can have several maximum elements. Comparison probabilities corresponding to all maximum elements will be essentially same because of STI. We define $\max(S)$ to be the maximum element with the least index, namely,

$$\max(S) \stackrel{\text{def}}{=} S\left(\min\{i : \tilde{p}(S(i), S(j)) \geq 0 \quad \forall j\}\right).$$

Lemma 3. KNOCKOUT-ROUND(S, ϵ, δ) uses $\frac{|S|}{4\epsilon^2} \log \frac{2}{\delta}$ comparisons and with probability $\geq 1 - \delta$,

$$\tilde{p}\left(\max(S), \max\left(\text{KNOCKOUT-ROUND}(S, \epsilon, \delta)\right)\right) \leq \gamma\epsilon.$$

3.2.2. KNOCKOUT

Now we present the main algorithm KNOCKOUT. KNOCKOUT takes an input set S and runs $\log n$ rounds of KNOCKOUT-ROUND halving the size of S at the end of each round. Recall that KNOCKOUT-ROUND makes sure that maximum element in the output set is comparable to

maximum element in the input set. Using this, KNOCKOUT makes sure that the output element is comparable to maximum element in the input set.

Since the size of S gets halved after each round, KNOCKOUT compares each pair more times in the latter rounds. Hence the bias between maximum element in input set and maximum element in output set is small in latter rounds.

Algorithm 3 KNOCKOUT

Input: Set S , bias ϵ , confidence δ , stochasticity γ .

Initialize: $i = 1$, $S =$ set of all elements, $c = 2^{1/3} - 1$.

while $|S| > 1$

1. $S = \text{KNOCKOUT-ROUND}\left(S, \frac{c\epsilon}{\gamma 2^{i/3}}, \frac{\delta}{2^i}\right)$.
2. $i = i + 1$.

Output: the unique element in S .

Note that KNOCKOUT uses only memory of set S and hence $\mathcal{O}(n)$ memory suffices.

Theorem 4 shows that KNOCKOUT outputs an ϵ -maximum with probability $\geq 1 - \delta$. It also bounds the number of comparisons used by the algorithm.

Theorem 4. KNOCKOUT(S, ϵ, δ) uses $\mathcal{O}\left(\frac{\gamma^4 |S|}{\epsilon^2} (1 + \log \frac{1}{\delta})\right)$ comparisons and with probability at least $1 - \delta$, outputs an ϵ -maximum.

4. Ranking

We propose a ranking algorithm that with probability at least $1 - \frac{1}{n}$ uses $\mathcal{O}\left(\frac{n \log n (\log \log n)^3}{\epsilon^2}\right)$ comparisons and outputs an ϵ -ranking.

Notice that we use only $\tilde{\mathcal{O}}\left(\frac{n \log n}{\epsilon^2}\right)$ comparisons for $\delta = \frac{1}{n}$ where as (Szörényi et al., 2015) uses $\mathcal{O}(n(\log n)^2/\epsilon^2)$ comparisons even for constant error probability δ . Furthermore (Szörényi et al., 2015) provided these guarantees only under Plackett-Luce model which is more restrictive compared to ours. Also, their algorithm uses $\mathcal{O}(n^2)$ memory compared to $\mathcal{O}(n)$ memory requirement of ours.

Our main algorithm BINARY-SEARCH-RANKING assumes the existence of a ranking algorithm RANK- x that with probability at least $1 - \delta$ uses $\mathcal{O}\left(\frac{n}{\epsilon^2} (\log n)^x \log \frac{n}{\delta}\right)$ comparisons and outputs an ϵ -ranking for any $\delta > 0$, $\epsilon > 0$ and some $x > 1$. We also present a RANK- x algorithm with $x = 3$.

Observe that we need RANK- x algorithm to work for any model that satisfies SST and STI. (Szörényi et al., 2015) showed that their algorithm works for Plackett-Luce model but not for more general model. So we present a RANK- x

algorithm that works for general model.

The main algorithm BINARY-SEARCH-RANKING randomly selects $\frac{n}{(\log n)^x}$ elements (anchors) and rank them using RANK- x . The algorithm has then effectively created $\frac{n}{(\log n)^x}$ bins, each between two successively ranked anchors. Then for each element, the algorithm identifies the bin it belongs to using a noisy binary search algorithm. The algorithm then ranks the elements within each bin using RANK- x .

We first present MERGE-RANK, a RANK-3 algorithm.

4.1. Merge Ranking

We present a simple ranking algorithm MERGE-RANK that uses $\mathcal{O}\left(\frac{n(\log n)^3}{\epsilon^2} \log \frac{n}{\delta}\right)$ comparisons, $\mathcal{O}(n)$ memory and with probability $\geq 1 - \delta$ outputs an ϵ -ranking. Thus MERGE-RANK is a RANK- x algorithm for $x = 3$.

Similar to Merge Sort, MERGE-RANK divides the elements into two sets of equal size, ranks them separately and combines the sorted sets. Due to the noisy nature of comparisons, MERGE-RANK compares two elements i, j sufficient times, so that the comparison output is correct with high probability when $|\tilde{p}(i, j)| \geq \frac{\epsilon}{\log n}$. Put differently, MERGE-RANK is same as the typical Merge Sort, except it uses COMPARE as the comparison function. Due to lack of space, MERGE-RANK is presented in Appendix A.

Let's define the error of an ordered set S as the maximum distance between two wrongly ordered items in S , namely,

$$\text{err}(S) \stackrel{\text{def}}{=} \max_{1 \leq i \leq j \leq |S|} \tilde{p}(S(i), S(j)).$$

We show that when we merge two ordered sets, the error of the resulting ordered set will be at most $\frac{\epsilon}{\log n}$ more than the maximum of errors of individual ordered sets.

Observe that MERGE-RANK is a recursive algorithm and the error of a singleton set is 0. Two singleton sets each containing a unique element from the input set merge to form a set with two elements with an error at most $\frac{2\epsilon}{\log n}$, then two sets with two elements merge to form a set with four elements with an error of at most $\frac{3\epsilon}{\log n}$ and henceforth. Thus the error of the output ordered set is bounded by ϵ .

Lemma 5 shows that MERGE-RANK can output an ϵ -ranking of S with probability $\geq 1 - \delta$. It also bounds the number of comparisons used by the algorithm.

Lemma 5. MERGE-RANK $\left(S, \frac{\epsilon}{\log |S|}, \frac{\delta}{|S|^2}\right)$ takes $\mathcal{O}\left(\frac{|S|(\log |S|)^3}{\epsilon^2} \log \frac{|S|}{\delta}\right)$ comparisons and with probability $\geq 1 - \delta$, outputs an ϵ -ranking. Hence, MERGE-RANK is a RANK-3 algorithm.

Now we present our main ranking algorithm.

4.2. BINARY-SEARCH-RANKING

We first sketch the algorithm outline below. We then provide a proof outline.

4.2.1. ALGORITHM OUTLINE

Our algorithm is stated in BINARY-SEARCH-RANKING. It can be summarized in three major parts.

Creating anchors: (Steps 1 to 3) BINARY-SEARCH-RANKING first selects a set S' of $\frac{n}{(\log n)^x}$ random elements (anchors) and ranks them using RANK- x . At the end of this part, there are $\frac{n}{(\log n)^x}$ ranked anchors. Equivalently, the algorithm creates $\frac{n}{(\log n)^x} - 1$ bins, each bin between two successively ranked anchors.

Coarse ranking: (Step 4) After forming the bins, the algorithm uses a random walk on a binary search tree, to find which bin each element belongs to. INTERVAL-BINARY-SEARCH is similar to the noisy binary search algorithm in (Feige et al., 1994). It builds a binary search tree with the bins as the leaves and it does a random walk over this tree. Due to lack of space the algorithm INTERVAL-BINARY-SEARCH is presented in Appendix B but more intuition is given later in this section.

Ranking within each bin: (Step 5) For each bin, we show that the number of elements far from both anchors is bounded. The algorithm checks elements inside a bin whether they are close to any of the bin's anchors. For the elements that are close to anchors, the algorithm ranks them close to the anchor. And for the elements that are away from both anchors the algorithm ranks them using RANK- x and outputs the resulting ranking.

4.2.2. ANALYSIS OF BINARY-SEARCH-RANKING

Creating anchors In Step 1 of the algorithm we select $n/(\log n)^x$ random elements. Since these are chosen uniformly random, they lie nearly uniformly in the set S . This intuition is formalized in the next lemma.

Lemma 6. *Consider a set S of n elements. If we select $\frac{n}{(\log n)^x}$ elements uniformly randomly from S and build an ordered set S' s.t. $\tilde{p}(S'(i), S'(j)) \geq 0 \forall i > j$, then with probability $\geq 1 - \frac{1}{n^4}$, for any $\epsilon > 0$ and all k ,*

$$|\{e \in S : \tilde{p}(e, S'(k)) > \epsilon, \tilde{p}(S'(k+1), e) > \epsilon\}| \leq 5(\log n)^{x+1}$$

In Step 2, we use RANK- x to rank S' . Lemma 7 shows the guarantee of ranking S' .

Lemma 7. *After Step 2 of the BINARY-SEARCH-RANKING with probability $\geq 1 - \frac{1}{n^6}$, S' is ϵ' -ranked.*

At the end of Step 2, we have $\frac{n}{(\log n)^x} - 1$ bins, each between two successively ranked anchors. Each bin has a left

Algorithm 4 BINARY-SEARCH-RANKING

Input: Set S , bias ϵ .

Initialize: $\epsilon' = \epsilon/16$, $\epsilon'' = \epsilon/15$, and $S^o = \emptyset$. $S_j = \emptyset$, $C_j = \emptyset$ and $B_j = \emptyset$, for $1 \leq j \leq \lfloor \frac{n}{(\log n)^x} \rfloor + 2$.

1. Form a set S' with $\lfloor \frac{n}{(\log n)^x} \rfloor$ random elements from S . Remove these elements from S .
2. Rank S' using RANK- x ($S', \epsilon', \frac{1}{n^6}$).
3. Add dummy element a at the beginning of S' such that $p(a, e) = 0 \forall e \in S \cup S'$. Add dummy element b at the end of S' such that $p(b, e) = 1 \forall e \in S \cup S'$.
4. **for** $e \in S$:
 - (a) $k = \text{INTERVAL-BINARY-SEARCH}(S', e, \epsilon'')$.
 - (b) Insert e in S_k .
5. **for** $j = 1$ to $\lfloor \frac{n}{(\log n)^x} \rfloor + 2$:
 - (a) **for** $e \in S_j$:
 - i. **if** $\text{COMPARE2}(e, S'(j), 10\epsilon''^{-2} \log n) \in [\frac{1}{2} - 6\epsilon'', \frac{1}{2} + 6\epsilon'']$, insert e in C_j .
 - ii. **else if** $\text{COMPARE2}(e, S'(j+1), 10\epsilon''^{-2} \log n) \in [\frac{1}{2} - 6\epsilon'', \frac{1}{2} + 6\epsilon'']$, then insert e in C_{j+1} .
 - iii. **else** insert e in B_j .
 - (b) Rank B_j using RANK- x ($B_j, \epsilon'', \frac{1}{n^4}$).
 - (c) Append $S'(j), C_j, B_j$ in order at the end of S^o .

Output: S^o

anchor and a right anchor. We say that an element belongs to a bin if it wins over the bin's left anchor with probability $\geq \frac{1}{2}$ and wins over the bin's right anchor with probability $\leq \frac{1}{2}$. Notice that some elements might win over $S'(1)$ with probability $< \frac{1}{2}$ and thus not belong to any bin. So in Step 3, we add a dummy element a at the beginning of S' where a loses to every element in $S \cup S'$ with probability 1. For similar reasons we add a dummy element b to the end of S' where every element in $S \cup S'$ loses to b with probability 1.

Coarse Ranking Note that $S'(i)$ and $S'(i+1)$ are respectively the left and right anchors of the bin S_i .

Algorithm 5 COMPARE2

Input: element i , element j , number of comparisons m .

1. Compare i and j for m times and return the fraction of times i wins over j .
-

Since S' is ϵ' -ranked and the comparisons are noisy, it is hard to find a bin S_i for an element e such that $p(e, S'(i)) \geq \frac{1}{2}$ and $p(S'(i+1), e) \geq \frac{1}{2}$. We call a bin S_i a ϵ'' -*nearly correct* bin for an element e if $p(e, S'(i)) \geq \frac{1}{2} - \epsilon''$ and $p(S'(i+1), e) \geq \frac{1}{2} - \epsilon''$ for some $\epsilon'' > \epsilon'$.

In Step 4, for each element we find an ϵ'' -*nearly correct* bin using INTERVAL-BINARY-SEARCH. Next we describe an outline of INTERVAL-BINARY-SEARCH.

INTERVAL-BINARY-SEARCH first builds a binary search tree of intervals (see Appendix B) as follows: the root node is the entire interval between the first and the last elements in S' . Each non-leaf node interval I has two children corresponding to the left and right halves of I . The leaves of the tree are the bins between two successively ranked anchors.

To find an ϵ'' -*nearly correct* bin for an element e , the algorithm starts at the root of the binary search tree and at every non-leaf node corresponding to interval I , it checks if e belongs to I or not by comparing e with I 's left and right anchors. If e loses to left anchor or wins against the right anchor, the algorithm backtracks to current node's parent.

If e wins against I 's left anchor and loses to its right one, the algorithm checks if e belongs to the left or right child by comparing e with the middle element of I and moves accordingly.

When at a leaf node, the algorithm checks if e belongs to the bin by maintaining a counter. If e wins against the bin's left anchor and loses to the bin's right anchor, it increases the counter by one or otherwise it decreases the counter by one. If the counter is less than 0 the algorithm backtracks to the bin's parent. By repeating each comparison several times, the algorithm makes a correct decision with probability $\geq \frac{19}{20}$.

Note that there could be several ϵ'' -*nearly correct* bins for e and even though at each step the algorithm moves in the direction of one of them, it could end up moving in a loop and never reaching one of them. We thus run the algorithm for $30 \log n$ steps and terminate.

If the algorithm is at a leaf node by $30 \log n$ steps and the counter is more than $10 \log n$ we show that the leaf node bin is a ϵ'' -*nearly correct* bin for e and the algorithm outputs the leaf node. If not, the algorithm puts in a set Q all the anchors visited so far and orders Q according to S' .

We select $30 \log n$ steps to ensure that if there is only one nearly correct bin, then the algorithm outputs that bin w.p. $\geq 1 - \frac{1}{n^6}$. Also we do not want too many steps so as to bound the size of Q .

By doing a simple binary search in Q using BINARY-SEARCH (see Appendix B) we find an anchor $f \in Q$ such that $|\tilde{p}(e, f)| \leq 4\epsilon''$. Since INTERVAL-BINARY-

SEARCH ran for at most $30 \log n$ steps, Q can have at most $60 \log n$ elements and hence BINARY-SEARCH can search effectively by repeating each comparison $\mathcal{O}(\log n)$ times to maintain high confidence. Next paragraph explains how BINARY-SEARCH finds such an element f .

BINARY-SEARCH first compares e with the middle element m of Q for $\mathcal{O}(\log n)$ times. If the fraction of wins for e is between $\frac{1}{2} - 3\epsilon''$ and $\frac{1}{2} + 3\epsilon''$, then w.h.p. $|\tilde{p}(e, m)| \leq 4\epsilon''$ and hence BINARY-SEARCH outputs m . If the fraction of wins for e is less than $\frac{1}{2} - 3\epsilon''$, then w.h.p. $\tilde{p}(e, m) \leq -2\epsilon''$ and hence it eliminates all elements to the right of m in Q . If the fraction of wins for e is more than $\frac{1}{2} + 3\epsilon''$, then w.h.p. $\tilde{p}(e, m) \geq 2\epsilon''$ and hence it eliminates all elements to the left of m in Q . It continues this process until it finds an element f such that the fraction of wins for e is between $\frac{1}{2} - 3\epsilon''$ and $\frac{1}{2} + 3\epsilon''$.

In next Lemma, we show that INTERVAL-BINARY-SEARCH achieves to find a $5\epsilon''$ -*nearly correct* bin for every element.

Lemma 8. *For any element $e \in S$, Step 4 of BINARY-SEARCH-RANKING places e in bin S_l such that $\tilde{p}(e, S'(l)) > -5\epsilon''$ and $\tilde{p}(S'(l+1), e) > -5\epsilon''$ with probability $\geq 1 - \frac{1}{n^5}$.*

Ranking within each bin Once we have identified the bins, we rank the elements inside each bin. By Lemma 6, inside each bin all elements are close to the bin's anchors except at most $5(\log n)^{x+1}$ of them.

The algorithm finds the elements close to anchors in Step 5a by comparing each element in the bin with the bin's anchors. If an element in bin S_j is close to bin's anchors $S'(j)$ or $S'(j+1)$, the algorithm moves it to the set C_j or C_{j+1} accordingly and if it is far away from both, the algorithm moves it to the set B_j . The following two lemmas state that this separating process happens accurately with high probability. The proofs of these results follow from the Chernoff bound and hence omitted.

Lemma 9. *At the end of Step 5a, for all j , $\forall e \in C_j$, $|\tilde{p}(e, S'(j))| < 7\epsilon''$ with probability $\geq 1 - \frac{1}{n^3}$.*

Lemma 10. *At the end of Step 5a, for all j , $\forall e \in B_j$, $\min(\tilde{p}(e, S'(j)), \tilde{p}(S'(j+1), e)) > 5\epsilon''$ with probability $\geq 1 - \frac{1}{n^3}$.*

Combining Lemmas 6, 7 and 10 next lemma shows that the size of B_j is bounded for all j .

Lemma 11. *At the end of Step 5a, $|B_j| \leq 5(\log n)^{x+1}$ for all j , with probability $\geq 1 - \frac{3}{n^3}$.*

Since all the elements in C_j are already close to an anchor, they need not be ranked. By Lemma 11 with probability $\geq 1 - \frac{3}{n^3}$ the number of elements in B_j is at most $5(\log n)^{x+1}$. We use RANK- x to rank each B_j and output the final ranking.

Lemma 12 shows that all B_j 's are ϵ'' -ranked at the end of Step 5b. Proof follows from properties of RANK- x and union bound.

Lemma 12. *At the end of Step 5b, all B_j s are ϵ'' -ranked with probability $\geq 1 - \frac{1}{n^3}$.*

Combining the above set of results yields our main result.

Theorem 13. *Given access to RANK- x , BINARY-SEARCH-RANKING with probability $\geq 1 - \frac{1}{n}$, uses $O\left(\frac{n \log n (\log \log n)^x}{\epsilon^2}\right)$ comparisons and outputs an ϵ -ranking.*

Using MERGE-RANK as a RANK- x algorithm with $x = 3$ leads to the following corollary.

Corollary 14. *BINARY-SEARCH-RANKING uses $O\left(\frac{n \log n (\log \log n)^3}{\epsilon^2}\right)$ comparisons and outputs an ϵ -ranking with probability $\geq 1 - \frac{1}{n}$.*

Using PALPAC-AMPRR (Szörényi et al., 2015) as a RANK- x algorithm with $x = 1$ leads to the following corollary over PL model.

Corollary 15. *Over PL model, BINARY-SEARCH-RANKING with probability $\geq 1 - \frac{1}{n}$ uses $O\left(\frac{n \log n \log \log n}{\epsilon^2}\right)$ comparisons and outputs an ϵ -ranking.*

It is well known that to rank a set of n values under the noiseless setting, $\Omega(n \log n)$ comparisons are necessary. We show that under the noisy model, $\Omega\left(\frac{n}{\epsilon^2} \log \frac{n}{\delta}\right)$ samples are necessary to output an ϵ -ranking and hence our algorithm is near-optimal.

Theorem 16. *For $\epsilon \leq \frac{1}{4}$, $\delta \leq \frac{1}{2}$, there exists a noisy model that satisfies SST and STI such that to output an ϵ -ranking with probability $\geq 1 - \delta$, $\Omega\left(\frac{n}{\epsilon^2} \log \frac{n}{\delta}\right)$ comparisons are necessary.*

5. Experiments

We compare the performance of our algorithms with that of others over simulated data. Similar to (Yue & Joachims, 2011), we consider the stochastic model where $p(i, j) = 0.6 \forall i < j$. Note that this model satisfies both SST and STI. We find 0.05-maximum with error probability $\delta = 0.1$. Observe that $i = 1$ is the only 0.05-maximum. We compare the sample complexity of KNOCKOUT with that of BTM-PAC (Yue & Joachims, 2011), MallowsMPI (Busa-Fekete et al., 2014a), and AR (Heckel et al., 2016). BTM-PAC is an (ϵ, δ) -PAC algorithm for the same model considered in this paper. MallowsMPI finds a Condorcet winner which exists under our general model. AR finds the maximum according to Borda scores. We also tried PLPAC (Szörényi et al., 2015), developed originally for PL model but the algorithm could not meet guarantees of $\delta = 0.1$ under this




Figure 1. Comparison of sample complexity for small input sizes, with $\epsilon = 0.05$, and $\delta = 0.1$




Figure 2. Comparison of sample complexity for large input size, with $\epsilon = 0.05$, and $\delta = 0.1$

model and hence omitted. Note that in all the experiments the reported numbers are averaged over 100 runs.

In Figure 1, we compare the sample complexity of algorithms when there are 7, 10 and 15 elements. Our algorithm outperforms all the others. BTM-PAC performs much worse in comparison to others because of high constants in the algorithm. Further BTM-PAC allows comparing an element with itself since the main objective in (Yue & Joachims, 2011) is to reduce the regret. We exclude BTM-PAC for further experiments with higher number of elements.

In Figure 2, we compare the algorithms when there are 50, 100, 200 and 500 elements. Our algorithm outperforms others for higher number of elements too. Performance of AR gets worse as the number of elements increases since Borda scores of the elements get closer to each other and hence AR takes more comparisons to eliminate an element. Notice that number of comparisons is in logarithmic scale and hence the performance of MallowsMPI appears to be close to that of ours.

As noted in (Szörényi et al., 2015), sample complexity of MallowsMPI gets worse as $\tilde{p}(i, j)$ gets close to 0. To




Figure 3. Sample complexity of KNOCKOUT and MallowsMPI for different values of \tilde{q} , with $\epsilon = 0.05$ and $\delta = 0.1$




Figure 4. Sample complexity of KNOCKOUT and MallowsMPI under Mallows model for various values of ϕ

show the pronounced effect, we use the stochastic model $p(1, j) = 0.6 \forall j > 1$, $p(i, j) = 0.5 + \tilde{q} \forall j > i, i > 1$ where $\tilde{q} < 0.1$, and the number of elements is 15. Here too we find 0.05-maximum with $\delta = 0.1$. Note that $i = 1$ is the only 0.05-maximum in this stochastic model. In Figure 3, we compare the algorithms for different values of \tilde{q} : 0.01, 0.005 and 0.001. As discussed above, the performance of MallowsMPI gets much worse whereas our algorithm's performance stays unchanged. The reason is that MallowsMPI finds the Condorcet winner using successive elimination technique and as \tilde{q} gets closer to 0, MallowsMPI takes more comparisons for each elimination. Our algorithm tries to find an alternative which defeats Condorcet winner with probability $\geq 0.5 - 0.05$ and hence for alternatives that are very close to each other, our algorithm declares either one of them as winner after comparing them for certain number of times.

Next we evaluate KNOCKOUT on Mallows model which does not satisfy STI. Mallows is a parametric model which is specified by single parameter ϕ . As in (Busa-Fekete et al., 2014a), we consider $n = 10$ elements and various values for ϕ : 0.03, 0.1, 0.3, 0.5, 0.7, 0.8, 0.9, 0.95 and 0.99. Here again we seek to find 0.05-maximum with $\delta = 0.05$.




Figure 5. Sample complexity of MERGE-RANK for different ϵ

As we can see in Figure 4, sample complexity of KNOCKOUT and MallowsMPI is essentially same under small values of ϕ but KNOCKOUT outperforms MallowsMPI as ϕ gets close to 1 since comparison probabilities grow closer to 1/2. Surprisingly, for all values of ϕ except for 0.99, KNOCKOUT returned Condorcet winner in all runs. For $\phi = 0.99$, KNOCKOUT returned second best element in 10 runs out of 100. Note that $\tilde{p}(1, 2) = 0.0025$ and hence KNOCKOUT still outputted a 0.05-maximum. Even though we could not show theoretical guarantees of KNOCKOUT under Mallows model, our simulations suggest that it can perform well even under this model.

For the stochastic model $p(i, j) = 0.6 \forall i < j$, we run our MERGE-RANK algorithm to find an ϵ -ranking with $\delta = 0.1$. Figure 5 shows that sample complexity does not increase a lot with decreasing ϵ . We attribute this to the subroutine COMPARE that finds the winner faster when the elements are more dissimilar.

Some more experiments are provided in Appendix G.

6. Conclusion

We studied maximum selection and ranking using noisy comparisons for broad comparison models satisfying SST and STI. For maximum selection we presented a simple algorithm with linear, hence optimal, sample complexity. For ranking we presented a framework that improves the performance of many ranking algorithms and applied it to merge ranking to derive a near-optimal algorithm.

We conducted several experiments showing that our algorithms perform well and out-perform existing algorithms on simulated data.

The maximum-selection experiments suggest that our algorithm performs well even without STI. It would be of interest to extend our theoretical guarantees to this case. For ranking, it would be interesting to close the $(\log \log n)^3$ ratio between the upper- and lower- complexity bounds.

7. Acknowledgements

We thank Yi Hao and Vaishakh Ravindrakumar for very helpful discussions and suggestions, and NSF for supporting this work through grants CIF-1564355 and CIF-1619448.

References

- Acharya, Jayadev, Jafarpour, Ashkan, Orlitsky, Alon, and Suresh, Ananda Theertha. Sorting with adversarial comparators and application to density estimation. In *ISIT*, pp. 1682–1686. IEEE, 2014a.
- Acharya, Jayadev, Jafarpour, Ashkan, Orlitsky, Alon, and Suresh, Ananda Theertha. Near-optimal-sample estimators for spherical gaussian mixtures. *NIPS*, 2014b.
- Acharya, Jayadev, Falahatgar, Moein, Jafarpour, Ashkan, Orlitsky, Alon, and Suresh, Ananda Theertha. Maximum selection and sorting with adversarial comparators and an application to density estimation. *arXiv preprint arXiv:1606.02786*, 2016.
- Ajtai, Miklós, Feldman, Vitaly, Hassidim, Avinatan, and Nelson, Jelani. Sorting and selection with imprecise comparisons. *ACM Transactions on Algorithms (TALG)*, 12(2):19, 2015.
- Bradley, Ralph Allan and Terry, Milton E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- Busa-Fekete, Róbert, Hüllermeier, Eyke, and Szörényi, Balázs. Preference-based rank elicitation using statistical models: The case of mallows. In *Proc. of the ICML*, pp. 1071–1079, 2014a.
- Busa-Fekete, Róbert, Szörényi, Balázs, and Hüllermeier, Eyke. Pac rank elicitation through adaptive sampling of stochastic pairwise preferences. In *AAAI*, 2014b.
- Feige, Uriel, Raghavan, Prabhakar, Peleg, David, and Upfal, Eli. Computing with noisy information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994.
- Heckel, Reinhard, Shah, Nihar B, Ramchandran, Kannan, and Wainwright, Martin J. Active ranking from pairwise comparisons and when parametric assumptions don't help. *arXiv preprint arXiv:1606.08842*, 2016.
- Herbrich, Ralf, Minka, Tom, and Graepel, Thore. Trueskill: a bayesian skill rating system. In *Proceedings of the 19th International Conference on Neural Information Processing Systems*, pp. 569–576. MIT Press, 2006.
- Jang, Minje, Kim, Sunghyun, Suh, Changho, and Oh, Sewoong. Top- k ranking from pairwise comparisons: When spectral ranking is optimal. *arXiv preprint arXiv:1603.04153*, 2016.
- Luce, R Duncan. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2005.
- Mallows, Colin L. Non-null ranking models. i. *Biometrika*, 44(1/2):114–130, 1957.
- Mukherjee, Sudipta. *Data structures using C: 1000 problems and solutions*. McGraw Hill Education, 2011.
- Negahban, Sahand, Oh, Sewoong, and Shah, Devavrat. Iterative ranking from pair-wise comparisons. In *NIPS*, pp. 2474–2482, 2012.
- Negahban, Sahand, Oh, Sewoong, and Shah, Devavrat. Rank centrality: Ranking from pairwise comparisons. *Operations Research*, 2016.
- Plackett, Robin L. The analysis of permutations. *Applied Statistics*, pp. 193–202, 1975.
- Radlinski, Filip and Joachims, Thorsten. Active exploration for learning rankings from clickthrough data. In *Proceedings of the 13th ACM SIGKDD*, pp. 570–579. ACM, 2007.
- Radlinski, Filip, Kurup, Madhu, and Joachims, Thorsten. How does clickthrough data reflect retrieval quality? In *Proceedings of the 17th ACM conference on Information and knowledge management*, pp. 43–52. ACM, 2008.
- Rajkumar, Arun and Agarwal, Shivani. A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *Proc. of the ICML*, pp. 118–126, 2014.
- Syrkkanis, Vasilis, Krishnamurthy, Akshay, and Schapire, Robert E. Efficient algorithms for adversarial contextual learning. *arXiv preprint arXiv:1602.02454*, 2016.
- Szörényi, Balázs, Busa-Fekete, Róbert, Paul, Adil, and Hüllermeier, Eyke. Online rank elicitation for plackett-luce: A dueling bandits approach. In *NIPS*, pp. 604–612, 2015.
- Urvoy, Tanguy, Clerot, Fabrice, Féraud, Raphael, and Naamane, Sami. Generic exploration and k-armed voting bandits. In *Proc. of the ICML*, pp. 91–99, 2013.
- Yue, Yisong and Joachims, Thorsten. Beat the mean bandit. In *Proc. of the ICML*, pp. 241–248, 2011.
- Zhou, Yuan and Chen, Xi. Optimal pac multiple arm identification with applications to crowdsourcing. 2014.