Efficient Top-K Ranking from Noisy Pairwise Comparisons

{ Rahul Parhi, Soheil Mohajer }   University of Minnesota

Introduction

Design of rank aggregation algorithms is an active area of research. These algorithms have applications in many settings such as social choice, web search, crowdsourcing, and recommendation systems. The main goal is to retrieve an ordering of a set of items when only partial and often noisy preference information are known. It is often the case that only a small subset of the top items, say the top-K, are desired. The focus of this work is on examining several existing algorithms both in the *passive* and *active* settings that rank via noisy *pairwise comparisons* to gain insight into a potentially new *active* top-K ranking algorithm.

Noise Models

Ranking via noisy measurements models the inconsistency that arises when humans are given the task to compare two objects. Hence, several probabilistic models have been proposed to characterize this phenomenon.

Let $i \succ j$ mean that item i is ranked higher than item j. For every item i, let $w_i \in \mathbb{R}$ be the underlying weight (or score) of item i so that the true ranking of a set of items is the non-increasing sorted order of the weights of the items, i.e., $i \succ j \iff w_i > w_j$.

Consider the three models below. The event $\{i \text{ beats } j\}$ refers to when i and j are compared a *single* time.

Bradley-Terry-Luce (BTL)

$$i \succ j \iff \mathbb{P}[i \text{ beats } j] = \frac{w_i}{w_i + w_j}$$

Uniform Noise

$$i \succ j \iff \mathbb{P}[i \text{ beats } j] = \frac{1}{2} + \gamma$$

where $\gamma \in (0, 0.5)$ is an arbitrary constant.

General Model

$$i \succ j \iff \mathbb{P}[i \text{ beats } j] > \frac{1}{2}$$

It is obvious that the General Model is a superset of the BTL and Uniform Noise models.

References

[1] Sahand Negahban, Sewoong Oh, and Devavrat Shah. Rank centrality: Ranking from pairwise comparisons. *Operations Research*, 2016.

[2] Mark Braverman, Jieming Mao, and S Matthew Weinberg. Parallel algorithms for select and partition with noisy comparisons. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 851–862. ACM, 2016.

[3] Soheil Mohajer and Changho Suh. Active top-K ranking from noisy comparisons. In *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*, pages 875–882. IEEE, 2016.

Mohajer-Suh

Mohajer-Suh [3] propose a top-K *active* algorithm with sample complexity $O(mn + mK \log K)$ under the general model. The algorithm has two components. The first is top-1 identification as shown in Figure 1,

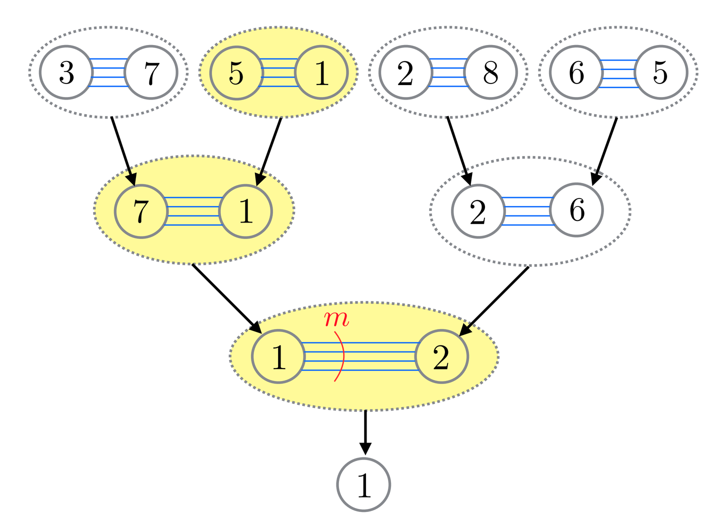

Figure 1: Binary tree to compute the top-1 item.

where comparisons between items are repeated m times. The second component is the top-K identification as shown in Figure 2.

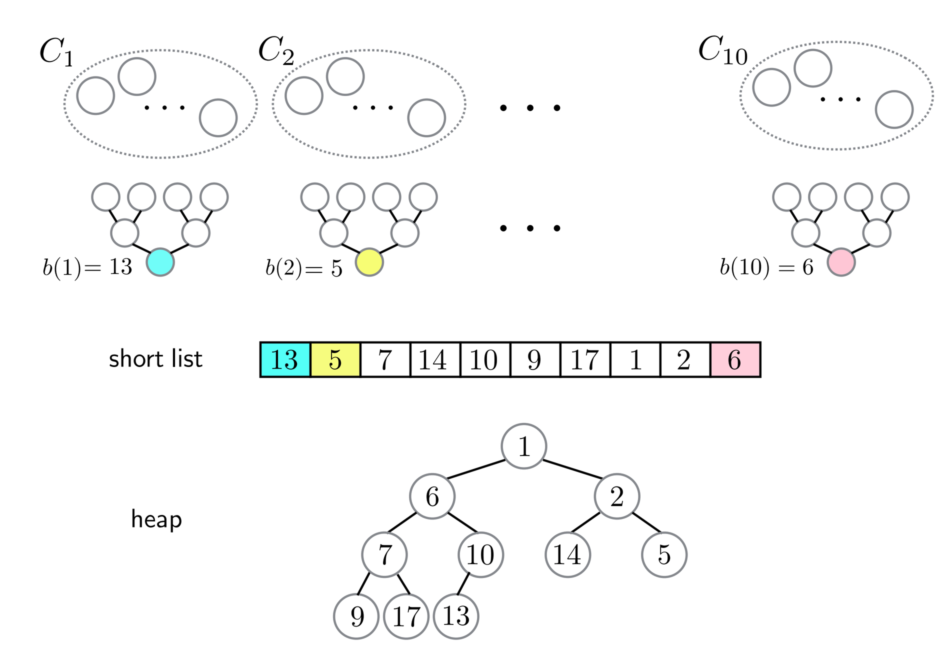

Figure 2: Top items from $K = 10$ subsets.

The set of n items, U, is divided into K disjoint subsets so that $C_1 \sqcup C_2 \sqcup \cdots \sqcup C_{\lceil n/K \rceil} = U$. The procedure of Figure 1 is repeated on each C_i to get the top item from each subset. These are then stored in a heap, H, as shown in Figure 2. Then Extract-Max(H) is stored and a new top is computed from the associated subset of the extracted item. This is repeated until K items have been stored. These are the top-K.

Braverman-Mao-Weinberg

Braverman-Mao-Weinberg [2] propose a top-1 *active* algorithm with minimal sample complexity, $O(n/\gamma^2)$, under the uniform noise model.
Of the set of items U, a random subset $S \subset U$ is found, the top item $x \in S$ is computed naïvely, then all elements in U are compared against x multiple times and thrown out if worse. Finally, $|U|$ will be small, and the top item $a \in U$ can be computed naïvely and returned.

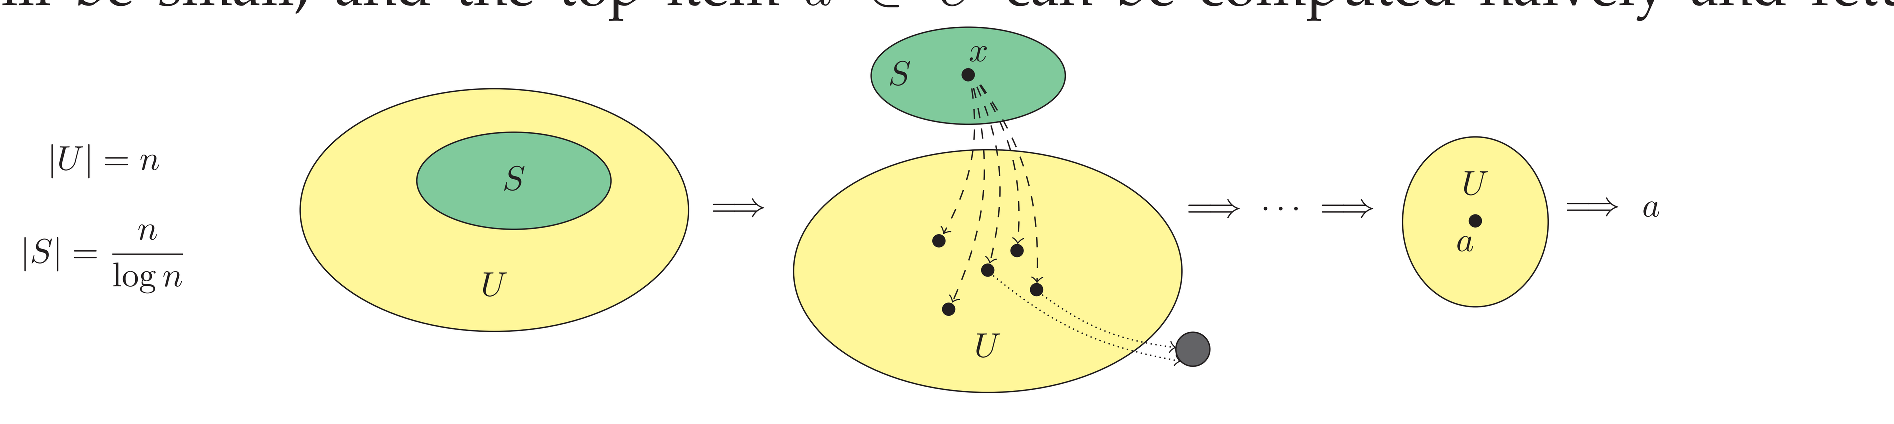

Figure 3: Braverman-Mao-Weinberg top-1 algorithm

Experimental Results

Since our focus is *active* ranking, we simulate Mohajer-Suh and Braverman-Mao-Weinberg and compare their performance. We simulate Mohajer-Suh under the Uniform Noise Model for $K = 10$, $n = 100000$, and various γ as shown in Figure 4. We compare the performance of the algorithms for the special case of $K = 1$ with $n = 1000$ and $\gamma = 0.1$ as shown in Figure 5.

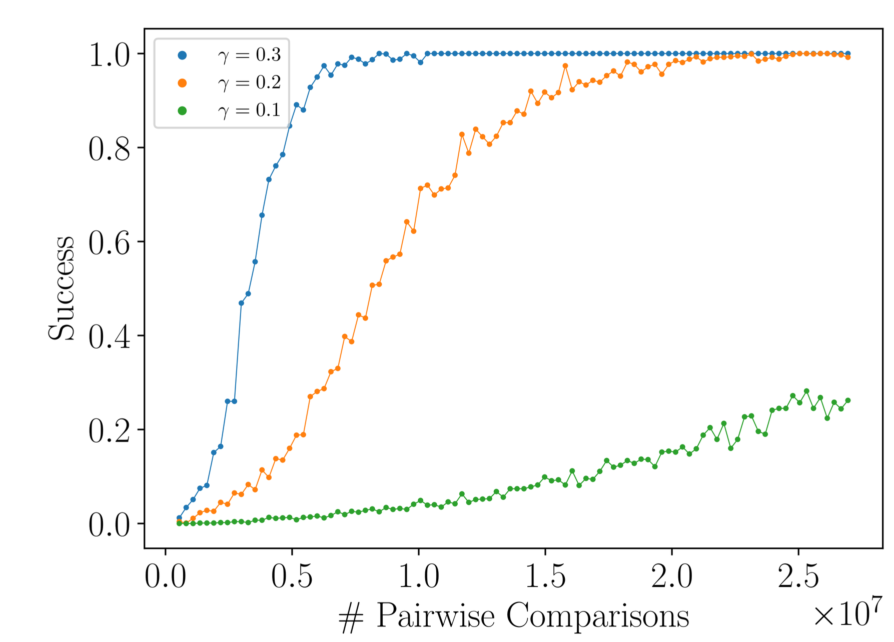

Figure 4: Mohajer-Suh top-10 under Uniform Noise.

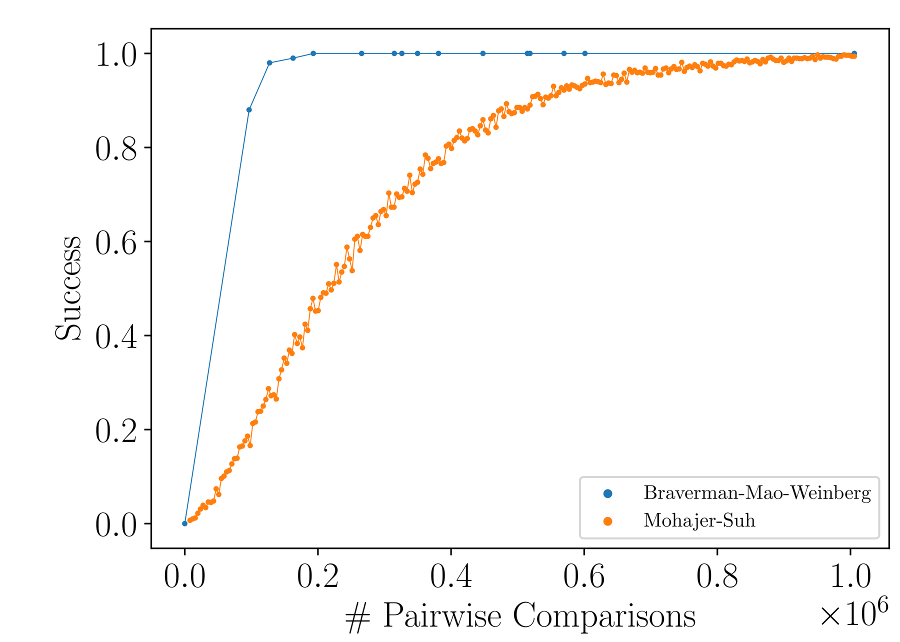

Figure 5: Mohajer-Suh vs. Braverman-Mao-Weinberg top-1.

Rank Centrality

Rank Centrality [1] is a *passive* ranking algorithm under the BTL model that assigns a total ranking as well as scores to a set of items. Essentially, a random walk is taken on the graph of items according to the transition matrix:

$$P(i,j) = \begin{cases} \frac{1}{d_{\max}} A(i,j), & i \neq j \\ 1 - \frac{1}{d_{\max}} \sum_{k: k \neq i} A(i,k), & i = j \end{cases}$$

where $A(i,j)$ is the fraction of times j beats i, and $d_{\max}$ is the maximum out-degree of any node. The total ranking is computed by sorting the stationary distribution of P. This style of approach is called *spectral ranking*.

Conclusion & Remarks

Braverman-Mao-Weinberg achieves a top-1 algorithm with very minimal sample complexity as seen in Figure 5 and succeeds in finding the top-1 item with probability $1 - e^{-O(c)}$.

Mohajer-Suh succeeds in finding top-K items with probability $1 - 2^{-O(\log \log n)}$, when $K = O(\log n)$, also with minimal sample complexity, though, as seen in Figure 5, Braverman-Mao-Weinberg significantly beats Mohajer-Suh for $K = 1$.

This gives insight to a potential new algorithm that extends the filtering technique used in Braverman-Mao-Weinberg to top-K ranking, and will the the focus of future work.