

Crowdsourced Top- k Algorithms: An Experimental Evaluation

Xiaohang Zhang Guoliang Li Jianhua Feng

Department of Computer Science, Tsinghua University, Beijing, China
zhangxiaohang12@mails.tsinghua.edu.cn; {liguoliang, fengjh}@tsinghua.edu.cn

ABSTRACT

Crowdsourced top- k computation has attracted significant attention recently, thanks to emerging crowdsourcing platforms, e.g., Amazon Mechanical Turk and CrowdFlower. Crowdsourced top- k algorithms ask the crowd to compare the objects and infer the top- k objects based on the crowdsourced comparison results. The crowd may return incorrect answers, but traditional top- k algorithms cannot tolerate the errors from the crowd. To address this problem, the database and machine-learning communities have independently studied the crowdsourced top- k problem. The database community proposes the heuristic-based solutions while the machine-learning community proposes the learning-based methods (e.g., maximum likelihood estimation). However, these two types of techniques have not been compared systematically under the same experimental framework. Thus it is rather difficult for a practitioner to decide which algorithm should be adopted. Furthermore, the experimental evaluation of existing studies has several weaknesses. Some methods assume the crowd returns high-quality results and some algorithms are only tested on simulated experiments. To alleviate these limitations, in this paper we present a comprehensive comparison of crowdsourced top- k algorithms. Using various synthetic and real datasets, we evaluate each algorithm in terms of result quality and efficiency on real crowdsourcing platforms. We reveal the characteristics of different techniques and provide guidelines on selecting appropriate algorithms for various scenarios.

1. INTRODUCTION

Given a set of objects, the top- k problem aims to find the top- k best objects from the set. Due to its widespread applications, top- k algorithms have been widely adopted in many real-world systems, e.g., search engines and advertisement systems. If the objects can be compared by machines, e.g., numerical values, existing deterministic top- k algorithms can efficiently compute top- k objects, e.g., the heap-based algorithms [5]. However in many real-world applications, the objects are rather hard to compare for ma-

chines. For example, in image search, comparing the picture clarity is hard for machines. As another example, in natural language processing, comparing the understanding difficulty between different sentences is also rather hard for machines. Obviously these comparisons are easy for human. Thus, crowdsourced top- k computation has been widely studied, thanks to emerging crowdsourcing platforms, e.g., Amazon Mechanical Turk and CrowdFlower. Crowdsourced top- k algorithms ask the crowd to compare objects and infer the top- k objects based on the crowdsourced comparison results.

As the crowd may return incorrect results, a natural problem is to obtain high-quality top- k answers by aggregating the inaccurate comparison results from the crowd. The database community and machine-learning community have independently studied the crowdsourced top- k problem. The database community takes top- k as the basic operation in databases and proposes heuristic-based solutions [9,18] while the machine-learning community focuses on ranking documents/images and proposes learning-based methods (e.g., maximal likelihood estimation [4,16], matrix decomposition [11]). However, these two categories of techniques have not been compared systematically under the same experimental framework. Thus it is difficult for a practitioner to decide which method should be adopted. Furthermore, the experimental evaluation of existing studies has weaknesses. Some methods assume the crowd returns high-quality results while some algorithms are only tested on simulated experiments.

To address these problems, this paper presents a comprehensive comparison on crowdsourced top- k algorithms. Most of existing studies utilize a comparison-based method and employ a two-step strategy. The first step (called pair selection) selects b object pairs and crowdsources them. And the second step (called result inference) infers the top- k pairs based on the comparison results of the crowdsourced pairs. We make a comprehensive survey on both pair selection and top- k inference algorithms. We also discuss other crowdsourced top- k algorithms. We evaluate all of the algorithms on real crowdsourcing platforms.

To summarize, we make the following contributions. (1) We provide a comprehensive survey on more than twenty crowdsourced top- k algorithms. (2) We compare existing crowdsourced top- k algorithms from different research communities through extensive experiments with a variety of synthetic and real datasets on real crowdsourcing platforms. (3) We report comprehensive findings obtained from the experiments. We provide new insights on the strengths and weaknesses of existing algorithms that can guide practitioners to select appropriate algorithms for various scenarios.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org.

Proceedings of the VLDB Endowment, Vol. 9, No. 8
Copyright 2016 VLDB Endowment 2150-8097/16/04.

2. PRELIMINARIES

2.1 Problem Formulation

DEFINITION 1 (CROWDSOURCED TOP- k PROBLEM).

Given an object set $\mathcal{O} = \{o_1, o_2, \dots, o_n\}$, where the objects are comparable but hard to compare by machines, find a k -size object set $\mathcal{R} = \{o_1, o_2, \dots, o_k\}$ where o_i is preferred to o_j (denoted by $o_i \succ o_j$) for $o_i \in \mathcal{R}$ and $o_j \in \mathcal{O} - \mathcal{R}$.

Any comparison operation (\succ) can be used to ask the crowd to compare objects. For example, given six images, we want to identify top-3 clearest ones. Obviously it is hard to use machines to decide which image is clearer and we need to utilize the crowd to compare the images. The goal of devising crowdsourced algorithms is to identify the top- k objects as accurate as possible and we will discuss two metrics to quantify crowdsourced algorithms in Section 6.1.

2.2 Workflow

To utilize the crowd to find top- k objects, we need to generate crowdsourced tasks. There are two widely-used ways to generate crowdsourced tasks. The first is pairwise comparison, which selects two objects and asks the crowd to choose the preferred one. The second is rating, which selects multiple objects and asks the crowd to assign a rate for each object. The rating-based method has some weaknesses. First, the crowd prefers pairwise comparisons to ratings as the former is much easier. Second, it is rather hard for the crowd to assign an accurate rate and the objects in different rating groups may not be correctly compared. Thus the rating-based method usually has a lower accuracy than pairwise comparison [12,14]. Due to these reasons, most of existing works use pairwise comparisons [13]. We first focus on pairwise comparisons, and then discuss hybrid methods that combine ratings and pairwise comparisons [12,20].

3. PAIRWISE COMPARISONS

Given a set with n objects, if all pairs are crowdsourced, there are $\binom{n}{2}$ pairs. Due to the high monetary cost for the crowd, it is not acceptable to enumerate all pairs for large object sets. An alternative is to ask \mathcal{B} pairs, where \mathcal{B} is a given budget. To achieve high recall, an iterative method is widely adopted, where b pairs are crowdsourced in each round. Based on the comparison results of the crowdsourced pairs, it decides how to select b pairs in next round. Iteratively, it crowdsources \mathcal{B} pairs in total and uses the comparison results on these pairs to infer the top- k objects.

For each selected pair, a microtask is generated, which asks the crowd to compare the pair. To reduce crowd errors, each microtask is assigned to multiple workers and the final result is aggregated based on the results of these workers, e.g., weighted majority vote. Thus for each pair (o_i, o_j) , two numbers are reported from the crowd: M_{ij}, M_{ji} , where M_{ij} (M_{ji}) is the number of workers who prefer o_i (o_j) to o_j (o_i). An aggregated weight w_{ij} is computed based on M_{ij} and M_{ji} , e.g., $w_{ij} = \frac{M_{ij}}{M_{ij} + M_{ji}}$. We can use a graph to model the comparison results of the crowdsourced pairs.

DEFINITION 2 (GRAPH MODEL). Given a set of crowdsourced pairs, a directed graph is constructed where nodes are objects and edges are comparison results. For each pair (o_i, o_j) , if the aggregated result is $o_i \succ o_j$ (e.g., $w_{ij} > \frac{1}{2}$), there is a directed edge from o_i to o_j with weight w_{ij} .

For example, consider the six objects in Figure 1, $o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5 \succ o_6$. Assume seven pairs are crowdsourced and each pair is answered by 3 workers (see Figure 1). $M_{12} = 2$ and $M_{21} = 1$ denote that two workers

Table 1: Notations Used In The Paper.

Notation	Description
o_i	an object in set \mathcal{O}
n	the number of objects in set \mathcal{O}
k	the number of reported objects
b	the budget of selected pairs in each round
\mathcal{L}	collected pairwise comparison results from the crowd
$o_i \succ o_j$	o_i is preferred to o_j
M_{ij}	the number of workers reporting $o_i \succ o_j$
M	voting matrix collected from the crowd
p_{ij}	the probability of $o_i \succ o_j$
\hat{s}	estimation matrix of the latent scores for all objects
s_i	$s_i = \hat{s}[i]$, estimated score for object o_i
δ_i	variance of score for object o_i
η_w	the estimated accuracy for worker w

	o_1	o_2	o_3	o_4	o_5	o_6
o_1	—	2	3	3	—	—
o_2	1	—	—	3	3	—
o_3	0	—	—	—	3	—
o_4	0	0	—	—	3	—
o_5	—	0	0	—	—	2
o_6	—	—	—	—	1	—




Figure 1: Pairwise Result. Figure 2: Graph Model.

Algorithm 1: Crowdsourced Top- k Algorithm

Input: Objects \mathcal{O} ; Budget \mathcal{B} ; Budget in each round b
Output: top- k highest ranked objects \mathcal{R}

- 1 $\mathcal{U} = \text{All pairs in } \mathcal{O}$; // unanswered pairs
- 2 $\mathcal{E} = \phi$; // answered pairs
- 3 **while** $\mathcal{B} > 0$ **do**
- 4 $\mathcal{Q} = \text{PAIRSELECTION}(\mathcal{O}, b, \mathcal{E}, \mathcal{U})$;
- 5 Publish questions \mathcal{Q} to the crowd ;
- 6 Collect answers from the crowd ;
- 7 $\mathcal{E} = \mathcal{E} + \mathcal{Q}$; $\mathcal{U} = \mathcal{U} - \mathcal{Q}$; $\mathcal{B} = \mathcal{B} - b$;
- 8 $\mathcal{R} = \text{RESULTINFERENCE}(\mathcal{O}, \mathcal{E})$;
- 9 **Return** \mathcal{R} ;

prefer o_1 to o_2 and one worker prefers o_2 to o_1 . If $o_i \succ o_j$, we says o_i wins (o_j loses) or o_i is better than o_j (o_j is worse than o_i). Based on the crowdsourced answers, a graph is constructed in Figure 2. Table 1 summarizes the notations.

The pairwise-comparison-based method contains two steps. The first is **result inference**, which infers the top- k answers based on the crowdsourced pairs. The second is **pair selection**, which selects next b pairs based on the current result. Algorithm 1 shows the pseudo code. It first selects b pairs (line 4). Then it publishes these b pairs to the crowd and collects the results (lines 5-6). Finally it infers the top- k results (line 8). We will discuss the details of result inference and pair selection in Sections 4 and 5 respectively.

4. RESULT INFERENCE

In this section, we discuss how to infer the top- k results. We first discuss the heuristic-based algorithms in Section 4.1 and then introduce the machine-learning methods in Section 4.2. Next we discuss how to extend heap-based methods to infer the top- k results in Section 4.3. Lastly, we present the hybrid methods in Section 4.4.

4.1 Heuristic-Based Methods

Guo et al. prove that finding the maximal (top-1) object is NP-Hard by a reduction from Kemeny rankings [9]. Thus inferring the top- k objects is also NP-Hard and many heuristics are proposed to infer top- k objects.

4.1.1 Score-Based Algorithms

The score-based algorithms assign each object o_i with a score s_i and the k objects with the largest scores are selected as top- k answers. Next we discuss how to compute scores.

(1) Election Algorithms. The election algorithms that determine the winner of an election can be used to compute the score of an object. There are two famous election algorithms BordaCount [1] and Copeland [17].

BordaCount [1]. The score of o_i is its out-degree $d^+(o_i)$ (the number of wins compared with its out-neighbors), i.e.,

$$s_i = d^+(o_i). \quad (1)$$

In our example, s_1 is 3 as o_1 has 3 out-neighbors. The scores of $\langle o_1, o_2, o_3, o_4, o_5, o_6 \rangle$ are $\langle 3, 2, 1, 1, 1, 0 \rangle$.

Copeland [17]. The score of object o_i is its out-degree $d^+(o_i)$ minus its in-degree $d^-(o_i)$ (the number of wins minus the number of losses), i.e.,

$$s_i = d^+(o_i) - d^-(o_i). \quad (2)$$

In our example, s_2 is 1 as o_2 has 2 out-neighbors and 1 in-neighbor. The scores of $\langle o_1, o_2, o_3, o_4, o_5, o_6 \rangle$ are $\langle 3, 1, 0, -1, -2, -1 \rangle$. The top-2 answers are o_1 and o_2 .

(2) Max Algorithms. Guo et al. [9] proposed several graph-based algorithms to compute the maximal object, which can be used to assign a score for each object.

Local [9]. The two-election methods only consider the neighbors of each object and cannot capture more information. Local algorithm [9] is proposed to solve the max problem by considering 2-hop neighbors and it can naturally be transformed to address the top- k problem. Obviously if an object has more 2-hop out-neighbors (i.e., its out-neighbors' out-neighbors), the object will beat more objects (based on transitivity), and thus the object has a larger score. Similarly, if an object has more 2-hop in-neighbors (i.e., its in-neighbors' in-neighbors), the object will be beaten by more objects, and thus the object has a lower score. Accordingly, the overall score of an object is computed as below.

$$s_i = M_{i*} - M_{*i} + \sum_{o_j \in \mathcal{D}^+(o_i)} M_{j*} - \sum_{o_j \in \mathcal{D}^-(o_i)} M_{*j}, \quad (3)$$

where $M_{i*} = \sum_j M_{ij}$, $M_{*i} = \sum_j M_{ji}$, and $\mathcal{D}^+(o_i)$ is the out-neighbor set of o_i and $\mathcal{D}^-(o_i)$ is the in-neighbor set of o_i . In our example, the score of o_2 is $(7-2) + (3+2) - 1 = 9$ and the score of o_5 is $(2-10) + 1 - (2+3+6) = -18$. The scores of $\langle o_1, o_2, o_3, o_4, o_5, o_6 \rangle$ are $\langle 20, 9, 1, -4, -18, -11 \rangle$.

Indegree [9]. It computes the score based on the bayesian model. The probability of $o_i \succ o_j$ given M_{ij} and M_{ji} is

$$\begin{aligned} P(o_i \succ o_j | M_{ij}, M_{ji}) &= \frac{P(M_{ij}, M_{ji} | o_i \succ o_j) P(o_i \succ o_j)}{P(M_{ij}, M_{ji})} \\ &= \frac{P(M_{ij}, M_{ji} | o_i \succ o_j)}{P(M_{ij}, M_{ji} | o_i \succ o_j) + P(M_{ij}, M_{ji} | o_j \succ o_i)}, \end{aligned} \quad (4)$$

$$\begin{aligned} P(M_{ij}, M_{ji} | o_i \succ o_j) &= \binom{M_{ij} + M_{ji}}{M_{ij}} p^{M_{ij}} (1-p)^{M_{ji}}; \\ P(M_{ij}, M_{ji} | o_j \succ o_i) &= \binom{M_{ij} + M_{ji}}{M_{ji}} p^{M_{ji}} (1-p)^{M_{ij}}. \end{aligned} \quad (5)$$

p is the estimated worker accuracy which is a fixed value. Besides, it assumes $P(o_i \succ o_j) = P(o_i \prec o_j)$. Then it

Top-k Algorithms		Details			
		Area	Input	Output	Parameters
Heuristic (Local)	BordaCount*	DB	Graph	Score	N/A
	Copeland*	DB	Graph	Score	N/A
	Local*	DB	Graph	Score	N/A
	Indegree*	DB	Graph	Score	N/A
	ELO*	ML	Comparison	Score	$H = 32$
	BRE*	ML	Comparison	Score	N/A
	URE*	ML	Comparison	Score	N/A
	SSCO	ML	Comparison	Score	$q = 0.1$
	SSSE	ML	Comparison	Score	$q = 0.1$
Heuristic (Global)	Iterative	DB	Graph	Subset	N/A
	PathRank	ML	Graph	Subset	N/A
	AdaptiveReduce	ML	Comparison	Subset	$p = 0.9$
	MPageRank*	DB	Graph	Score	$c = 1$
	RankCentrality*	ML	Graph	Score	N/A
	SSRW*	ML	Graph	Score	$q = 0.1$
ML	TrueSkill*	ML	Comparison	Score	$\varepsilon = 0$
	CrowdBT*	ML	Comparison	Score	$\lambda = 0.5$
	CrowdGauss*	ML	Comparison	Score	N/A
	HodgeRank*	ML	Comparison	Score	N/A
Extension	TwoStageHeap	DB	Comparison	Subset	$X = \frac{0.1n}{k^2}$
Rate + Compare	Combine*	ML	Rate+Compare	Score	N/A
	Hybrid*	DB	Rate+Compare	Score	$\delta = 0.1$
	HybridMPR*	DB	Rate+Graph	Score	$\beta = 0.6$

Table 2: Details for top- k Inference Algorithms (Methods with * can also be used for sorting.)

computes the score of object o_i as below.

$$s_i = \sum_{j \neq i} P(o_i \succ o_j | M_{ij}, M_{ji}). \quad (6)$$

Note that if o_i and o_j are not compared, $P(o_i \succ o_j | M_{ij}, M_{ji}) = 0.5$. In our example, if we set p to 0.55 [9], the scores for $\langle o_1, o_2, o_3, o_4, o_5, o_6 \rangle$ are $\langle 2.84, 2.74, 2.5, 2.35, 2.11, 2.45 \rangle$.

Modified PageRank(MPageRank) [9]. It extends the original PageRank by considering the crowdsourced comparisons and computes the score of each object as below.

$$s_i = pr_k[i] = \frac{1-c}{n} + c \sum_{j \neq i} \frac{M_{ij}}{M_{*j}} pr_{k-1}[j], \quad (7)$$

where $pr_k[i]$ is the score of o_i in the k -th iteration (initially $pr_1[i] = \frac{1}{n}$) and c is a damping factor and set to 1 by default). The scores are $\langle 0.5, 0.5, 5e^{-11}, 5e^{-11}, 1.67e^{-11}, 1.67e^{-11} \rangle$.

(3) Ranking Algorithms. There are some ranking-based algorithms that focus on ranking documents/images.

RankCentrality [15]. The algorithm adopts an iterative method based on random walk. Its core component is to construct an $n \times n$ transition matrix P , in which

$$P_{ij} = \begin{cases} \frac{1}{d_{max}} w_{ji} & \text{if } i \neq j \\ 1 - \frac{1}{d_{max}} \sum_{k \neq i} w_{ki} & \text{if } i = j \end{cases}, \quad (8)$$

where w_{ji} is the weight of edge from o_j to o_i and d_{max} is the maximum in-degree of a node (i.e., $d_{max} = \max_{o_i} d^-(o_i)$). Then it initializes a $1 \times n$ matrix \hat{s} and computes $\hat{s} \times P \times P \times \dots$ until convergence. The score of o_i is computed as

$$s_i = \hat{s}[i] \text{ where } \hat{s} = \lim_{t \rightarrow \infty} \hat{s} \times P^t. \quad (9)$$

In our example, the scores for $\langle o_1, o_2, o_3, o_4, o_5, o_6 \rangle$ are $\langle 2.26, 1.68, 0.93, 0.63, 0.38, 0.91 \rangle$.

Balanced Rank Estimation(BRE)/Unbalanced Rank Estimation(URE) [19]. The score is computed based on the probability theory. The balanced rank estimation (BRE) considers both incoming and outgoing edges. To compute the score s_j , it computes the relative difference of the number of objects proceeding and succeeding o_j :

$$s_j = \frac{\sum_{i \neq j} b_{ij}(2w_{ij} - 1)}{2\alpha n} \propto \sum_{i \neq j} b_{ij}(2w_{ij} - 1). \quad (10)$$

where α is the selection rate (i.e., $\alpha \binom{n}{2}$ pairs will be compared); $w_{ij} = 1$ if $o_i \succ o_j$; $w_{ij} = 0$ otherwise; $b_{ij} = 1$ if o_i and o_j are compared by the crowd; $b_{ij} = 0$ otherwise.

The unbalanced rank estimation (URE) computes the score of o_i only based on its incoming edges. To compute the score s_j , it computes the number of objects proceeding o_j :

$$s_j = \frac{1}{\alpha n} \sum_{i \neq j} b_{ij} w_{ij}^n \propto \sum_{i \neq j} b_{ij} w_{ij}^n, \quad (11)$$

In our example, if we set the value of α to 1, then the scores are $\langle -0.25, -0.08, 0, 0.08, 0.17, 0.08 \rangle$ for BRE. And the scores are $\langle 0, 0.17, 0.17, 0.33, 0.5, 0.17 \rangle$ for URE.

ELO [6]. The ELO method is a chess ranking system, which can be applied to address the pairwise top- k problem. The basic idea is that, if object o_i with higher ranking beats another lower one o_j , only a few points will be added to s_i ; on the contrary, if o_j wins, a lot of points will be added to s_j . Formally, each object o_i is randomly assigned a score s_i initially. When o_i is compared with o_j , the two scores s_i and s_j will be updated as below.

$$\begin{aligned} s_i &= s_i + H(C_i - \frac{1}{1 + 10^{(s_j - s_i)/400}}); \\ s_j &= s_j + H(1 - C_i - \frac{1}{1 + 10^{(s_i - s_j)/400}}), \end{aligned} \quad (12)$$

where H is a tuning parameter (set to 32 by default). $C_i = 1$ if $o_i \succ o_j$; $C_i = 0$ otherwise. The scores are $\langle 97.3, 65.3, 50.7, 36.2, 23.6, 32.8 \rangle$ if the initial score is 50.

4.1.2 Iterative Reduce Top- k Algorithms

They iteratively eliminate lowly ranked objects that have small possibilities in the top- k results, until k objects left.

Iterative [9]. Guo et al. [9] improve the max algorithm and propose the iterative algorithm to improve the quality. It first utilizes the score-based methods to compute the scores of each object and then removes the half of the objects with small scores. Next it re-computes the scores on the survived objects and repeats the iterations until k objects left.

Eriksson et al. [7] propose two iterative reduce algorithms which are devised to rank documents and images.

PathRank [7]. The main idea of PathRank is to perform a “reverse” depth first search (DFS) for each node, which traverses the graph by visiting the in-neighbors of each node. If it finds a path with length larger than k , it eliminates the object as k objects have already been better than the object. However this method is not robust [7], because if not all the comparisons are observed, due to missing edges, objects ranked far from the top- k answers could potentially have no more than k -paths observed, and non-top- k objects can also be erroneously returned as an estimated top- k result.

AdaptiveReduce [7]. Initially there are n objects. It first randomly selects $(16(p - \frac{1}{2})^{-2} + 32) \log n$ objects as the chosen set X_c , where p is the worker accuracy. Then it wants to select an informative set and utilizes the set to eliminate objects with small possibilities in the top- k answers. To this end, it computes the voting count v_i for each object o_i in the chosen set, i.e., $v_i = \sum_{j=1}^n M_{ij}$. Next, it computes a subset of X_c as the informative set X_v , with voting count between $\frac{n}{4}$ and $\frac{3n}{4}$, i.e., $X_v = \{x \in X_c : \frac{n}{4} \leq v_x \leq \frac{3n}{4}\}$.

Based on X_v , it computes the winning count t_i of each object o_i , i.e., $t_i = \sum_{x \in X_v} M_{ix}$. Then objects with winning counts larger than $\frac{|X_v|}{2}$ will be kept, because they win many objects in the informative set. Then, it repeats the above steps using the survived objects until finding top- k objects.

Fekete et al. [3] propose bound-based methods for ranking in recommendation systems and image search, which build a stochastic matrix Q where $Q_{ij} = \frac{M_{ij}}{M_{ij} + M_{ji}}$, and devise several algorithms to identify top- k results from the matrix. **Sampling Strategy with Copeland’s Ranking (SSCO)** [3] selects the top- k rows with most entries above 0.5 as the top- k results. **Sampling Strategy with Sum of Expectations (SSSE)** [3] selects the top- k rows with the largest average value as the top- k results. **Sampling Strategy based on Random Walk (SSRW)** [3] first transfers the matrix Q to the stochastic matrix S where $S_{ij} = \frac{Q_{ij}}{\sum_i Q_{li}}$ and computes the principal eigenvectors (that belong to the eigenvalue 1). Then it identifies the top- k rows with the largest eigenvalues as the top- k answers.

4.2 Machine-Learning Methods

These methods assume that each object has a latent score which follows a certain distribution. Then they utilize machine-learning techniques to estimate the scores. Lastly, they select k objects with the largest scores as the top- k results.

4.2.1 Maximum Likelihood Estimation

CrowdBT with Bradley-Terry Model [2]. This method uses Bradley-Terry (BT) model to estimate the latent score [2]. In the BT model, the probability of $o_i \succ o_j$ is assumed as $\frac{e^{s_i}}{e^{s_i} + e^{s_j}}$. Then based on the crowdsourced comparisons, it computes the latent scores by maximizing

$$\sum_{o_i \succ o_j \in \mathcal{L}} \log\left(\frac{e^{s_i}}{e^{s_i} + e^{s_j}}\right), \quad (13)$$

where \mathcal{L} is a set of crowdsourced comparison results.

However, the BT model does not consider the crowd quality. As different workers have different quality, it is very important to tolerate errors in estimating the scores. To address this problem, Chen et al. [4] propose the CrowdBT model, which assumes that each worker w has a quality η_w . That is, if worker w returns $o_i \succ o_j$, the probability of $o_i \succ o_j$ is η_w and the probability of $o_j \succ o_i$ is $1 - \eta_w$. In addition, if the pairwise comparison graph is not strongly connected, Equation 13 may not get an accurate estimation for the objects. To solve this problem, Chen et al. [4] add a virtual object to the comparison graph with comparisons to other objects. The score for the virtual object is represented by s_0 . Then it computes the scores by maximizing

$$\begin{aligned} &\sum_w \sum_{o_i \succ o_j \in \mathcal{L}_w} \log\left(\eta_w \frac{e^{s_i}}{e^{s_i} + e^{s_j}} + (1 - \eta_w) \frac{e^{s_j}}{e^{s_i} + e^{s_j}}\right) + \\ &\lambda \sum_{i=1}^n \left(\log\left(\frac{e^{s_0}}{e^{s_0} + e^{s_i}}\right) + \log\left(\frac{e^{s_i}}{e^{s_0} + e^{s_i}}\right)\right), \end{aligned} \quad (14)$$

where \mathcal{L}_w is the set of comparison results by worker w .

CrowdGauss with Gaussian Model [16]. Pfeiffer et al. [16] propose CrowdGauss that uses the gaussian model to estimate the score. It assumes that the score follows the gaussian distribution, where the score is the mean of the distribution. The probability of $o_i \succ o_j$, i.e., $P(o_i \succ o_j)$, can be computed by the cumulative distribution function (Φ) of the two standard gaussian distributions, i.e.,

$$P(o_i \succ o_j) = \Phi(s_i - s_j). \quad (15)$$

Then CrowdGauss computes the scores by maximizing

$$\sum_{o_i \succ o_j \in \mathcal{L}} M_{ij} \cdot \log(\Phi(s_i - s_j)). \quad (16)$$

4.2.2 Matrix Decomposition Based Method

HodgeRank [11]. In order to estimate a global order for n objects, HodgeRank [11] utilizes a matrix decomposition based technique. It first computes a preference matrix Y based on the crowd's answers, where

$$Y_{ij} = \begin{cases} \ln(\frac{M_{ji}}{M_{ij}}) & \text{if } i \neq j \text{ and } M_{ij} \neq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

Then, it constructs another matrix X where $X_{ij} = s_j - s_i$. Lastly, it computes the score by minimizing

$$\sum_{o_i \succ o_j \in \mathcal{L}} (X_{ij} - Y_{ij})^2 = \sum_{o_i \succ o_j \in \mathcal{L}} (s_j - s_i - Y_{ij})^2. \quad (18)$$

4.2.3 Others

TrueSkill [10]. TrueSkill improves ELO by reducing the number of iterations as ELO needs to repeat many times to convergence. Different from ELO, the score of each object o_i is represented by a Gaussian distribution $N(s_i, \delta_i)$, where s_i is the estimated score for o_i and δ_i is the deviation of s_i . Let $\beta = \frac{\sum_i \delta_i}{n}$ denote the average deviation for all objects. \mathcal{U} and \mathcal{V} are two monotonically decreasing functions, which are used to update s_i and δ_i . $\mathcal{U}(\frac{s_i - s_j}{c}, \frac{\varepsilon}{c})$ is the reward score added to o_i or the punishment score subtracted from o_j , where ε is a predefined parameter. $\mathcal{V}(\frac{s_i - s_j}{c}, \frac{\varepsilon}{c})$ indicates the reduction of the deviation. For each comparison result $o_i \succ o_j$, it updates the scores and deviations as below.

$$\begin{aligned} s_i &\leftarrow s_i + \frac{\delta_i^2}{c} \cdot \mathcal{U}(\frac{s_i - s_j}{c}, \frac{\varepsilon}{c}); & s_j &\leftarrow s_j - \frac{\delta_j^2}{c} \cdot \mathcal{U}(\frac{s_i - s_j}{c}, \frac{\varepsilon}{c}) \\ \delta_i^2 &\leftarrow \delta_i^2 \cdot [1 - \frac{\delta_i^2}{c^2} \cdot \mathcal{V}(\frac{s_i - s_j}{c}, \frac{\varepsilon}{c})]; & \delta_j^2 &\leftarrow \delta_j^2 \cdot [1 - \frac{\delta_j^2}{c^2} \cdot \mathcal{V}(\frac{s_i - s_j}{c}, \frac{\varepsilon}{c})], \end{aligned} \quad (19)$$

where $\beta = \frac{\sum_i \delta_i}{n}$, $c = \sqrt{2\beta^2 + \delta_i^2 + \delta_j^2}$, ε is a predefined parameter, and s_i, δ_i are learned from the crowdsourced pairs.

4.3 Extensions of Heap-Based Methods

TwoStageHeap [5]. A two-stage heap algorithm [5] is proposed to extend heap-based algorithms to support crowdsourced top- k computation in databases. In the first phase, the objects are divided into $\frac{n}{X}$ buckets (where $X = \frac{kn}{k^2}$) such that the probability of two top- k objects appearing in the same bucket is at most x (set to 0.1 by default). In each bucket, a tournament based max algorithm [5] is conducted to select the best object in the bucket. Each pair on top levels of the tournament is compared multiple times and each pair on low levels of the tournament is compared only once. The second phase utilizes a heap-based method [8] to identify the top- k results from these best objects. To tolerate errors, when constructing and re-heapifying the heap, each pair is compared by multiple workers and the algorithm uses the majority voting to obtain a combined preference. After popping an object from the heap, the algorithm asks next pairs following the re-heapifying order.

4.4 Combining Rating And Comparison

There are two algorithms [12,20] that combine rating and comparison. The methods in [12] are initially designed to find the maximal object, and they can be extended to address the top- k problem. They utilize the rating and comparison results to learn the score for each object. For rating, they pre-define τ categories and each category χ_c has a range $(\gamma_{c-1}, \gamma_c]$, where $\gamma_0 < \gamma_1 < \dots < \gamma_\tau$. If the score of object o_i falls in $(\gamma_{c-1}, \gamma_c]$, o_i is in the category χ_c .

Combine [20]. It first selects some rating and comparison tasks, and then infers the scores based on these results. The score for each object o_i is modeled by $s_i + \varepsilon_i$, where $\varepsilon_i \sim N(0, \delta^2)$, which is utilized to model crowd errors.

For rating, o_i is in the category χ_c with probability of

$$P(\chi_c | s_i) = Pr(\gamma_{\chi_c-1} < s_i + \varepsilon_i \leq \gamma_{\chi_c}) = \Phi(\frac{\gamma_{\chi_c} - s_i}{\delta}) - \Phi(\frac{\gamma_{\chi_c-1} - s_i}{\delta}) \quad (20)$$

where Φ represents the Cumulative Density Function(CDF) of standard Gaussian distribution. Then it constructs a voting matrix V based on the voting result, where V_{ij} is the number of times that o_i is observed in the j -category.

It computes the probability of observing V given s by:

$$\begin{aligned} P(V|s) &= \prod_{i=1}^n Pr(V_{i1}, V_{i2}, \dots, V_{i\tau} | s_i) \\ &= c_1 \prod_{i=1}^n \prod_{c=1}^{\tau} (\Phi(\frac{\gamma_c - s_i}{\delta}) - \Phi(\frac{\gamma_{c-1} - s_i}{\delta}))^{V_{i,c}} \end{aligned} \quad (21)$$

where c_1 is a constant.

As ε_i has a normal distribution, the pairwise preference can be modeled as,

$$P(o_i \succ o_j) = Pr(s_i + \varepsilon_i > s_j + \varepsilon_j) = \Phi(\frac{s_i - s_j}{\sqrt{2}\delta}) \quad (22)$$

For comparison, it constructs the comparison matrix M and computes the probability of observing M by:

$$P(M|s) = \prod_{i < j \in 1, \dots, n} Pr(M_{ij}, M_{ji} | s_i, s_j) = c_2 \prod_{i \neq j} \Phi(\frac{s_i - s_j}{\sqrt{2}\delta})^{M_{ij}} \quad (23)$$

where c_2 is a constant.

Based on $Pr(V|s)$ and $Pr(M|s)$, it computes $Pr(s|V, M)$ based on the bayesian theory and estimates the score by

$$\hat{s} = \arg \max_s Pr(s|V, M) \quad (24)$$

Monte Carlo methods can be used to compute Equation 24, but the computation cost is rather high. [20] approximates $Pr(s|V, M)$ by $Pr(s|V, M, \hat{\gamma})$, where $\hat{\gamma} = \arg \max_{\gamma} Pr(V, M|\gamma)$. Utilizing the Laplace approximation, an analytical form can be formulated for $Pr(V, M|\gamma)$, and the maximum likelihood estimation of γ can be computed by gradient based optimization methods. After obtaining γ , s can be estimated by maximizing a posteriori probability (MAP) estimation.

Hybrid [12]. It first crowdsources all rating tasks and then selects some candidate objects with higher ratings. Then it chooses some pairs from the candidates as comparison tasks. The score of each object o_i is modeled as a normal variable $N(s_i, \delta^2)$. Given the rating results E_R and comparison results E_C , it estimates the score by

$$\hat{s} = \arg \max_s P(E_R, E_C | s) \quad (25)$$

If the questions are answered independently, the above function can be decomposed to

$$\begin{aligned} P(E_R, E_C | s) &= P(E_R | s) P(E_C | s) \\ P(E_R | s) &= \prod_{i=1}^n \prod_{j=1}^{|R_i^*|} P(R_i^j | s_i) \\ P(E_C | s) &= \prod_{i=1}^{n-1} \prod_{j=i+1}^n \prod_{t=1}^{|M_{ij} + M_{ji}|} P(C_{ij}^t | s_i, s_j) \end{aligned} \quad (26)$$

where $|R_i^*|$ is the number of ratings for object o_i , R_i^j is the assigned category in the j -th rating (by assuming each category has equal width and $\gamma_0 = 0, \gamma_\tau = 1$). Thus, the probability that object o_i lies in category χ_c is computed as

$$P(R_i^j = \chi_c | s_i) = \begin{cases} F_{s_i, \delta^2}(\frac{c}{\tau}) & c = 1 \\ F_{s_i, \delta^2}(\frac{c}{\tau}) - F_{s_i, \delta^2}(\frac{c-1}{\tau}) & 1 < c < \tau \\ 1 - F_{s_i, \delta^2}(\frac{c-1}{\tau}) & c = \tau \end{cases} \quad (27)$$

where $F_{s_i, \delta^2}(x)$ is the CDF of s_i and δ .

Let C_{ij}^t denote the comparison result in the t -th comparison for pair (o_i, o_j) . $C_{ij}^t = 0$ if o_j is superior to o_i , and $P(C_{ij} = 0) + P(C_{ij} = 1) = 1$. Since o_i and o_j both follow the normal distribution, $o_i - o_j \sim N(s_i - s_j, 2\delta^2)$,

$$P(C_{i,j}^t = 0 | s_i, s_j) = F_{s_i - s_j, 2\delta^2}(0). \quad (28)$$

HybridMPR. However, the maximum likelihood estimation is rather expensive, thus a modified PageRank approximation [12] is proposed to estimate the score for each object. It defines the average rating score r_i for each object o_i as,

$$r_i = \begin{cases} dv & \text{if } |R_i^*| = 0 \\ \frac{\sum_j R_{ij}^*}{\tau |R_i^*|} - \frac{1}{2\tau} & \text{otherwise} \end{cases} \quad (29)$$

where dv is a default value (e.g., 0.5) for the case of no rating on o_i . Then the modified PageRank is defined as below,

$$\rho_i = \beta \sum_{j: w_j \neq 0} \frac{w_{j,i}}{w_j} \rho_j + (1 - \beta) \frac{r_i}{\sum_j r_j} \quad (30)$$

where β is the fraction of comparison questions to the total questions, $w_j = \sum_i w_{j,i}$, and $w_{j,i}$ is defined as

$$w_{j,i} = \begin{cases} 1 & \text{if } M_{ij} + M_{ji} = 0 \text{ and } r_i \geq r_j \\ 0 & \text{if } M_{ij} + M_{ji} = 0 \text{ and } r_i < r_j \\ \frac{M_{ij}}{M_{ij} + M_{ji}} & \text{otherwise} \end{cases} \quad (31)$$

5. PAIR SELECTION

This section discusses different pair selection methods. The random selection method that randomly selects next b pairs is not effective [19], because different comparison pairs have different importance to infer the results. For example, in Figure 2, suppose we select pairs (o_1, o_2) , (o_2, o_5) , (o_1, o_5) . As $o_1 \succ o_2$ and $o_2 \succ o_5$, (o_1, o_5) does not need to be asked as its result can be deduced. Similarly, (o_1, o_6) and (o_2, o_6) do not need to be crowdsourced. Thus it is important to select high-quality pairs. We first introduce the heuristic-based methods in Section 5.1 and then discuss the bound-based methods in Section 5.2. Lastly, we discuss the active-learning methods in Section 5.3.

5.1 Heuristic-Based Methods

Guo et al. [9] prove that selecting the pairs to maximize the probability of obtaining the top- k results (given the comparison results of arbitrary crowdsourced pairs) is NP-Hard and propose four heuristics, which are designed for selecting the max (top-1) result. The proposed algorithms first compute a score s_i for object o_i (see Section 4.1). Suppose the sorted objects based on the scores are o_1, o_2, \dots, o_n . Then, the algorithms select the next b pairs as follows.

Max. It selects b pairs: (o_1, o_2) , (o_1, o_3) , \dots , (o_1, o_{b+1}) .

Group. It groups the i -th object with the $(i+1)$ -th object and the selected pairs are (o_1, o_2) , (o_3, o_4) , \dots , (o_{2b-1}, o_{2b}) .

Greedy. It selects the pairs based on $s_i \times s_j$ in descending order, and selects b pairs with the largest value.

Complete. It has two phases. The first phase finds the maximal number of objects and makes a pairwise comparison on them. Thus the first phase finds top- x objects with the highest scores, where x is the largest integer satisfying $\frac{x*(x-1)}{2} \leq b$. The first phase compares every two objects among the first x objects. If the budget is not exhausted in the first phase (i.e., $b > \frac{x*(x-1)}{2}$), then the second phase compares the first $b - \frac{x*(x-1)}{2}$ objects with the $(x+1)$ -th object, i.e., (o_1, o_{x+1}) , (o_2, o_{x+1}) , \dots , $(o_{b - \frac{x*(x-1)}{2}}, o_{x+1})$.

5.2 Bound-Based Methods

SSCO and SSSE estimate a bound for each pair and utilize the bound to select next pairs [3]. Formally, SSCO and SSSE first compute a confidence interval $[l_{ij}, u_{ij}]$, where l_{ij} (u_{ij}) is the lower (upper) bound of the probability of $o_i \succ o_j$, which can be computed as below.

$$u_{ij} = \frac{M_{ij}}{M_{ij} + M_{ji}} + c_{ij}; \quad l_{ij} = \frac{M_{ij}}{M_{ij} + M_{ji}} - c_{ij}, \quad (32)$$

where $c_{ij} = \sqrt{\frac{1}{2(M_{ij} + M_{ji})} \log \frac{2n_{max}^2}{q}}$, n_{max} is a predefined maximal number of crowdsourced tasks for each pair and q is a confidence factor that top- k objects can be found with probability $1-q$. Based on the confidence interval, they select a set S and discard a set D . Since the relationships between pairs in $S \cup D$ have been effectively captured, these pairs do not need to be compared. They propose two algorithms to select pairs that are not in $S \cup D$.

Sampling Strategy with Copeland's Ranking (SSCO) [3]. Based on the confidence interval $[l_{ij}, u_{ij}]$, it computes a bound $L_i = |\{o_j | l_{ij} > 1/2\}|$ of object o_i , which is the number of objects with high probability of $o_i \succ o_j$, and a bound $U_i = |\{o_j | u_{ij} < 1/2\}|$ of object o_i , which is the number of objects with high probability of $o_j \succ o_i$. Given two objects o_i and o_j , if $L_i > n - U_j$, o_i should be preferred to o_j . If $|\{j | L_i > n - U_j\}| > n - k$ (i.e., the number of objects beaten by o_i is larger than $n - k$), o_i will have high probability in the top- k results and is added into set S , i.e.,

$$S = \{o_i : |\{j | L_i > n - U_j\}| > n - k\}. \quad (33)$$

Similarly, if $U_i > n - L_j$, o_j should be preferred to o_i . If $|\{j | U_i > n - L_j\}| > k$ (i.e., the number of objects preferred to o_i is larger than k), o_i has small probability in the top- k results and is added into the set D , i.e.,

$$D = \{o_i : |\{j | U_i > n - L_j\}| > k\}. \quad (34)$$

Sampling Strategy with Sum of Expectations (SSSE) [3].

This method utilizes the sum of expectations to compute the two bounds, and $L_i = \frac{1}{n-1} \sum_{j \neq i} l_{ij}$ is the expectation of lower bound of o_i and $U_i = \frac{1}{n-1} \sum_{j \neq i} u_{ij}$ is the expectation of upper bound of o_i . If $L_i > U_j$, o_i should be preferred to o_j . If $|\{j | L_i > U_j\}| > n - k$ (i.e., the number of objects beaten by o_i is larger than $n - k$), o_i will have high probability in the top- k results and is added into set S , i.e.,

$$S = \{o_i : |\{j | L_i > U_j\}| > n - k\}. \quad (35)$$

Similarly, if $U_i < L_j$, o_j should be preferred to o_i . If $|\{j | U_i < L_j\}| > k$ (i.e., the number of objects preferred to

o_i is larger than k), o_i will have small probability in the top- k results and is added into the set D , i.e.,

$$D = \{o_i : |\{j | U_i < L_j\}| > k\}. \quad (36)$$

Sampling Strategy based on Random Walk(SSRW) [3]. SSRW uses a random walk based method to select next pairs. It defines a matrix \bar{Y} , where each of its element $\bar{y}_{i,j} (\bar{y}_{i,j} = \frac{M_{i,j}}{M_{i,j} + M_{j,i}})$ is the estimated probability that o_i is preferred to o_j . The real probabilities for pairwise comparisons are represented by matrix Y , and the accurate value for matrix Y is unknown. To utilize random walk to address the top- k problem, the method constructs a stochastic matrix \bar{S} for \bar{Y} . Each element in \bar{S} is computed as $\bar{s}_{i,j} = \frac{\bar{y}_{i,j}}{\sum_l \bar{y}_{l,i}}$. Besides, the stochastic matrix for Y is defined as S and $s_{i,j} = \frac{y_{i,j}}{\sum_l y_{l,i}}$. Assuming the principal eigenvectors for \bar{S} and S are \bar{v} and v . The method makes a statement that

$$\|v - \bar{v}\|_{max} \leq \|S - \bar{S}\|_1 \|\bar{A}^\# \|_{max} \quad (37)$$

where $\bar{A}^\# = (I - \bar{S} + 1\bar{v}^T)^{-1} - 1\bar{v}^T$. And $\|\bar{A}^\# \|_{max}$ is converged to $\|A^\# \|_{max}$. As $\|A^\# \|_{max}$ is bounded, $\|S - \bar{S}\|_1 \leq \frac{n}{3} \arg \max_{i,j} c_{i,j} \sum_l \bar{y}_{l,i}$. To minimize Equation 37, a sampling strategy is used to minimize $\|S - \bar{S}\|_1$, and it selects pairs $(i, j) = \arg \max_{i,j} c_{i,j} \sum_l \bar{y}_{l,i}$, where $c_{ij} = \frac{1}{2(M_{ij} + M_{ji})} \log \frac{2n^2 n_{max}}{q}$.

5.3 Active-Learning Methods

Active-learning methods select next pairs by maximizing the information gain based on the estimated latent scores.

5.3.1 Active Learning

CrowdGauss [16]. The scores can be modeled by a multivariate Gaussian distribution $N(\hat{s}, \mathcal{C})$, where \hat{s} is a $1 \times n$ matrix indicating the score for all the objects (initialized by random values), and \mathcal{C} is the covariance matrix of \hat{s} (\mathcal{C}_{ij} is the value at i -th row and j -th column). In each round of pair selection, the expected information gain g_{ij} for each pair (o_i, o_j) is computed. For pair (o_i, o_j) , $N(\hat{s}^{ij}, \mathcal{C}^{ij})$ is the reevaluation distribution by assuming $o_i \succ o_j$, and p_{ij} is the probability of $o_i \succ o_j$ estimated in the previous iteration. The expected information gain is computed as:

$$g_{ij} = p_{ij} \cdot \text{KL}(N(\hat{s}^{ij}, \mathcal{C}^{ij}), N(\hat{s}, \mathcal{C})) + p_{ji} \cdot \text{KL}(N(\hat{s}^{ji}, \mathcal{C}^{ji}), N(\hat{s}, \mathcal{C})), \quad (38)$$

where $p_{ij} = \Phi\left(\frac{s_i - s_j}{\sqrt{1 + \mathcal{C}_{ii} + \mathcal{C}_{jj} - 2\mathcal{C}_{ij}}}\right)$ and KL is the Kullback-Leibler divergence. At each iteration, it selects the pair with the largest expected information gain and updates \hat{s} and \mathcal{C} .

CrowdBT [4]. The above method does not consider the worker quality. Chen et al. [4] propose an active-learning method by taking into account the worker quality. The score of object o_i is modeled by a Gaussian distribution $N(s_i, \delta_i)$ and the quality of worker w is modeled by a Beta distribution $\text{Beta}(\alpha_w, \beta_w)$. Each time it computes the probability of worker w on $o_i \succ o_j$, denoted by $P(o_i \succ_w o_j)$, as below.

$$P(o_i \succ_w o_j) = \eta_w \frac{e^{s_i}}{e^{s_i} + e^{s_j}} + (1 - \eta_w) \frac{e^{s_j}}{e^{s_i} + e^{s_j}}. \quad (39)$$

It uses a parameter γ to tune the comparison result and worker quality. It assigns (o_i, o_j) to worker w by maximizing the probability of w reporting a correct result on (o_i, o_j) .

5.3.2 Combining Rating and Pairwise Comparisons Combine [20]. Instead of just utilizing pairwise comparisons, Ye et al. [20] propose an active-learning strategy by

combining rating and comparison together. Each time with a budget b , it selects a subset of rating-based questions and some comparison-based questions to maximize the expected information gain Z^* . The rating-based method asks the crowd to assign an object to a rate, and the rating-based information gain is computed as below.

$$gr(i) = \mathbb{E}\left(\sum_{r=1}^M p_{ir} \log\left(\frac{p_{ir}}{p(x_i = r)}\right)\right) \quad (40)$$

where p_{ir} is the probability that the crowd assigns o_i with rate r . $p(x_i = r)$ is the prior probability for o_i with rate r .

The comparison based information gain is computed as

$$gc(i, j) = \mathbb{E}\left(p_{ij} \log\left(\frac{p_{ij}}{p(x_{ij} = 1)}\right) + q_{ij} \log\left(\frac{q_{ij}}{p(x_{ij} = 0)}\right)\right), \quad (41)$$

where p_{ij} is the probability of $o_i \succ o_j$ based on crowdsourced comparisons, $q_{ij} = 1 - p_{ij}$, $p(x_{ij} = 1)(p(x_{ij} = 0))$ is the prior probability of $o_i \succ o_j$ ($o_j \succ o_i$). Assuming the cost for rating (comparison) is C_r (C_c), which are set to 1, it selects tasks with the maximal information gain.

6. EXPERIMENTAL STUDY

We evaluated all inference and selection methods. All the source code and real datasets were available at <http://dbgroup.cs.tsinghua.edu.cn/ligl/crowdtopk/>.

6.1 Experimental Setting

6.1.1 Real Experiments

Existing works selected datasets based on the following characters: difficulty (easy or hard for workers), task types (images or texts), objective/subjective. We utilized these features to select datasets and Table 3 showed the details. **PeopleAge**¹. It included 50 human photos with ages from 50 to 100. We asked the crowd to judge which one was younger. **PeopleNum [12]**. It contained 39 images taken in a mall. Each image contained multiple people, where the number of people varied from 13 to 53. We asked the crowd to judge which one contained more people. **EventTime**. It contained 100 history events, e.g., “Germany invades Poland starting World War II” and “Napoleon defeated at Waterloo”. Each event happened in different years. The crowd was asked to judge which event happened earlier. **ImageClarity**. It contained 100 images with different clarities. The crowd was asked to judge which image was clearer.

The datasets had different levels of difficulties. **ImageClarity** was the easiest, and **EventTime** was the hardest. **PeopleAge** and **PeopleNum** had medium difficulty. **EventTime** used textual tasks while **ImageClarity**, **PeopleAge** and **PeopleNum** used image tasks. **PeopleAge** was subjective because different workers might have different judgements on people age while **EventTime**, **PeopleNum** and **ImageClarity** were objective as the tasks had clear answers.

We conducted the real experiments on a real crowdsourcing platform, CrowdFlower(www.crowdflower.com). The price of each microtask was 0.05\$. To obtain high-quality workers, we used 5 qualification tests on each dataset.

6.1.2 Simulation Experiments

To further evaluate different methods on larger datasets, we generated a dataset with 1000 objects with scores from 1 to 1000 and there were totally 499,500 pairs. (Note that if we used 10000 objects, there would be 49,995,000 pairs and most of algorithms cannot support such large dataset.) We varied the worker accuracy in {60%, 70%, 80%, 90%}.

¹<http://www.edouardjanssens.com/art/1-to-100-years/men/>




Figure 3: Recall on Simulation for Selected Inference Methods: Varying Selection Rate (1000 Objects, $k = 10$).




Figure 4: ACC^k on Simulation for Selected Inference Methods: Varying Selection Rate (1000 Objects, $k = 10$).

Table 3: Worker Accuracy on Real Datasets (c-Accuracy: Comparison; r-Accuracy: Rating).

Datasets	#Objects	#Pairs	c-Accuracy	#Ratings	r-Accuracy
PeopleAge	50	1225	85.2%	50	100%
PeopleNum	39	741	88.4%	39	71.8%
EventTime	100	4950	76.6%	100	89.4%
ImageClarity	100	4950	98.6%	100	97%

6.1.3 Setting

Microtasks. For comparison-based methods, we generated comparison tasks to compare each object pair. For hybrid methods that combined rating and comparison, we used two types of tasks: comparison and rating. The rating task asked the worker to select a category for each object. We used five categories for each dataset. For example, the five categories on the **PeopleAge** dataset were ‘50-60’, ‘60-70’, ‘70-80’, ‘80-90’, and ‘90-100’. Each task was assigned to 3 workers. For comparison, we combined the comparison results of the 3 workers using majority vote and computed voted comparison accuracy, denoted by c-accuracy. For rating, we used the majority vote to combine the result and computed voted rating accuracy, denoted by r-accuracy.

Quality Metrics. We used two metrics to evaluate the quality. Let A^k and T^k respectively denote the top- k results by an algorithms and the real top- k results. (1) Recall. It was computed by $\frac{|A^k \cap T^k|}{k}$. (2) ACC^k . It considered the orders of the returned results and was computed as

$$ACC^k = \frac{\sum_{o_i, o_j \in T^k \cap A^k} \mathbb{I}(a_i < a_j \wedge t_i < t_j)}{k(k+1)/2} \quad (42)$$

where t_i was the real ranking position for object o_i and a_i was the ranking position by the algorithm. \mathbb{I} was a function, which returned 1 if the condition was true; 0 otherwise.

Setting. All algorithms were implemented by Python 2.7.6. The parameters were used the same as in the original paper as shown in Table 2. All experiments were conducted on a MacBook Pro with 2.3GHz Intel i7 CPU and 16GB RAM.

6.2 Evaluation of Inference Methods

The number of crowdsourced pairs could affect the quality of inference methods. We varied the selection rate which was the ratio of the number of selected pairs to the total number of pairs, to compare the inference algorithms.

6.2.1 Simulation – Varying Selection Rate

We first compared the inference methods by randomly selecting crowdsourced pairs, and then evaluated different

selection strategies in Section 6.3. Here we showed the results of selected representative methods from each category and the results of all the methods were in the full version². For the heuristic-based algorithms, we selected Local, URE, SSRW, and Iterative as they achieved higher performance. For the machine-learning methods, we selected CrowdBT which outperformed other methods. For hybrid methods, we selected Hybrid. We also compared with TwoStageHeap. Thus we showed the results of these seven algorithms. Figure 3 showed the recall and Figure 4 showed ACC^k .

We had the following observations. First, for low worker quality (e.g., the worker accuracy of 60%), all the algorithms had low recall. For example, in Figure 3(a), the recall for most of the algorithms were smaller than 60%, because for low worker accuracy, there were many incorrect comparisons and it was hard to correct many errors.

Second, for high worker quality (e.g., the worker quality larger than 70%), the machine-learning methods outperformed most of the heuristic-based methods, as the machine-learning methods utilized the global comparison results to infer the top- k answers while the heuristic-based approaches only used the local comparison results, e.g., in-degree and out-degree. If the comparison results on local pairs were incorrect, they would significantly affect the performance while the machine-learning methods tolerated local errors.

Third, among the heuristic-based methods, the iterative inference algorithms (e.g., Iterative, AdaptiveReduce) were better than the other heuristic inference algorithms (e.g., BordaCount, Copeland, Indegree, ELO), and even outperformed the machine-learning algorithms, because the iterative algorithms iteratively reduced objects that were unlikely in the top- k answers. In addition, Local was better than BordaCount and Copeland, as Local utilized 2-hop neighbors while BordaCount and Copeland only utilized the direct neighbors.

Fourth, TwoStageHeap was worse than others when the worker accuracy was low, while it had a good performance when the accuracy was high. The main reason was that for low worker quality, it might miss some top- k candidates in the first step due to the comparison errors from the crowd. URE had a better performance than BRE when the average worker accuracy was above 70%. URE was more robust than BRE as URE computed the number of losses for each object, rather than computing the relative difference of the number

²<http://dbgroup.cs.tsinghua.edu.cn/ligl/crowdtopk/topk.pdf>




Figure 5: Recall on Real Datasets for Selected Inference Methods: Varying Selection Rate ($k = 10$).




Figure 6: ACC^k on Real Datasets for Selected Inference Methods: Varying Selection Rate ($k = 10$).

of wins and losses for each specific object. As the number of pairwise comparisons was not sufficient, the losses for each object could bring more information than the relative difference of the wins and losses. Among the three bound-based algorithms, SSRW and SSSE had poor performance as they utilized local information to update the score matrix. SSRW outperformed SSRW and SSSE, as SSRW utilized random walk to capture global comparison results.

Fifth, for the machine-learning methods, CrowdBT outperformed other methods. CrowdBT had better performance than CrowdGauss, because CrowdBT not only considered the worker quality, but also utilized the graph structure to facilitate the inference. HodgeRank and SSRW had lower recall, because the matrix decomposition had deviations and led to errors to compute latent scores.

Sixth, AdaptiveReduce had lower recall than other methods for small selection rates but achieved comparable and even higher recall for large selection rates. This was because for small selection rates, AdaptiveReduce cannot get high-quality informative set (which required to find a subset with voting account between $\frac{n}{4}$ and $\frac{3n}{4}$), and there were not enough comparisons for small rates; for large selection rates, AdaptiveReduce identified a high-quality informative set based on lots of comparison results. PathRank had poor recall because it required that the graph was strongly connected and had no circles, which made it not robust.

Seventh, for the hybrid algorithms, Hybrid achieved higher performance for a low selection rate, because it greatly reduced the candidates with little cost and utilized global pairwise comparisons in the second stage to infer the top- k answers. Hybrid achieved a good performance for high worker accuracy, because its quality was greatly affected by the worker accuracy and if the worker accuracy was low, its first stage could not guarantee that all the top- k objects were selected into the candidates. Combine was worse than Hybrid, as Combine did not consider worker accuracy. HybridMPR was worse than Hybrid because its estimated scores were worse than the scores learned by Hybrid.

Eighth, with the increase of the selection rates, the comparison algorithms achieved higher recall, as more crowd-sourced pairs were used to infer the final answers. However for the hybrid method, the quality first increased and then decreased, because for a larger selection rate, it involved more candidates and the second phase had to consider more

pairs to infer the top- k results. The results were also consistent with those in the original paper [12].

Ninth, ACC^k was smaller than recall, as ACC^k considered the orders of returned results. The methods computing scores for each object were better than the adaptive-reduce methods (e.g., AdaptiveReduce, PathRank, TwoStageHeap), because the former used the score to infer an order for the objects while the latter could not get any order.

Tenth, with the increase of selection rates, the recall increased, but the improvement after 30% selection rates was not significant. In other words, 30% selection rates had similar recall with 100% selection rates. Thus if we had a monetary budget, we could select 30% pairs.

Although machine-learning methods had higher recall than some heuristic-based algorithms, the superiority was not significant and they were even worse than Iterative, as some heuristics also considered global information. The heuristics achieved as high recall as the machine-learning methods if they effectively utilized global comparison results.

6.2.2 Real Experiments – Varying Selection Rate

We compared different algorithms on real datasets by varying selection rates. Figures 5 and 6 illustrated the results for the selected algorithms. We had the following observations.

First, the observations on the simulation experiments were still true on the real datasets. On the **EventTime** dataset, CrowdBT, CrowdGauss, TrueSkill had similar performance, which outperformed other methods. Among all heuristic-based methods, URE and SSRW had the highest recall, as URE could tolerate more noise and SSRW could utilize global information based on random walk. Second, on the **PeopleAge** dataset, URE was slightly superior to the other methods, because if the noise mainly came from the dataset (as workers had diverse judgements on people’s ages), the simple heuristics could have a competitive performance with the machine-learning methods. Third, on the **ImageClarity** dataset, the algorithms based on random walk and maximum likelihood estimation had much better recall than the heuristic-based methods. For AdaptiveReduce, when the selection rate was low, it got low recall; however with a high selection rate, its performance was significantly improved, because if more pairwise results were collected, AdaptiveReduce could utilize more comparison results to do effective pruning. Fourth, on the **PeopleNum** and **ImageClarity** datasets, as the workers’ accuracy was high, all the algorithms could get higher recall.




Figure 7: Recall on Real Datasets for Selected Inference Methods: Varying k .




Figure 8: Recall on Real Datasets for Selected Inference Methods: Varying Object No ($k = 10$).

For example, on the *ImageClarity* dataset, the recall could reach 90%, even for small rates. Fifth, the hybrid method achieved higher quality for small selection rates, as it used rating to eliminate lowly ranked objects with little cost and utilized the pairwise comparisons to compare highly ranked objects. With the increase of the selection rate, more pairs would be selected into the second stage, and more noise were brought by these pairs. Sixth, for subjective tasks, e.g., *PeopleAge*, we found that the quality could not be improved even with a very high selection rate, because workers had diverse judgements on images with close ages. For objective tasks, the quality was consistently improved with the increase of selection rates.

6.2.3 Real Datasets – Varying k

We varied the number of reported results, k . Figure 7 showed the results. We found that a large k had higher recall than a small k , as the crowd usually returned inaccurate results for hard comparison pairs (with close scores) and correct results for easy comparison pairs (with significantly different scores). The ratio of incorrect pairs to k decreased as k increased. Thus a small k was harder than a large k .

6.2.4 Real Datasets – Varying Object No

We varied the number of objects. Figure 8 showed the results. We found that with the increase of the object number, the recall of some algorithms decreased, because they had to consider more pairs and it became harder to get top- k results from a larger dataset (than a smaller dataset). This was consistent with the case of decreasing k .

6.2.5 Efficiency And Scalability

We evaluated the efficiency of different algorithms. We only reported the time used to infer the top- k answers and the time spent by workers was not included. We generated 1000 objects as discussed in Section 6.1.2. Figure 9 showed the results. First, the heuristics had higher efficiency and scalability than the machine-learning algorithms, because the latter involved huge computations of utilizing the global comparison results to infer latent scores. Second, the local inference heuristics had better efficiency and scalability than the global inference heuristics. Third, with the increase of selection rates, the time increased as they considered more pairs. The machine-learning methods slightly increased as they considered all pairs for every selection rate.

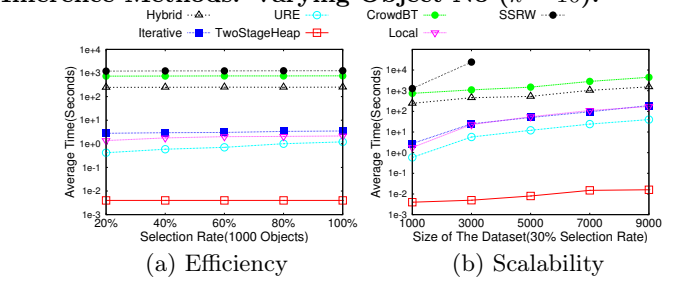


Figure 9: Efficiency/Scalability for Inference.

6.2.6 Takeaways: Inference Algorithm Suggestions

We had the following suggestions that can guide practitioners to select appropriate inference algorithms.

- (1) If the user preferred high quality with small selection rates, the Hybrid method was recommended; If the user preferred a higher quality with larger selection rates, we recommended the machine-learning algorithms, e.g., CrowdBT.
- (2) If the user preferred high quality with high efficiency, we recommended iterative inference heuristics, e.g., Iterative.
- (3) If efficiency was crucial, we recommended TwoStageHeap.
- (4) We recommended 30% selection rate.

6.3 Evaluation on Selection Methods

6.3.1 Real Datasets – Varying Selection Rate

We first varied the selection rates. For each selection rate, we utilized different methods to select pairs and adopted different inference algorithms to infer the top- k answers based on comparison results on these selected pairs. Figures 10-13 showed the results for using CrowdBT and Iterative as inference methods. We had the following observations.

First, the active-learning methods had higher quality than random selection on all the datasets, because they judiciously selected the pairs with high possibility in the top- k results, which were better than the randomly chosen ones. Second, the four heuristics were even worse than the random method, as these four methods were specifically devised to select the top-1 object and cannot effectively find the top- k objects. Even for large selection rates, they still could not achieve high quality, as they selected many duplicated pairs. Third, SSCO, SSSE and SSRW were slightly better than Random, because they could use the estimated scores to remove some unnecessary pairs but they could not estimate an accurate bound with limited number of comparisons. Fourth, Combine achieved high quality as it utilized




Figure 10: Recall of Selection Method (Inference Method: CrowdBT).




Figure 11: ACC^k of Selection Methods (Inference Method: CrowdBT).




Figure 12: Recall of Selection Methods (Inference Method: Iterative).




Figure 13: ACC^k of Selection Methods (Inference Method: Iterative).

the rating tasks to select high-quality rating or comparison tasks. Fifth, the recall of the selection methods became stable when the selection rates were larger than 30% on all the datasets, because the selection strategy was not important for large selection rates as they could utilize more pairs. Sixth, on the *PeopleNum* and *ImageClarity* datasets, the recall of all the methods were high even with a low selection rate due to the high worker accuracy on these datasets.

6.3.2 Real Datasets – Varying k

We varied the number of reported results, k . We used CrowdBT as the inference method. Figure 14 showed the results. With a larger k , the average recall was higher, because a large k made the problem easy as the comparisons on many distinctively different objects were easier than the comparisons on few similar objects. In addition, the active-learning methods and bound-based methods still outperformed Random, which was better than the heuristics.

6.3.3 Real Datasets – Varying Object No

We varied the number of objects and still used CrowdBT as the inference method. The selection rate was 30%. Figure 15 showed the results. With the increase of the object number, the average recall was slightly decreased, because more objects were involved, which made the problem harder. This was also similar to the case of increasing k .

6.3.4 Efficiency and Scalability

We evaluated the efficiency and scalability of the selection algorithms. In each iteration, we selected 10 pairs. Figure 16 showed the results. The four heuristics had much higher efficiency and scalability than other methods, as the heuristics only sorted the objects based on their scores and returned the pairs with the highest scores. The active-learning algorithms had the worst efficiency and scalability, because they had to evaluate all the unselected pairs. Specifically, CrowdGauss involved huge matrix computation in each iteration. CrowdBT enumerated all possible combinations of tasks and workers. Combine did not consider workers and identified the tasks with the largest information gain. Among the three bound-based methods, SSCO and SSSE had higher efficiency as they only found the rows with high bounds, while SSRW took longer time than SSCO and SSSE because SSRW evaluated all the unselected pairs.

6.3.5 Takeaways: Selection Algorithm Suggestions

We had the following suggestions that can guide practitioners to select appropriate selection algorithms. (1) If the user preferred high quality, the active-learning method CrowdBT was recommended. (2) If the user preferred high efficiency, we recommended SSCO/SSSE. (3) For large selection rates ($>30\%$), the random method was acceptable.




Figure 14: Recall on Real Datasets for Selection Methods: Varying k (Inference Method: CrowdBT).




Figure 15: Recall on Real Datasets for Selection Methods: Varying Object No ($k = 10$, Inference: CrowdBT).

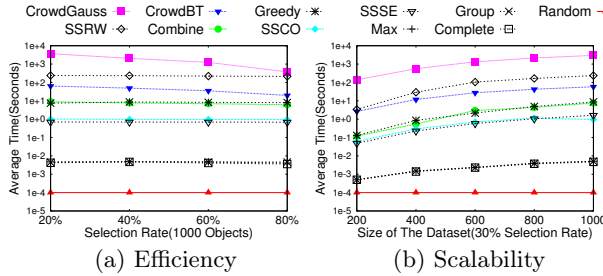


Figure 16: Efficiency/Scalability for Selection.

7. CONCLUSION

We provided a detailed survey on the crowdsourced top- k problem, discussed all of the comparison-based algorithms, rating-based algorithms and hybrid algorithms. Based on the experimental results, we also provided guidelines on selecting appropriate algorithms for various scenarios.

- (1) The inference and selection algorithms were important to achieve high quality, especially for small selection rates.
- (2) For result inference, the machine-learning methods achieved high quality. The global inference heuristics that utilized global comparison results achieved comparable and even higher quality than the machine-learning methods. The local inference heuristics had poor quality. However, the heuristics achieved higher efficiency and scalability. We had the following suggestions to select appropriate algorithms. (i) If the user preferred high efficiency, we recommended **TwoStage-Heap**. (ii) If the user preferred high quality with low selection rates, we recommended **Hybrid**; If the user preferred a high quality with large selection rates, we recommended the machine-learning algorithms, e.g., **CrowdBT**; (iii) If the user preferred high quality with acceptable efficiency, we recommended the iterative inference heuristics, e.g., **Iterative**.
- (3) For pair selection, the active-learning methods and bound-based methods achieved higher quality than heuristics. However the active-learning methods had lower efficiency and cannot meet the online requirement. We had the following suggestions. (i) We recommended 30% selection rate; (ii) If the user preferred high efficiency, we recommended **SSCO/SSSE**; (iii) If the user preferred quality, the active-learning method **CrowdBT** was recommended. (iv) For large selection rates, the random method was acceptable.

8. REFERENCES

- [1] R. M. Adelman and A. B. Whinston. Sophisticated voting with information for two voting functions. *Journal of Economic Theory*, 15(1):145–159, 1977.

- [2] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, pages 324–345, 1952.
- [3] R. Busa-Fekete, B. Szorenyi, W. Cheng, P. Weng, and E. Hullermeier. Top- k selection based on adaptive sampling of noisy preferences. In *ICML*, pages 1094–1102, 2013.
- [4] X. Chen, P. N. Bennett, K. Collins-Thompson, and E. Horvitz. Pairwise ranking aggregation in a crowdsourced setting. In *WSDM*, pages 193–202, 2013.
- [5] S. B. Davidson, S. Khanna, T. Milo, and S. Roy. Using the crowd for top- k and group-by queries. In *ICDT*, pages 225–236, 2013.
- [6] A. E. Elo. *The rating of chessplayers, past and present*, volume 3. Batsford London, 1978.
- [7] B. Eriksson. Learning to top- k search using pairwise comparisons. In *AISTATS*, pages 265–273, 2013.
- [8] U. Feige, P. Raghavan, D. Peleg, and E. Upfal. Computing with noisy information. *SIAM J. Comput.*, 23(5):1001–1018, 1994.
- [9] S. Guo, A. G. Parameswaran, and H. Garcia-Molina. So who won?: dynamic max discovery with the crowd. In *SIGMOD*, pages 385–396, 2012.
- [10] R. Herbrich, T. Minka, and T. Graepel. Trueskill: A bayesian skill rating system. In *NIPS*, pages 569–576, 2006.
- [11] X. Jiang, L. Lim, Y. Yao, and Y. Ye. Statistical ranking and combinatorial hodge theory. *Math. Program.*, 127(1):203–244, 2011.
- [12] A. R. Khan and H. Garcia-Molina. Hybrid strategies for finding the max with the crowd. Technical report, 2014.
- [13] G. Li, J. Wang, Y. Zheng, and M. Franklin. Crowdsourced data management: A survey. *IEEE Trans. Knowl. Data Eng.*, 2016.
- [14] A. Marcus, E. Wu, D. R. Karger, S. Madden, and R. C. Miller. Human-powered sorts and joins. *PVLDB*, 5(1):13–24, 2011.
- [15] S. Negahban, S. Oh, and D. Shah. Iterative ranking from pair-wise comparisons. In *NIPS*, pages 2483–2491, 2012.
- [16] T. Pfeiffer, X. A. Gao, Y. Chen, A. Mao, and D. G. Rand. Adaptive polling for information aggregation. In *AAAI*, 2012.
- [17] J.-C. Pomerol and S. Barba-Romero. *Multicriterion decision in management: principles and practice*, volume 25. Springer Science & Business Media, 2012.
- [18] P. Venetis, H. Garcia-Molina, K. Huang, and N. Polyzotis. Max algorithms in crowdsourcing environments. In *WWW*, pages 989–998, 2012.
- [19] F. L. Wauthier, M. I. Jordan, and N. Jojic. Efficient ranking from pairwise comparisons. In *ICML*, pages 109–117, 2013.
- [20] P. Ye and D. Doermann. Combining preference and absolute judgements in a crowd-sourced setting. In *Proc. of Intl. Conf. on Machine Learning*, pages 1–7, 2013.

APPENDIX

A. EXPERIMENTAL RESULTS FOR RANKING

We showed the experimental results for top- k computations on all algorithms.

(a) c-Accuracy=63.85% r-Accuracy=66.9%.							(b) c-Accuracy=77.45% r-Accuracy=73.9%.						
Top-k Algorithms		Sample Rate					Top-k Algorithms		Sample Rate				
		20%	40%	60%	80%	100%			20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.12	0.16	0.2	0.23	0.3	Heuristic (Local)	BordaCount	0.18	0.24	0.3	0.36	0.4
	Copeland	0.18	0.24	0.28	0.32	0.35		Copeland	0.27	0.35	0.4	0.43	0.42
	Local	0.14	0.18	0.2	0.22	0.24		Local	0.25	0.32	0.38	0.43	0.46
	Indegree	0.13	0.18	0.18	0.18	0.2		Indegree	0.24	0.3	0.34	0.38	0.4
	ELO	0.17	0.19	0.2	0.18	0.19		ELO	0.27	0.29	0.3	0.29	0.3
	BRE	0.17	0.24	0.29	0.3	0.3		BRE	0.22	0.3	0.35	0.41	0.45
	URE	0.15	0.2	0.25	0.27	0.37		URE	0.26	0.3	0.34	0.37	0.43
	SSCO	0.15	0.19	0.23	0.26	0.27		SSCO	0.19	0.23	0.25	0.29	0.31
Heuristic (Global)	SSSE	0.13	0.16	0.21	0.25	0.3	Heuristic (Global)	SSSE	0.15	0.19	0.22	0.25	0.35
	Iterative	0.3	0.44	0.54	0.62	0.67		Iterative	0.6	0.73	0.8	0.84	0.88
	PathRank	0.01	0.0	0.01	0.01	0.02		PathRank	0.01	0.01	0.01	0.01	0.0
	AdaptiveReduce	0.01	0.16	0.16	0.39	0.65		AdaptiveReduce	0.01	0.25	0.28	0.76	0.83
	MPageRank	0.15	0.24	0.29	0.38	0.38		MPageRank	0.33	0.42	0.57	0.59	0.7
	RankCentrality	0.15	0.21	0.26	0.33	0.35		RankCentrality	0.22	0.35	0.45	0.46	0.55
ML	SSRW	0.12	0.13	0.13	0.1	0.2	ML	SSRW	0.28	0.32	0.37	0.38	0.51
	TrueSkill	0.15	0.25	0.29	0.35	0.45		TrueSkill	0.25	0.4	0.45	0.55	0.62
	CrowdBT	0.23	0.42	0.51	0.55	0.64		CrowdBT	0.55	0.75	0.8	0.82	0.9
	CrowdGauss	0.26	0.31	0.35	0.38	0.41		CrowdGauss	0.36	0.43	0.5	0.52	0.58
Extension	HodgeRank	0.23	0.28	0.31	0.34	0.37	Extension	HodgeRank	0.47	0.52	0.56	0.61	0.63
	TwoStageHeap	0.04	0.12	0.13	0.15	0.25		TwoStageHeap	0.17	0.22	0.34	0.41	0.58
Rate + Compare	Combine	0.12	0.2	0.23	0.25	0.3	Rate + Compare	Combine	0.15	0.24	0.27	0.3	0.35
	Hybrid	0.5	0.33	0.33	0.3	0.26		Hybrid	0.7	0.64	0.6	0.6	0.58
	HybridMPR	0.5	0.42	0.32	0.3	0.3		HybridMPR	0.72	0.61	0.58	0.56	0.4
(c)c-Accuracy=88.70% r-Accuracy=83.5%.							(d)c-Accuracy=96.43% r-Accuracy=91.9%.						
Top-k Algorithms		Sample Rate					Top-k Algorithms		Sample Rate				
		20%	40%	60%	80%	100%			20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.22	0.3	0.38	0.47	0.63	Heuristic (Local)	BordaCount	0.24	0.33	0.41	0.51	0.78
	Copeland	0.32	0.42	0.48	0.53	0.55		Copeland	0.37	0.48	0.58	0.69	0.8
	Local	0.32	0.42	0.5	0.58	0.65		Local	0.35	0.47	0.55	0.65	0.85
	Indegree	0.22	0.34	0.36	0.48	0.5		Indegree	0.33	0.41	0.53	0.63	0.9
	ELO	0.39	0.48	0.51	0.54	0.51		ELO	0.47	0.57	0.65	0.71	0.75
	BRE	0.31	0.36	0.42	0.48	0.61		BRE	0.37	0.49	0.6	0.7	0.86
	URE	0.38	0.46	0.51	0.56	0.6		URE	0.62	0.7	0.75	0.79	0.77
	SSCO	0.19	0.25	0.3	0.35	0.41		SSCO	0.22	0.3	0.36	0.43	0.56
Heuristic (Global)	SSSE	0.18	0.26	0.29	0.34	0.42	Heuristic (Global)	SSSE	0.24	0.3	0.38	0.44	0.6
	Iterative	0.75	0.85	0.89	0.92	0.95		Iterative	0.81	0.89	0.93	0.96	0.98
	PathRank	0.01	0.01	0.01	0.01	0.01		PathRank	0.03	0.01	0.01	0.0	0.0
	AdaptiveReduce	0.01	0.35	0.73	0.89	0.93		AdaptiveReduce	0.01	0.34	0.78	0.93	0.93
	MPageRank	0.56	0.69	0.72	0.8	0.9		MPageRank	0.67	0.8	0.85	0.93	1.0
	RankCentrality	0.36	0.46	0.56	0.62	0.8		RankCentrality	0.4	0.51	0.65	0.71	0.8
ML	SSRW	0.36	0.42	0.46	0.57	0.8	ML	SSRW	0.48	0.64	0.66	0.68	0.8
	TrueSkill	0.36	0.47	0.59	0.67	0.75		TrueSkill	0.45	0.5	0.65	0.73	0.78
	CrowdBT	0.8	0.8	0.92	0.95	0.98		CrowdBT	0.9	0.9	0.96	0.98	1.0
	CrowdGauss	0.41	0.52	0.6	0.66	0.72		CrowdGauss	0.44	0.55	0.65	0.75	0.89
Extension	HodgeRank	0.49	0.58	0.64	0.66	0.7	Extension	HodgeRank	0.59	0.69	0.72	0.77	0.8
	TwoStageHeap	0.45	0.53	0.68	0.72	0.71		TwoStageHeap	0.73	0.74	0.82	0.82	0.89
Rate + Compare	Combine	0.3	0.4	0.45	0.56	0.6	Rate + Compare	Combine	0.35	0.52	0.62	0.65	0.7
	Hybrid	0.8	0.75	0.7	0.6	0.56		Hybrid	0.9	0.87	0.85	0.8	0.77
	HybridMPR	0.74	0.72	0.65	0.62	0.6		HybridMPR	0.8	0.76	0.76	0.72	0.7

Figure 17: Recall on Simulation for All Inference Methods: Varying Selection Rate (1000 Objects).

(a) c-Accuracy=63.85% r-Accuracy=66.9%.

Top-k Algorithms		Sample Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.0	0.01	0.01	0.02	0.02
	Copeland	0.01	0.02	0.04	0.05	0.07
	Local	0.01	0.01	0.01	0.02	0.01
	Indegree	0.01	0.01	0.01	0.01	0.0
	ELO	0.01	0.01	0.01	0.01	0.0
	BRE	0.01	0.02	0.03	0.03	0.03
	URE	0.01	0.01	0.02	0.02	0.03
	SSCO	0.01	0.01	0.01	0.02	0.02
	SSSE	0.0	0.01	0.02	0.03	0.04
Heuristic (Global)	Iterative	0.04	0.08	0.12	0.16	0.18
	PathRank	0.0	0.0	0.0	0.0	0.0
	AdaptiveReduce	0.0	0.01	0.01	0.06	0.2
	MPageRank	0.01	0.01	0.01	0.02	0.04
	RankCentrality	0.0	0.02	0.04	0.05	0.05
	SSRW	0.0	0.01	0.02	0.05	0.06
ML	TrueSkill	0.0	0.01	0.02	0.02	0.03
	CrowdBT	0.02	0.04	0.06	0.09	0.13
	CrowdGauss	0.01	0.02	0.03	0.04	0.06
	HodgeRank	0.01	0.01	0.03	0.05	0.08
Extension	TwoStageHeap	0.01	0.02	0.02	0.02	0.03
Rate + Compare	Combine	0.01	0.03	0.05	0.04	0.06
	Hybrid	0.15	0.13	0.03	0.02	0.03
	HybridMPR	0.14	0.14	0.05	0.02	0.02

(b) c-Accuracy=77.45% r-Accuracy=73.9%.

Top-k Algorithms		Sample Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.01	0.02	0.03	0.05	0.08
	Copeland	0.03	0.05	0.06	0.07	0.07
	Local	0.03	0.04	0.06	0.07	0.07
	Indegree	0.04	0.06	0.1	0.12	0.09
	ELO	0.03	0.03	0.04	0.04	0.02
	BRE	0.02	0.03	0.05	0.06	0.07
	URE	0.02	0.03	0.04	0.04	0.05
	SSCO	0.02	0.04	0.06	0.06	0.05
	SSSE	0.01	0.01	0.02	0.02	0.06
Heuristic (Global)	Iterative	0.14	0.22	0.26	0.29	0.31
	PathRank	0.0	0.0	0.0	0.0	0.0
	AdaptiveReduce	0.0	0.03	0.03	0.31	0.4
	MPageRank	0.08	0.11	0.15	0.22	0.22
	RankCentrality	0.02	0.05	0.08	0.08	0.25
	SSRW	0.04	0.04	0.05	0.06	0.24
ML	TrueSkill	0.02	0.02	0.02	0.03	0.06
	CrowdBT	0.04	0.05	0.2	0.28	0.42
	CrowdGauss	0.04	0.03	0.05	0.09	0.12
	HodgeRank	0.03	0.05	0.03	0.1	0.11
Extension	TwoStageHeap	0.05	0.06	0.08	0.07	0.08
Rate + Compare	Combine	0.03	0.04	0.05	0.08	0.1
	Hybrid	0.27	0.25	0.24	0.2	0.12
	HybridMPR	0.2	0.15	0.14	0.13	0.13

(c) c-Accuracy=88.70% r-Accuracy=83.5%.

Top-k Algorithms		Sample Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.02	0.03	0.06	0.09	0.2
	Copeland	0.04	0.06	0.09	0.13	0.16
	Local	0.04	0.07	0.1	0.15	0.14
	Indegree	0.04	0.06	0.08	0.14	0.18
	ELO	0.03	0.09	0.1	0.07	0.08
	BRE	0.03	0.05	0.06	0.07	0.06
	URE	0.05	0.08	0.11	0.14	0.2
	SSCO	0.02	0.04	0.05	0.09	0.1
	SSSE	0.01	0.02	0.03	0.03	0.08
Heuristic (Global)	Iterative	0.22	0.29	0.33	0.33	0.36
	PathRank	0.0	0.0	0.0	0.0	0.0
	AdaptiveReduce	0.0	0.04	0.26	0.45	0.5
	MPageRank	0.13	0.18	0.27	0.35	0.44
	RankCentrality	0.04	0.08	0.14	0.16	0.35
	SSRW	0.01	0.12	0.1	0.15	0.36
ML	TrueSkill	0.04	0.08	0.09	0.11	0.15
	CrowdBT	0.29	0.4	0.45	0.48	0.53
	CrowdGauss	0.06	0.07	0.13	0.08	0.22
	HodgeRank	0.05	0.12	0.16	0.21	0.16
Extension	TwoStageHeap	0.16	0.23	0.25	0.28	0.35
Rate + Compare	Combine	0.08	0.11	0.14	0.15	0.21
	Hybrid	0.33	0.32	0.27	0.25	0.16
	HybridMPR	0.32	0.25	0.16	0.15	0.14

(d) c-Accuracy=96.43% r-Accuracy=91.9%.

Top-k Algorithms		Sample Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.02	0.04	0.07	0.11	0.35
	Copeland	0.05	0.09	0.14	0.22	0.32
	Local	0.05	0.09	0.13	0.19	0.4
	Indegree	0.06	0.12	0.14	0.26	0.42
	ELO	0.08	0.17	0.19	0.29	0.25
	BRE	0.05	0.09	0.14	0.2	0.3
	URE	0.13	0.22	0.26	0.31	0.29
	SSCO	0.02	0.03	0.05	0.06	0.1
	SSSE	0.03	0.06	0.09	0.1	0.12
Heuristic (Global)	Iterative	0.27	0.32	0.36	0.37	0.39
	PathRank	0.0	0.0	0.0	0.0	0.0
	AdaptiveReduce	0.0	0.04	0.31	0.59	0.63
	MPageRank	0.24	0.36	0.5	0.54	0.65
	RankCentrality	0.07	0.11	0.18	0.23	0.35
	SSRW	0.16	0.19	0.23	0.23	0.24
ML	TrueSkill	0.1	0.12	0.21	0.25	0.27
	CrowdBT	0.32	0.51	0.55	0.68	0.77
	CrowdGauss	0.06	0.09	0.12	0.19	0.29
	HodgeRank	0.17	0.24	0.35	0.35	0.36
Extension	TwoStageHeap	0.27	0.3	0.35	0.39	0.45
Rate + Compare	Combine	0.06	0.08	0.11	0.16	0.24
	Hybrid	0.42	0.37	0.32	0.28	0.2
	HybridMPR	0.38	0.35	0.23	0.22	0.21

Figure 18: ACC^k on Simulation for All Inference Methods: Varying Selection Rate (1000 Objects).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.34	0.4	0.44	0.48	0.6
	Copeland	0.57	0.68	0.75	0.8	0.8
	Local	0.54	0.66	0.72	0.77	0.8
	Indegree	0.32	0.37	0.42	0.45	0.5
	ELO	0.57	0.71	0.78	0.83	0.86
	BRE	0.57	0.68	0.75	0.8	0.8
	URE	0.77	0.84	0.88	0.9	0.9
	SSCO	0.33	0.35	0.37	0.39	0.4
	SSSE	0.33	0.35	0.36	0.38	0.4
Heuristic (Global)	Iterative	0.58	0.67	0.73	0.8	0.9
	PathRank	0.1	0.11	0.11	0.12	0.13
	AdaptiveReduce	0.1	0.46	0.5	0.51	0.75
	MPageRank	0.44	0.53	0.58	0.67	0.7
	RankCentrality	0.56	0.68	0.78	0.8	0.8
	SSRW	0.69	0.79	0.83	0.87	0.88
ML	TrueSkill	0.72	0.83	0.87	0.89	0.9
	CrowdBT	0.77	0.81	0.88	0.89	0.9
	CrowdGauss	0.74	0.85	0.88	0.88	0.9
	HodgeRank	0.64	0.77	0.84	0.86	0.9
Extension	TwoStageHeap	0.43	0.48	0.56	0.63	0.73
Rate + Compare	Combine	0.56	0.57	0.58	0.65	0.7
	Hybrid	0.8	0.67	0.62	0.63	0.6
	HybridMPR	0.75	0.65	0.58	0.57	0.56

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.7	0.78	0.84	0.89	0.9
	Copeland	0.78	0.85	0.89	0.91	0.9
	Local	0.71	0.83	0.88	0.91	0.9
	Indegree	0.67	0.76	0.83	0.88	0.9
	ELO	0.76	0.84	0.88	0.91	0.9
	BRE	0.78	0.85	0.89	0.92	0.9
	URE	0.81	0.86	0.9	0.92	0.9
	SSCO	0.73	0.8	0.84	0.86	0.9
	SSSE	0.74	0.81	0.84	0.86	0.9
Heuristic (Global)	Iterative	0.8	0.87	0.9	0.94	1.0
	PathRank	0.36	0.37	0.38	0.39	0.41
	AdaptiveReduce	0.31	0.83	0.86	0.89	0.9
	MPageRank	0.77	0.84	0.85	0.85	0.8
	RankCentrality	0.77	0.84	0.88	0.89	0.9
	SSRW	0.87	0.88	0.89	0.89	0.9
ML	TrueSkill	0.81	0.88	0.91	0.92	0.9
	CrowdBT	0.84	0.89	0.92	0.91	0.9
	CrowdGauss	0.84	0.9	0.91	0.91	0.9
	HodgeRank	0.75	0.85	0.88	0.9	0.9
Extension	TwoStageHeap	0.8	0.81	0.82	0.85	0.9
Rate + Compare	Combine	0.69	0.82	0.86	0.88	0.89
	Hybrid	0.9	0.9	0.9	0.9	0.9
	HybridMPR	0.9	0.86	0.85	0.83	0.81

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.59	0.65	0.68	0.7	0.7
	Copeland	0.66	0.69	0.7	0.72	0.7
	Local	0.62	0.67	0.7	0.72	0.7
	Indegree	0.56	0.63	0.67	0.7	0.7
	ELO	0.63	0.68	0.7	0.71	0.7
	BRE	0.66	0.69	0.7	0.71	0.7
	URE	0.7	0.72	0.72	0.72	0.7
	SSCO	0.62	0.67	0.7	0.72	0.7
	SSSE	0.62	0.67	0.69	0.72	0.7
Heuristic (Global)	Iterative	0.68	0.7	0.7	0.7	0.7
	PathRank	0.26	0.27	0.28	0.28	0.3
	AdaptiveReduce	0.25	0.68	0.69	0.7	0.7
	MPageRank	0.64	0.68	0.72	0.71	0.7
	RankCentrality	0.64	0.69	0.7	0.7	0.7
	SSRW	0.69	0.7	0.7	0.69	0.7
ML	TrueSkill	0.68	0.71	0.7	0.71	0.7
	CrowdBT	0.68	0.7	0.71	0.71	0.7
	CrowdGauss	0.67	0.71	0.71	0.71	0.7
	HodgeRank	0.61	0.67	0.7	0.72	0.7
Extension	TwoStageHeap	0.68	0.68	0.69	0.68	0.68
Rate + Compare	Combine	0.57	0.65	0.67	0.69	0.7
	Hybrid	0.71	0.7	0.7	0.7	0.68
	HybridMPR	0.7	0.68	0.67	0.65	0.63

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.52	0.62	0.7	0.78	0.9
	Copeland	0.65	0.75	0.83	0.89	1.0
	Local	0.64	0.75	0.82	0.88	1.0
	Indegree	0.49	0.6	0.68	0.76	0.8
	ELO	0.65	0.78	0.86	0.92	0.98
	BRE	0.66	0.75	0.83	0.88	1.0
	URE	0.8	0.87	0.91	0.95	1.0
	SSCO	0.54	0.63	0.71	0.81	0.95
	SSSE	0.54	0.63	0.71	0.8	0.9
Heuristic (Global)	Iterative	0.81	0.88	0.93	0.96	1.0
	PathRank	0.45	0.46	0.48	0.55	0.58
	AdaptiveReduce	0.14	0.76	0.77	0.96	1.0
	MPageRank	0.72	0.82	0.88	0.92	1.0
	RankCentrality	0.65	0.78	0.88	0.93	1.0
	SSRW	0.86	0.93	0.97	0.99	1.0
ML	TrueSkill	0.78	0.89	0.93	0.94	0.99
	CrowdBT	0.84	0.92	0.97	0.99	1.0
	CrowdGauss	0.85	0.95	0.95	0.99	1.0
	HodgeRank	0.73	0.8	0.85	0.9	0.9
Extension	TwoStageHeap	0.79	0.85	0.87	0.88	0.93
Rate + Compare	Combine	0.76	0.9	0.93	0.98	1.0
	Hybrid	1.0	0.98	0.97	0.97	0.96
	HybridMPR	0.96	0.96	0.95	0.94	0.93

Figure 19: Recall on Real Datasets for All Inference Methods: Varying Selection Rate.

(a) **EventTime** (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.04	0.05	0.07	0.1	0.16
	Copeland	0.13	0.21	0.29	0.36	0.43
	Local	0.13	0.2	0.27	0.33	0.4
	Indegree	0.11	0.17	0.23	0.28	0.29
	ELO	0.14	0.23	0.34	0.42	0.49
	BRE	0.13	0.21	0.29	0.36	0.43
	URE	0.25	0.39	0.48	0.53	0.54
	SSCO	0.03	0.03	0.03	0.03	0.02
	SSSE	0.03	0.03	0.03	0.03	0.02
Heuristic (Global)	Iterative	0.13	0.18	0.21	0.25	0.33
	PathRank	0.0	0.0	0.01	0.0	0.0
	AdaptiveReduce	0.0	0.09	0.11	0.13	0.27
	MPageRank	0.07	0.14	0.18	0.18	0.23
	RankCentrality	0.13	0.22	0.31	0.38	0.45
	SSRW	0.23	0.37	0.44	0.49	0.69
ML	TrueSkill	0.24	0.35	0.46	0.5	0.54
	CrowdBT	0.26	0.35	0.43	0.44	0.56
	CrowdGauss	0.26	0.43	0.49	0.55	0.58
	HodgeRank	0.22	0.29	0.37	0.48	0.49
Extension	TwoStageHeap	0.07	0.1	0.12	0.15	0.23
Rate + Compare	Combine	0.13	0.15	0.21	0.33	0.45
	Hybrid	0.33	0.23	0.2	0.19	0.18
	HybridMPR	0.28	0.25	0.2	0.18	0.16

(c) **PeopleNum** (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.19	0.27	0.35	0.44	0.53
	Copeland	0.27	0.38	0.45	0.52	0.53
	Local	0.23	0.35	0.44	0.51	0.51
	Indegree	0.27	0.37	0.45	0.51	0.49
	ELO	0.26	0.38	0.47	0.52	0.54
	BRE	0.27	0.38	0.45	0.52	0.53
	URE	0.28	0.41	0.48	0.54	0.53
	SSCO	0.23	0.32	0.4	0.47	0.55
	SSSE	0.23	0.32	0.4	0.47	0.55
Heuristic (Global)	Iterative	0.26	0.31	0.33	0.36	0.41
	PathRank	0.05	0.05	0.06	0.05	0.05
	AdaptiveReduce	0.04	0.37	0.42	0.49	0.52
	MPageRank	0.2	0.33	0.38	0.44	0.56
	RankCentrality	0.29	0.41	0.48	0.5	0.58
	SSRW	0.31	0.46	0.49	0.5	0.5
ML	TrueSkill	0.32	0.45	0.52	0.55	0.57
	CrowdBT	0.39	0.49	0.55	0.57	0.59
	CrowdGauss	0.34	0.43	0.54	0.61	0.64
	HodgeRank	0.28	0.39	0.46	0.5	0.53
Extension	TwoStageHeap	0.39	0.42	0.45	0.47	0.52
Rate + Compare	Combine	0.29	0.45	0.49	0.52	0.54
	Hybrid	0.53	0.52	0.5	0.49	0.49
	HybridMPR	0.56	0.52	0.51	0.48	0.47

(b) **PeopleAge** (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.13	0.18	0.21	0.25	0.29
	Copeland	0.18	0.23	0.26	0.28	0.29
	Local	0.17	0.23	0.26	0.29	0.27
	Indegree	0.19	0.23	0.26	0.28	0.29
	ELO	0.18	0.23	0.27	0.29	0.3
	BRE	0.19	0.23	0.26	0.28	0.29
	URE	0.19	0.25	0.27	0.29	0.29
	SSCO	0.15	0.19	0.21	0.23	0.24
	SSSE	0.15	0.19	0.21	0.23	0.24
Heuristic (Global)	Iterative	0.18	0.19	0.19	0.19	0.19
	PathRank	0.03	0.03	0.04	0.03	0.04
	AdaptiveReduce	0.02	0.23	0.24	0.28	0.33
	MPageRank	0.15	0.24	0.29	0.3	0.31
	RankCentrality	0.2	0.25	0.27	0.29	0.33
	SSRW	0.27	0.27	0.27	0.27	0.36
ML	TrueSkill	0.21	0.26	0.28	0.3	0.32
	CrowdBT	0.24	0.28	0.3	0.31	0.31
	CrowdGauss	0.22	0.27	0.29	0.29	0.29
	HodgeRank	0.2	0.23	0.26	0.29	0.29
Extension	TwoStageHeap	0.24	0.25	0.27	0.32	0.35
Rate + Compare	Combine	0.12	0.17	0.21	0.24	0.26
	Hybrid	0.3	0.3	0.29	0.28	0.27
	HybridMPR	0.33	0.32	0.29	0.29	0.27

(d) **ImageClarity** (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	0.1	0.16	0.21	0.29	0.46
	Copeland	0.17	0.26	0.35	0.47	0.72
	Local	0.19	0.28	0.38	0.49	0.73
	Indegree	0.18	0.26	0.35	0.48	0.6
	ELO	0.19	0.31	0.42	0.55	0.71
	BRE	0.17	0.26	0.36	0.47	0.73
	URE	0.27	0.42	0.53	0.62	0.7
	SSCO	0.1	0.15	0.2	0.26	0.35
	SSSE	0.1	0.15	0.2	0.26	0.35
Heuristic (Global)	Iterative	0.26	0.31	0.35	0.38	0.4
	PathRank	0.21	0.23	0.24	0.27	0.33
	AdaptiveReduce	0.01	0.29	0.3	0.65	0.75
	MPageRank	0.18	0.37	0.51	0.61	0.78
	RankCentrality	0.21	0.32	0.47	0.59	0.65
	SSRW	0.55	0.6	0.7	0.74	0.76
ML	TrueSkill	0.28	0.46	0.54	0.65	0.72
	CrowdBT	0.4	0.57	0.68	0.76	0.78
	CrowdGauss	0.27	0.56	0.59	0.65	0.77
	HodgeRank	0.17	0.24	0.38	0.53	0.78
Extension	TwoStageHeap	0.44	0.45	0.47	0.53	0.56
Rate + Compare	Combine	0.25	0.51	0.53	0.62	0.65
	Hybrid	0.76	0.75	0.72	0.7	0.7
	HybridMPR	0.69	0.62	0.61	0.6	0.57

Figure 20: ACC^k on Real Datasets for All Inference Methods: Varying Selection Rate.

Top-k Algorithms		Selection Rate (1000 Objects)					Top-k Algorithms		Object No (30% Selection Rate)				
		20%	40%	60%	80%	100%			1000	3000	5000	7000	9000
Heuristic (Local)	BordaCount	0.11	0.2	0.31	0.36	0.44	Heuristic (Local)	BordaCount	0.15	0.54	1.57	4.38	5.95
	Copeland	0.11	0.21	0.31	0.37	0.43		Copeland	0.15	0.53	1.50	4.22	5.89
	Local	1.37	1.78	2.01	2.06	2.15		Local	1.78	22.9	55.0	110	172
	Indegree	2.87	3.55	4.47	5.56	6.46		Indegree	3.47	72.9	130	286	418
	ELO	0.78	2.92	4.67	7.77	7.89		ELO	2.74	19.5	44.8	98.4	144
	BRE	0.6	0.78	0.86	1.19	0.94		BRE	0.76	6.88	16.0	43.3	66.7
	URE	0.42	0.59	0.71	1.01	0.89		URE	0.59	5.73	12.1	24.2	40.0
	SSCO	2.25	2.57	2.8	3.03	3.27		SSCO	2.61	28.8	67.8	118	214
	SSSE	2.95	3.15	3.28	3.22	3.36		SSSE	3.01	18.1	54.2	110	201
Heuristic (Global)	Iterative	2.79	2.96	3.03	3.36	3.5	Heuristic (Global)	Iterative	2.78	24.9	50.5	95.5	191
	PathRank	5.04	6.78	10.4	14.5	17.9		PathRank	5.66	1435	4188	8207	16365
	AdaptiveReduce	4.79	8.33	10.77	15.27	18.14		AdaptiveReduce	6.80	12.4	16.6	41.2	71.3
	MPageRank	22.4	22.5	23.5	25.9	26.6		MPageRank	22.4	65.6	215	363	717
	RankCentrality	18.6	18.5	18.5	19.5	19.2		RankCentrality	17.4	76.4	221	447	742
	SSRW	1199	1223	1229	1244	1251		SSRW	1276	24294	—	—	—
ML	TrueSkill	73	151	200	267	337	ML	TrueSkill	108	813	3462	9639	25168
	CrowdBT	735	742	750	751	757		CrowdBT	745	1095	1487	2815	4460
	CrowdGauss	506	519	534	540	547		CrowdGauss	513	889	1043	2003	3385
	HodgeRank	46.8	47.2	48.0	47.3	47.1		HodgeRank	46.4	249	738	1429	2191
Extension	TwoStageHeap	4e-3	4e-3	4e-3	4e-3	4e-3	Extension	TwoStageHeap	4e-3	5e-3	8e-3	0.015	0.016
Rate + Compare	Combine	476	653	847	1056	1320	Rate + Compare	Combine	553	923	1256	2447	3862
	Hybrid	246	248	251	249	252		Hybrid	245	457	532	1043	1535
	HybridMPR	12.4	22.9	40.4	55.6	62.8		HybridMPR	15.6	42.3	60.6	112	143

(a) Efficiency

(b) Scalability

Figure 21: Efficiency and Scalability of Inference Methods (Seconds).

Top-k Algorithms		Selection Methods					
		Input	Output	EventTime	PeopleAge	PeopleNum	ImageClarity
Heuristic (Local)	BordaCount	Graph	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
	Copeland	Graph	Score	CrowdBT	CrowdGauss	CrowdGauss	Combine
	Local	Graph	Score	CrowdGauss	CrowdBT	Combine	CrowdGauss
	Indegree	Graph	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
	ELO	Comparison	Score	Combine	Combine	Combine	Combine
	BRE	Comparison	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
	URE	Comparison	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
	SSCO	Comparison	Score	SSCO	SSCO	SSCO	SSCO
	SSSE	Comparison	Score	SSSE	SSSE	SSSE	SSSE
Heuristic (Global)	Iterative	Graph	Subset	CrowdGauss	CrowdGauss	CrowdGauss	Combine
	PathRank	Graph	Subset	CrowdBT	CrowdBT	CrowdGauss	CrowdGauss
	AdaptiveReduce	Comparison	Subset	Combine	CrowdBT	Combine	CrowdBT
	MPageRank	Graph	Score	CrowdBT	CrowdGauss	CrowdBT	CrowdBT
	RankCentrality	Graph	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
	SSRW	Graph	Score	SSRW	SSRW	SSRW	SSRW
ML	TrueSkill	Comparison	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
	CrowdBT	Comparison	Score	CrowdBT	CrowdBT	CrowdBT	CrowdBT
	CrowdGauss	Comparison	Score	CrowdGauss	CrowdGauss	CrowdGauss	CrowdGauss
	HodgeRank	Comparison	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
Extension	TwoStageHeap	Comparison	Subset	TwoStageHeap	TwoStageHeap	TwoStageHeap	TwoStageHeap
Rate + Compare	Combine	Rate+Compare	Score	Combine	Combine	Combine	Combine
	Hybrid	Rate+Compare	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss
	HybridMPR	Rate+Graph	Score	CrowdBT	CrowdGauss	CrowdGauss	CrowdGauss

Figure 22: The Best Combination of Inference and Selection Algorithms on Each Dataset.

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.4	0.35	0.32	0.35	0.33
	Max	0.6	0.6	0.58	0.61	0.63
	Greedy	0.49	0.48	0.49	0.46	0.45
	Complete	0.53	0.51	0.52	0.53	0.53
	SSCO	0.72	0.82	0.85	0.87	0.9
	SSSE	0.7	0.82	0.82	0.87	0.9
	SSRW	0.68	0.81	0.87	0.9	0.94
ML	CrowdGauss	0.7	0.8	0.84	0.88	0.89
	CrowdBT	0.78	0.86	0.88	0.88	0.92
	Combine	0.7	0.78	0.88	0.9	0.9
Random	Random	0.71	0.8	0.83	0.87	0.9

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.69	0.65	0.63	0.62	0.62
	Max	0.73	0.72	0.72	0.71	0.7
	Greedy	0.69	0.68	0.68	0.68	0.68
	Complete	0.71	0.7	0.69	0.68	0.68
	SSCO	0.85	0.9	0.91	0.91	0.9
	SSSE	0.85	0.9	0.92	0.91	0.9
	SSRW	0.84	0.9	0.92	0.91	0.9
ML	CrowdGauss	0.9	0.9	0.9	0.9	0.9
	CrowdBT	0.88	0.9	0.9	0.9	0.91
	Combine	0.88	0.9	0.9	0.9	0.9
Random	Random	0.83	0.9	0.91	0.91	0.9

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.58	0.54	0.53	0.53	0.53
	Max	0.61	0.61	0.61	0.61	0.61
	Greedy	0.6	0.59	0.58	0.58	0.58
	Complete	0.6	0.6	0.6	0.58	0.59
	SSCO	0.68	0.71	0.71	0.71	0.7
	SSSE	0.68	0.7	0.71	0.71	0.7
	SSRW	0.67	0.7	0.7	0.7	0.7
ML	CrowdGauss	0.7	0.7	0.7	0.7	0.7
	CrowdBT	0.71	0.7	0.7	0.7	0.71
	Combine	0.69	0.7	0.7	0.7	0.7
Random	Random	0.68	0.7	0.71	0.71	0.7

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.57	0.55	0.58	0.56	0.56
	Max	0.7	0.68	0.67	0.6	0.6
	Greedy	0.65	0.64	0.63	0.59	0.57
	Complete	0.66	0.67	0.66	0.65	0.65
	SSCO	0.86	0.93	0.95	0.98	1.0
	SSSE	0.85	0.94	0.96	0.98	1.0
	SSRW	0.77	0.83	0.9	0.96	1.0
ML	CrowdGauss	1.0	1.0	1.0	1.0	0.99
	CrowdBT	0.91	1.0	1.0	0.99	0.99
	Combine	0.86	1.0	1.0	1.0	1.0
Random	Random	0.84	0.93	0.97	0.98	1.0

Figure 23: Recall on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:CrowdBT).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.05	0.02	0.07	0.05	0.07
	Max	0.16	0.13	0.16	0.2	0.18
	Greedy	0.09	0.11	0.08	0.04	0.1
	Complete	0.11	0.15	0.12	0.1	0.1
	SSCO	0.24	0.36	0.43	0.49	0.47
	SSSE	0.24	0.37	0.42	0.48	0.47
	SSRW	0.24	0.36	0.49	0.51	0.48
ML	CrowdGauss	0.26	0.36	0.45	0.53	0.51
	CrowdBT	0.28	0.39	0.49	0.51	0.51
	Combine	0.24	0.36	0.49	0.51	0.5
Random	Random	0.24	0.35	0.43	0.49	0.47

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.22	0.18	0.19	0.2	0.2
	Max	0.24	0.25	0.21	0.25	0.25
	Greedy	0.22	0.22	0.25	0.21	0.23
	Complete	0.23	0.25	0.21	0.22	0.22
	SSCO	0.41	0.52	0.54	0.55	0.54
	SSSE	0.41	0.53	0.55	0.55	0.54
	SSRW	0.4	0.51	0.54	0.51	0.53
ML	CrowdGauss	0.45	0.53	0.55	0.51	0.52
	CrowdBT	0.49	0.53	0.53	0.53	0.54
	Combine	0.48	0.53	0.53	0.53	0.53
Random	Random	0.39	0.5	0.55	0.55	0.54

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.15	0.15	0.11	0.13	0.12
	Max	0.18	0.18	0.17	0.19	0.17
	Greedy	0.16	0.17	0.16	0.17	0.16
	Complete	0.15	0.17	0.18	0.15	0.15
	SSCO	0.26	0.3	0.3	0.3	0.31
	SSSE	0.26	0.3	0.29	0.3	0.31
	SSRW	0.23	0.27	0.29	0.29	0.29
ML	CrowdGauss	0.26	0.27	0.31	0.29	0.29
	CrowdBT	0.29	0.3	0.3	0.3	0.3
	Combine	0.26	0.3	0.3	0.3	0.3
Random	Random	0.24	0.28	0.29	0.3	0.31

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.16	0.14	0.15	0.14	0.15
	Max	0.22	0.21	0.23	0.23	0.24
	Greedy	0.21	0.18	0.19	0.23	0.2
	Complete	0.18	0.21	0.22	0.19	0.23
	SSCO	0.43	0.59	0.67	0.74	0.78
	SSSE	0.42	0.6	0.69	0.75	0.78
	SSRW	0.34	0.53	0.69	0.67	0.8
ML	CrowdGauss	0.78	0.76	0.78	0.8	0.78
	CrowdBT	0.58	0.76	0.76	0.76	0.75
	Combine	0.49	0.79	0.79	0.79	0.79
Random	Random	0.41	0.6	0.69	0.74	0.78

Figure 24: ACC^k on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:CrowdBT).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.48	0.48	0.48	0.48	0.48
	Max	0.5	0.49	0.49	0.49	0.48
	Greedy	0.45	0.46	0.45	0.44	0.45
	Complete	0.51	0.52	0.5	0.49	0.5
	SSCO	0.59	0.69	0.74	0.8	0.83
	SSSE	0.6	0.69	0.74	0.8	0.83
	SSRW	0.56	0.66	0.73	0.8	0.83
ML	CrowdGauss	0.6	0.8	0.8	0.81	0.83
	CrowdBT	0.56	0.71	0.81	0.82	0.84
	Combine	0.6	0.79	0.82	0.82	0.83
Random	Random	0.58	0.68	0.74	0.8	0.82

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.74	0.74	0.75	0.74	0.75
	Max	0.69	0.69	0.69	0.69	0.69
	Greedy	0.68	0.68	0.69	0.69	0.68
	Complete	0.7	0.69	0.7	0.7	0.7
	SSCO	0.82	0.89	0.9	0.94	1.0
	SSSE	0.82	0.89	0.91	0.94	1.0
	SSRW	0.8	0.87	0.9	0.92	1.0
ML	CrowdGauss	0.9	0.95	1.0	1.0	0.99
	CrowdBT	0.77	0.82	0.95	0.99	1.0
	Combine	0.85	0.93	0.98	1.0	1.0
Random	Random	0.81	0.87	0.91	0.94	1.0

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.62	0.63	0.63	0.63	0.63
	Max	0.63	0.62	0.63	0.62	0.62
	Greedy	0.6	0.6	0.59	0.6	0.6
	Complete	0.6	0.6	0.61	0.61	0.61
	SSCO	0.68	0.7	0.7	0.7	0.7
	SSSE	0.68	0.7	0.7	0.7	0.7
	SSRW	0.67	0.69	0.7	0.7	0.7
ML	CrowdGauss	0.7	0.7	0.7	0.7	0.7
	CrowdBT	0.66	0.68	0.7	0.7	0.7
	Combine	0.7	0.7	0.7	0.7	0.7
Random	Random	0.68	0.7	0.7	0.7	0.7

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.72	0.72	0.72	0.71	0.72
	Max	0.71	0.72	0.73	0.72	0.71
	Greedy	0.66	0.69	0.68	0.67	0.71
	Complete	0.7	0.69	0.69	0.7	0.7
	SSCO	0.84	0.9	0.92	0.96	1.0
	SSSE	0.83	0.9	0.93	0.96	1.0
	SSRW	0.79	0.89	0.92	0.96	1.0
ML	CrowdGauss	0.9	0.9	1.0	1.0	1.0
	CrowdBT	0.78	1.0	1.0	1.0	1.0
	Combine	0.84	1.0	1.0	1.0	1.0
Random	Random	0.81	0.88	0.93	0.96	1.0

Figure 25: Recall on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:Iterative).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.09	0.08	0.09	0.09	0.09
	Max	0.1	0.09	0.1	0.09	0.09
	Greedy	0.08	0.09	0.08	0.08	0.08
	Complete	0.1	0.11	0.09	0.1	0.1
	SSCO	0.15	0.2	0.22	0.26	0.3
	SSSE	0.15	0.2	0.22	0.25	0.3
	SSRW	0.12	0.15	0.2	0.23	0.3
ML	CrowdGauss	0.15	0.27	0.29	0.31	0.3
	CrowdBT	0.16	0.24	0.29	0.28	0.31
	Combine	0.14	0.25	0.28	0.27	0.27
Random	Random	0.13	0.18	0.22	0.26	0.29

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.22	0.22	0.23	0.22	0.22
	Max	0.19	0.19	0.19	0.19	0.19
	Greedy	0.19	0.19	0.19	0.19	0.19
	Complete	0.19	0.19	0.19	0.2	0.2
	SSCO	0.28	0.33	0.33	0.36	0.41
	SSSE	0.28	0.32	0.33	0.36	0.41
	SSRW	0.26	0.32	0.34	0.38	0.42
ML	CrowdGauss	0.29	0.36	0.39	0.4	0.41
	CrowdBT	0.24	0.27	0.37	0.4	0.41
	Combine	0.29	0.35	0.39	0.41	0.41
Random	Random	0.26	0.31	0.34	0.36	0.41

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.15	0.15	0.16	0.15	0.15
	Max	0.15	0.15	0.16	0.15	0.15
	Greedy	0.14	0.14	0.14	0.14	0.14
	Complete	0.14	0.14	0.15	0.15	0.15
	SSCO	0.18	0.19	0.19	0.19	0.19
	SSSE	0.18	0.19	0.19	0.19	0.19
	SSRW	0.18	0.19	0.19	0.19	0.2
ML	CrowdGauss	0.18	0.19	0.2	0.2	0.2
	CrowdBT	0.17	0.18	0.19	0.19	0.19
	Combine	0.19	0.19	0.19	0.19	0.19
Random	Random	0.18	0.19	0.19	0.19	0.19

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.2	0.22	0.21	0.2	0.2
	Max	0.2	0.21	0.2	0.21	0.21
	Greedy	0.18	0.2	0.19	0.19	0.2
	Complete	0.2	0.19	0.19	0.19	0.2
	SSCO	0.29	0.34	0.34	0.37	0.41
	SSSE	0.29	0.34	0.35	0.38	0.41
	SSRW	0.25	0.33	0.35	0.38	0.42
ML	CrowdGauss	0.31	0.37	0.38	0.4	0.42
	CrowdBT	0.23	0.42	0.43	0.4	0.4
	Combine	0.27	0.41	0.41	0.41	0.42
Random	Random	0.27	0.31	0.35	0.38	0.41

Figure 26: ACC^k on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:Iterative).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.29	0.26	0.27	0.23	0.23
	Max	0.3	0.27	0.27	0.24	0.29
	Greedy	0.31	0.35	0.36	0.37	0.32
	Complete	0.27	0.34	0.35	0.39	0.33
	SSCO	0.53	0.61	0.64	0.68	0.69
	SSSE	0.53	0.61	0.65	0.68	0.7
	SSRW	0.53	0.6	0.69	0.68	0.71
ML	CrowdGauss	0.5	0.65	0.67	0.69	0.69
	CrowdBT	0.6	0.64	0.68	0.69	0.71
	Combine	0.49	0.61	0.7	0.72	0.72
Random	Random	0.52	0.57	0.65	0.67	0.7

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.56	0.48	0.44	0.47	0.48
	Max	0.61	0.48	0.43	0.41	0.43
	Greedy	0.62	0.58	0.57	0.54	0.57
	Complete	0.59	0.57	0.58	0.57	0.57
	SSCO	0.79	0.87	0.89	0.91	0.9
	SSSE	0.79	0.87	0.89	0.91	0.9
	SSRW	0.76	0.84	0.89	0.9	0.91
ML	CrowdGauss	0.8	0.9	0.9	0.9	0.91
	CrowdBT	0.77	0.85	0.91	0.91	0.91
	Combine	0.84	0.9	0.91	0.91	0.9
Random	Random	0.77	0.86	0.89	0.91	0.9

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.45	0.39	0.41	0.42	0.41
	Max	0.48	0.36	0.36	0.36	0.38
	Greedy	0.54	0.51	0.47	0.47	0.49
	Complete	0.51	0.51	0.51	0.46	0.48
	SSCO	0.66	0.7	0.7	0.7	0.7
	SSSE	0.66	0.69	0.7	0.71	0.7
	SSRW	0.65	0.68	0.7	0.71	0.7
ML	CrowdGauss	0.7	0.7	0.7	0.7	0.7
	CrowdBT	0.65	0.7	0.7	0.7	0.7
	Combine	0.68	0.7	0.7	0.7	0.7
Random	Random	0.64	0.69	0.7	0.7	0.7

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.49	0.43	0.41	0.45	0.44
	Max	0.54	0.43	0.39	0.39	0.38
	Greedy	0.57	0.55	0.54	0.54	0.53
	Complete	0.54	0.52	0.54	0.53	0.52
	SSCO	0.74	0.88	0.92	0.95	0.98
	SSSE	0.74	0.87	0.92	0.95	0.98
	SSRW	0.67	0.82	0.85	0.9	0.94
ML	CrowdGauss	0.9	0.9	1.0	1.0	0.99
	CrowdBT	0.73	0.94	0.96	0.96	0.97
	Combine	0.74	0.97	1.0	0.98	0.98
Random	Random	0.72	0.85	0.93	0.95	0.98

Figure 27: Recall on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:Hybrid).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.03	0.03	0.03	0.02	0.02
	Max	0.04	0.03	0.03	0.02	0.02
	Greedy	0.04	0.05	0.05	0.07	0.04
	Complete	0.02	0.04	0.06	0.06	0.03
	SSCO	0.13	0.16	0.18	0.21	0.22
	SSSE	0.13	0.16	0.18	0.21	0.23
	SSRW	0.11	0.16	0.19	0.22	0.23
ML	CrowdGauss	0.11	0.17	0.19	0.2	0.22
	CrowdBT	0.16	0.18	0.2	0.21	0.23
	Combine	0.1	0.15	0.24	0.24	0.22
Random	Random	0.11	0.13	0.18	0.2	0.22

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.13	0.09	0.1	0.09	0.09
	Max	0.18	0.11	0.04	0.06	0.08
	Greedy	0.19	0.14	0.15	0.13	0.16
	Complete	0.17	0.13	0.13	0.14	0.14
	SSCO	0.31	0.43	0.5	0.53	0.53
	SSSE	0.31	0.43	0.5	0.54	0.53
	SSRW	0.3	0.4	0.47	0.49	0.54
ML	CrowdGauss	0.38	0.53	0.52	0.52	0.53
	CrowdBT	0.33	0.46	0.54	0.54	0.54
	Combine	0.39	0.53	0.54	0.54	0.53
Random	Random	0.3	0.42	0.5	0.54	0.53

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.11	0.08	0.09	0.06	0.07
	Max	0.09	0.04	0.05	0.04	0.04
	Greedy	0.13	0.13	0.12	0.1	0.1
	Complete	0.14	0.1	0.1	0.09	0.13
	SSCO	0.22	0.26	0.28	0.29	0.3
	SSSE	0.22	0.27	0.28	0.29	0.3
	SSRW	0.21	0.25	0.3	0.29	0.3
ML	CrowdGauss	0.22	0.29	0.29	0.29	0.3
	CrowdBT	0.21	0.28	0.29	0.3	0.3
	Combine	0.23	0.29	0.3	0.3	0.29
Random	Random	0.2	0.25	0.28	0.29	0.3

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.05	0.06	0.02	0.05	0.05
	Max	0.03	0.02	0.04	0.04	0.05
	Greedy	0.1	0.11	0.1	0.06	0.11
	Complete	0.09	0.08	0.09	0.06	0.1
	SSCO	0.27	0.46	0.58	0.66	0.74
	SSSE	0.27	0.45	0.58	0.66	0.73
	SSRW	0.18	0.4	0.48	0.5	0.7
ML	CrowdGauss	0.56	0.6	0.75	0.74	0.74
	CrowdBT	0.31	0.64	0.7	0.7	0.72
	Combine	0.32	0.7	0.74	0.75	0.74
Random	Random	0.27	0.44	0.59	0.66	0.73

Figure 28: ACC^k on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:Hybrid).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.4	0.41	0.42	0.4	0.41
	Max	0.46	0.46	0.45	0.46	0.47
	Greedy	0.39	0.4	0.42	0.4	0.43
	Complete	0.41	0.42	0.42	0.46	0.48
	SSCO	0.55	0.68	0.72	0.76	0.8
	SSSE	0.56	0.67	0.72	0.76	0.8
	SSRW	0.58	0.67	0.76	0.79	0.83
ML	CrowdGauss	0.6	0.8	0.8	0.8	0.8
	CrowdBT	0.65	0.75	0.78	0.8	0.8
	Combine	0.45	0.64	0.81	0.8	0.83
Random	Random	0.53	0.66	0.72	0.77	0.8

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.73	0.75	0.74	0.75	0.75
	Max	0.69	0.68	0.69	0.68	0.69
	Greedy	0.69	0.7	0.7	0.69	0.7
	Complete	0.68	0.68	0.68	0.67	0.68
	SSCO	0.74	0.85	0.88	0.91	0.9
	SSSE	0.74	0.84	0.88	0.9	0.9
	SSRW	0.75	0.83	0.88	0.91	0.9
ML	CrowdGauss	0.7	0.9	0.9	0.9	0.9
	CrowdBT	0.8	0.88	0.91	0.91	0.91
	Combine	0.82	0.91	0.91	0.9	0.9
Random	Random	0.72	0.83	0.88	0.91	0.9

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.63	0.62	0.63	0.63	0.63
	Max	0.57	0.57	0.57	0.58	0.57
	Greedy	0.59	0.59	0.59	0.59	0.59
	Complete	0.58	0.57	0.58	0.58	0.58
	SSCO	0.64	0.7	0.7	0.72	0.7
	SSSE	0.64	0.69	0.7	0.72	0.7
	SSRW	0.64	0.69	0.7	0.7	0.7
ML	CrowdGauss	0.6	0.7	0.7	0.7	0.7
	CrowdBT	0.68	0.71	0.7	0.7	0.7
	Combine	0.65	0.7	0.7	0.7	0.7
Random	Random	0.62	0.67	0.7	0.72	0.7

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.69	0.68	0.65	0.65	0.7
	Max	0.54	0.53	0.53	0.52	0.55
	Greedy	0.58	0.53	0.55	0.6	0.58
	Complete	0.53	0.55	0.5	0.54	0.54
	SSCO	0.67	0.77	0.82	0.88	1.0
	SSSE	0.66	0.76	0.82	0.88	1.0
	SSRW	0.6	0.77	0.85	0.85	0.98
ML	CrowdGauss	0.9	0.9	1.0	1.0	1.0
	CrowdBT	0.85	0.86	0.98	0.98	0.98
	Combine	0.52	0.72	0.98	1.0	1.0
Random	Random	0.64	0.75	0.82	0.87	1.0

Figure 29: Recall on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:Local).

(a) EventTime (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.08	0.06	0.08	0.03	0.04
	Max	0.1	0.06	0.05	0.13	0.12
	Greedy	0.03	0.07	0.06	0.2	0.04
	Complete	0.09	0.06	0.06	0.07	0.02
	SSCO	0.14	0.22	0.27	0.32	0.4
	SSSE	0.14	0.22	0.27	0.33	0.4
	SSRW	0.14	0.23	0.33	0.36	0.4
ML	CrowdGauss	0.22	0.35	0.35	0.35	0.38
	CrowdBT	0.2	0.26	0.3	0.36	0.4
	Combine	0.09	0.17	0.4	0.37	0.41
Random	Random	0.12	0.2	0.26	0.34	0.4

(c) PeopleNum (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.24	0.26	0.3	0.29	0.28
	Max	0.22	0.22	0.17	0.23	0.21
	Greedy	0.23	0.22	0.25	0.22	0.24
	Complete	0.25	0.23	0.2	0.23	0.23
	SSCO	0.25	0.37	0.44	0.5	0.51
	SSSE	0.25	0.36	0.43	0.5	0.51
	SSRW	0.25	0.36	0.41	0.44	0.51
ML	CrowdGauss	0.2	0.56	0.56	0.53	0.52
	CrowdBT	0.34	0.49	0.54	0.54	0.54
	Combine	0.35	0.49	0.54	0.52	0.51
Random	Random	0.23	0.35	0.43	0.51	0.51

(b) PeopleAge (c-Accuracy = 85.2% r-Accuracy = 100%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.16	0.18	0.19	0.19	0.19
	Max	0.16	0.16	0.14	0.15	0.14
	Greedy	0.15	0.16	0.15	0.15	0.18
	Complete	0.15	0.16	0.15	0.14	0.16
	SSCO	0.19	0.25	0.26	0.29	0.27
	SSSE	0.19	0.25	0.26	0.29	0.27
	SSRW	0.18	0.26	0.27	0.3	0.27
ML	CrowdGauss	0.16	0.25	0.31	0.29	0.29
	CrowdBT	0.22	0.29	0.29	0.29	0.29
	Combine	0.19	0.25	0.29	0.3	0.29
Random	Random	0.18	0.22	0.26	0.29	0.27

(d) ImageClarity (c-Accuracy = 98.6% r-Accuracy = 97%).

Top-k Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic	Group	0.21	0.24	0.24	0.24	0.26
	Max	0.12	0.12	0.14	0.11	0.14
	Greedy	0.18	0.13	0.14	0.15	0.15
	Complete	0.22	0.11	0.14	0.15	0.15
	SSCO	0.21	0.3	0.38	0.48	0.73
	SSSE	0.2	0.3	0.37	0.49	0.73
	SSRW	0.15	0.33	0.44	0.56	0.73
ML	CrowdGauss	0.55	0.51	0.75	0.73	0.72
	CrowdBT	0.43	0.44	0.71	0.72	0.71
	Combine	0.13	0.37	0.61	0.72	0.72
Random	Random	0.19	0.29	0.38	0.47	0.73

Figure 30: ACC^k on Real Datasets for All Selection Methods: Varying Selection Rate(Inference Method:Local).

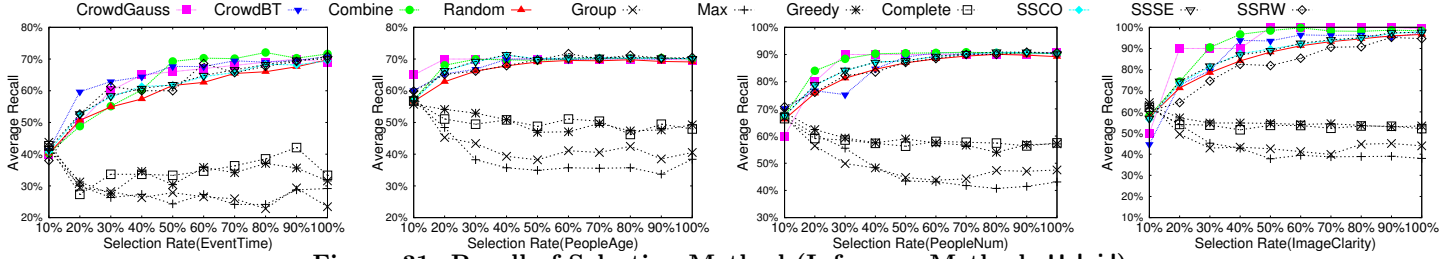


Figure 31: Recall of Selection Method (Inference Method: Hybrid).

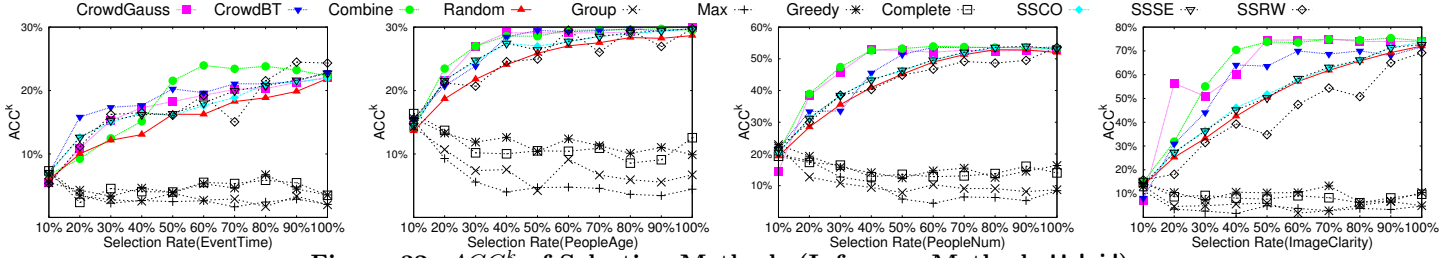


Figure 32: ACC^k of Selection Methods (Inference Method: Hybrid).




Figure 33: Recall of Selection Methods (Inference Method: Local).

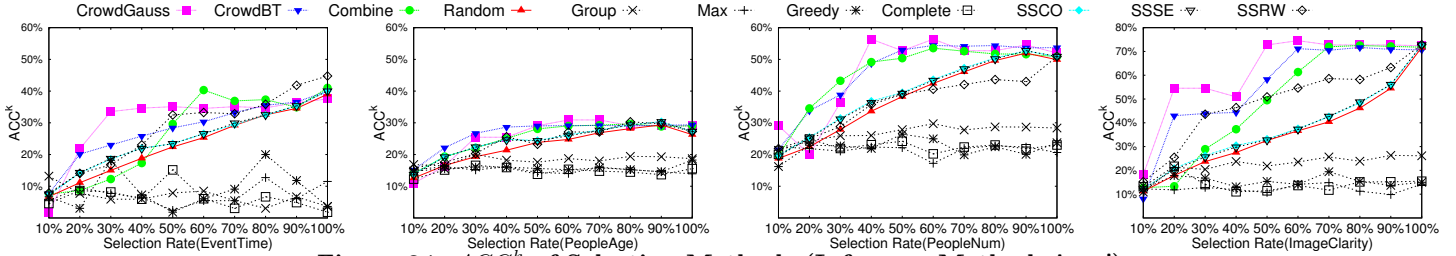


Figure 34: ACC^k of Selection Methods (Inference Method: Local).

B. EXPERIMENTAL RESULTS FOR SORTING

The top- k algorithms that estimated a score for each object can be used to sort the objects based on the scores while the iterative reduce based methods cannot. Table 2 showed the algorithms that can be used to the sorting problem. As the inference time and selection time of these methods on the sorting problem were similar to those on the top- k problem, we did not show the efficiency of the methods and we only compared the quality.

B.1 Experimental Setting

The ACC metric could be utilized to evaluate the sorting quality, which was defined as below,

$$ACC = \frac{\sum_{i,j} \mathbb{I}(a_i < a_j \wedge t_i < t_j)}{\sum_{i,j} \mathbb{I}(t_i < t_j)} \quad (43)$$

where t_i was the sorting position for object o_i and a_i was the sorting position by an algorithm. \mathbb{I} was a function, which returned 1 if the condition was true; 0 otherwise. However, we noticed that if we utilized ACC as the evaluation metric, all the algorithms nearly had the same ACC . For example, in Figure 35(c), if the selection rate was 40%, then the average deviation ($|t_i - a_i|$) of each object for CrowdBT was 6.7, and the deviation was 27.4 for BordaCount. This implied that even BordaCount was utilized for sorting, the average deviation for each object was not large, and a highly ranked object would not be arranged at the lowly ranked position. Thus the number of inverse pairs among different algorithms would not have too much difference. Moreover, the denominator of Equation 43 was $\binom{1000}{2}$ on the simulated dataset. Note that the average number of the inverse pairs for CrowdBT was twenty thousand less than BordaCount. Thus if this difference was divided by $\binom{1000}{2}$, then the difference of ACC value between CrowdBT and BordaCount was trivial.

To address this problem, we used another widely-used metric, Spearman’s footrule distance, which was defined as below

$$dis(\pi) = \sum_{i=1}^n |\pi(i) - \pi^*(i)| \quad (44)$$

where π was the estimated permutation in a sorting method, and π^* was the true permutation. For example, consider six objects and the true permutation was $\langle 1, 2, 3, 4, 5, 6 \rangle$. An algorithm reported a permutation $\langle 2, 1, 4, 3, 5, 6 \rangle$. Then the Spearman’s footrule distance was 4.

B.2 Evaluation of Inference Methods

B.2.1 Simulation – Varying Selection Rate

We compared the inference methods by randomly selecting crowdsourced pairs with different selection rates, and we evaluated different selection strategies in Section B.3. Figure 35 showed the results of all methods and Figure 37 showed the results of some selected representative methods from each category. For the heuristic-based algorithms, we selected Local and SSRW as they achieved higher performance. For the machine-learning methods, we selected TrueSkill, CrowdBT, CrowdGauss and HodgeRank. For hybrid methods, we selected Combine. Thus we showed the results of these seven algorithms in Figure 37.

We had the following observations. First, for low worker quality (e.g., the worker accuracy lower than 70%), all the methods had a bad performance. This was because the collected results had too much noise, and neither the heuristics nor the machine learning based methods could recover the original order. And we had a similar observation in the top- k experiments on the simulated datasets.

Second, for high worker quality (e.g., the worker quality larger than 80%), all the methods achieved a better performance. CrowdBT had a tremendous advantage over the other methods. And this was also consistent with the observations when we utilized CrowdBT to find the top- k objects, because CrowdBT focused on ranking the objects and could assign a reasonable score for each object.

Third, among the heuristic-based methods, for the local comparison heuristics, BordaCount, ELO and URE got a high spearman’s footrule distance compared with other methods, because BordaCount only utilized the out-degree of the pairwise graph and could not get a reasonable result. For the same reason, the recall of BordaCount in the top- k experiments was also low.

For ELO, the objects had different weights, and updating the scores depended on the weights. In addition, ELO involved multiple iterations to finally convergence. So we found that the spearman’s footrule distance for ELO was high. Besides, the performance of ELO in the top- k experiments was also bad due to the same reason.

For URE, this method was mainly designed for obtaining the highest ranked objects. When the accuracy was high, URE could obtain a reasonable top- k result. But URE could not perform well for sorting, since URE only utilized the incoming edges for each object. And we observed that it had a similar performance with BordaCount.

For the global comparison heuristics, SSRW and Rank-Centrality had a similar performance and they were both superior to MPageRank. In MPageRank, based on the update equation 7, if the number of iterative time was large, some objects had zero PageRank value. And these objects could not be sorted. However, in MPageRank, the highest ranked objects could accumulate a relative high score based on Equation 7 even with a random sample of the comparisons. And we found that MPageRank could still be used to obtain a reasonable top- k result.

Fourth, for the machine-learning methods, CrowdBT outperformed other methods. CrowdBT outperformed CrowdGauss as CrowdBT considered the worker quality. HodgeRank had a worse performance, because the matrix decomposition had deviations and led to errors to compute latent scores. And this was also consistent with the observation in the top- k experiments.

Fifth, for the hybrid methods, the performance for Hybrid and HybridMPR was poor when the selection rate was low, due to their ineffective pruning strategy. Hybrid and HybridMPR both had two stages. In the first stage, many objects would be filtered and they could not participate in the second comparison stage. For these filtered objects, they would be randomly added to the final ordered result and could not be correctly sorted. However, both Hybrid and HybridMPR achieved a high recall in the top- k experiments when the selection rate was low, because both of the methods filtered objects that were not likely in the top- k results in the first stage.

B.2.2 Real Experiments – Varying Selection Rate

We compared different algorithms on real datasets by varying selection rates. Figures 38 illustrated the results for the selected algorithms and the detailed results of all the methods could be found in Figure 36.

We had the following observations. First, the observations on the simulation experiments were still true on the real datasets. On the real datasets, **CrowdBT** and **CrowdGauss** had similar performance, which outperformed other methods. As the size of the real datasets was small, even with a limited selection rate, we could get a strongly-connected graph, and the improvement of **CrowdBT** over **CrowdGauss** was not obvious. This was also consistent with the observation in the top- k experiments conducted on the real datasets.

Second, in the **PeopleAge** and **PeopleNum** datasets, most methods had a similar performance, and the advantages of **CrowdBT** or **CrowdGauss** were not obvious. For the **PeopleAge** dataset, the noise mainly came from the dataset itself, so even with a high selection rate, the spearman’s footrule distance was still high and the recall could not be further improved. For the **PeopleNum** dataset, the accuracy was higher than in the **PeopleAge** dataset, so the spearman’s footrule distance on **PeopleNum** was smaller and the average recall was higher. In the **PeopleAge** and **PeopleNum** datasets, all the methods converged very quickly even with a small selection rate.

B.2.3 Takeaways: Inference Algorithm Suggestions

We had the following suggestions that can guide practitioners to select appropriate inference algorithms for the sorting problem.

- (1) If the user preferred quality, we recommended **CrowdBT**.
- (2) If the user preferred efficiency, although **Iterative** was the best choice for the top- k problem, it could not be used to support the sorting problem, as it used an adaptive reduce based method, and we recommended **Local**.
- (3) We recommended 30% selection rate.

B.3 Evaluation on Selection Methods

B.3.1 Real Datasets – Varying Selection Rate

We compared the performance for the selection methods on the real datasets. For the inference methods, we selected **CrowdGauss**, **CrowdBT**, **HodgeRank** and **TrueSkill**. The other three selection methods **SSCO**, **SSSE** and **SSRW** were not utilized in this section since they were mainly designed for resolving top- k issues. For all the selection methods used for sorting, Figures 39- 42 showed the results of all methods.

We had the following summaries. First, all the four heuristics had a very bad performance as in the top- k experiments, because they focused on finding the maximal object. Second, all the three active-learning methods were superior to **Random** selection method which was consistent in top- k experiments, as they would always select pairs with high information gain. Third, all the three active learning methods had a similar performance on the **PeopleAge**, **PeopleNum** and **ImageClarity** datasets. But in the **EventTime** dataset, **CrowdBT** had a advantage compared to **CrowdGauss** and **Combine**, as **CrowdBT** not only selected pairs but also selected high quality workers. So the average accuracy of the selected pairs by **CrowdBT** was higher than **CrowdGauss** and **Combine**. The above observation was also consistent with

the results in the top- k experiments. Fourth, with the increase of the selection rates, the spearman’s footrule distances of **CrowdBT**, **CrowdGauss**, **Combine**, and **Random** were decreased (i.e., the quality was improved, as the smaller distance the better), because they could utilize more crowd comparisons to improve the quality. However the spearman’s footrule distances of the heuristics, e.g., **Max**, **Group**, **Greedy**, and **Complete**, kept stable and even increased, because the heuristics selected many duplicated pairs and no more information was obtained with the increase of the selection rates. Especially if **TrueSkill** was used as the inference method, the spearman’s footrule distances largely grew with the increase of the selection rates, because if many duplicated pairs were selected, the scores of the lose objects would significantly decrease (even if they are not lowly ranked objects) compared with the objects that were not selected.

B.3.2 Takeaways: Selection Algorithm Suggestions

We had the following suggestions that can guide practitioners to select appropriate selection algorithms for the sorting problem.

- (1) If the user preferred high quality, **CrowdBT** was recommended.
- (2) If the user preferred high efficiency, although **SSCO** and **SSSE** were the best choice for the top- k problem, they could not be used to support the sorting problem, as they used adaptive reduce methods. We recommended **Combine** and **Random**. For low selection rate, **Combine** was recommended; while for high selection rate, **Random** was acceptable.

(a) c-Accuracy=63.85% r-Accuracy=66.9%.

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	108.7	77.1	60.8	50.1	42.0
	Copeland	85.2	63.0	52.5	46.2	42.0
	Local	82.5	59.8	49.4	42.8	38.3
	Indegree	78.8	56.7	46.4	40.5	36.4
	ELO	101	94.9	94.4	94.5	94.3
	BRE	84.7	61.4	50.6	43.9	39.2
	URE	108	75.9	59.2	47.7	38.3
Heuristic (Global)	MPageRank	108.7	77.4	61.5	51.4	43.5
	RankCentrality	147	87.5	65.3	54.2	44.0
	SSRW	105	74.4	58.4	47.3	38.7
ML	TrueSkill	84.8	60.2	50.9	45.4	43
	CrowdBT	81.5	58.3	48.3	41.7	37.7
	CrowdGauss	78.2	56.6	46.0	40.1	35.7
	HodgeRank	85.0	59.9	49.9	43.8	39.0
Rate + Compare	Combine	88.8	64.2	54.9	49.4	47.0
	Hybrid	189	134	85.0	45.2	28.5
	HybridMPR	188	126	86.9	50.3	35.0

(c) c-Accuracy=88.70% r-Accuracy=83.5%.

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	42.2	27.4	19.7	14.3	9.5
	Copeland	31.9	21.0	15.8	12.2	9.5
	Local	28.3	18.9	14.6	11.7	9.4
	Indegree	30.9	20.6	15.3	11.7	8.9
	ELO	30.0	24.5	22.8	22.2	21.7
	BRE	32.2	21.5	16.3	12.8	10.2
	URE	42.2	27.5	19.8	14.4	9.3
Heuristic (Global)	MPageRank	75.9	53.2	41.6	33.9	28.3
	RankCentrality	42.5	29.3	21.5	15.2	9.5
	SSRW	41.5	27.2	19.7	14.5	9.8
ML	TrueSkill	27.4	19.3	14.7	12.6	11.7
	CrowdBT	10.5	6.7	5.5	5.1	4.0
	CrowdGauss	31.0	20.5	15.3	11.7	9.0
	HodgeRank	31.3	21.3	15.8	12.7	9.2
Rate + Compare	Combine	28.9	20.8	16.2	14.1	11.2
	Hybrid	152	92.3	40.4	25.3	6.87
	HybridMPR	151	85.8	55.0	27.2	7.56

(b) c-Accuracy=77.45% r-Accuracy=73.9%.

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	60.0	40.2	30.4	23.6	18.4
	Copeland	45.6	31.8	25.3	21.2	18.4
	Local	42.9	30.1	24.2	20.5	17.9
	Indegree	42.6	29.5	23.4	19.4	16.5
	ELO	50.2	45.1	44.0	43.7	43.6
	BRE	45.1	31.3	24.8	20.7	17.8
	URE	60.0	40.4	30.6	24.0	18.7
Heuristic (Global)	MPageRank	79.4	57.3	46.3	39.2	34.0
	RankCentrality	62.3	42.5	33.4	25.2	16.9
	SSRW	57.5	38.9	29.4	22.8	17.6
ML	TrueSkill	43.4	30.6	24.7	21.6	20
	CrowdBT	27.6	23.7	10.0	8.1	9.5
	CrowdGauss	42.8	29.6	23.4	19.4	16.5
	HodgeRank	45.3	31.5	24.8	20.6	18.3
Rate + Compare	Combine	45.4	32.6	26.7	23.6	22.5
	Hybrid	166	113	54.0	23.5	10.2
	HybridMPR	165	106	73.1	32.5	10.3

(d) c-Accuracy=96.43% r-Accuracy=91.9%.

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	34.6	21.9	15.1	9.9	4.3
	Copeland	26.2	16.5	11.5	7.8	4.3
	Local	21.7	13.9	9.9	7.0	4.5
	Indegree	26.1	16.6	11.5	7.9	4.3
	ELO	20.1	14.3	12.2	11.2	10.5
	BRE	26.2	16.6	11.6	7.9	4.5
	URE	34.6	21.9	15.1	10.0	4.3
Heuristic (Global)	MPageRank	40.3	42.0	37.0	31.3	25.8
	RankCentrality	30.5	17.6	12.7	7.2	4.3
	SSRW	33.6	21.4	14.8	9.8	4.3
ML	TrueSkill	17.9	11.6	9.2	7.4	6.7
	CrowdBT	7.5	4.6	3.9	2.9	2.1
	CrowdGauss	26.0	16.5	11.5	7.8	4.3
	HodgeRank	26.0	16.7	11.7	8.0	4.2
Rate + Compare	Combine	18.9	12.6	10.2	8.4	6.9
	Hybrid	145	62.7	29.8	15.6	4.85
	HybridMPR	137	68.9	38.0	17.6	4.95

Figure 35: Spearman's Footrule Distance ($\times 1000$) on Simulation for All Inference Methods: Varying Selection Rate (1000 Objects).

(a) **EventTime** (c-Accuracy = 76.6% r-Accuracy = 89.4%).

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	1518	1182	994	855	740
	Copeland	1195	919	792	716	658
	Local	1065	809	697	630	584
	Indegree	1159	923	818	751	698
	ELO	1186	878	742	666	616
	BRE	1199	923	793	715	658
	URE	1408	1052	872	758	678
Heuristic (Global)	MPageRank	1247	876	703	614	576
	RankCentrality	1218	930	809	728	700
	SSRW	1535	1000	818	712	645
ML	TrueSkill	1038	744	655	582	551
	CrowdBT	986	766	674	617	582
	CrowdGauss	843	648	569	525	488
	HodgeRank	948	721	625	573	544
	Combine	2687	1412	906	644	568
Rate + Compare	Hybrid	1381	1204	1004	974	888
	HybridMPR	1361	1278	1201	1133	1048

(c) **PeopleNum** (c-Accuracy = 88.4% r-Accuracy = 71.8%).

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	244	190	161	139	119
	Copeland	204	163	142	128	119
	Local	211	163	138	125	118
	Indegree	203	164	144	131	124
	ELO	212	164	141	128	119
	BRE	204	163	142	128	119
	URE	233	179	152	133	119
Heuristic (Global)	MPageRank	247	175	155	145	144
	RankCentrality	207	165	144	129	123
	SSRW	433	218	155	133	120
ML	TrueSkill	188	147	132	123	119
	CrowdBT	165	134	124	120	119
	CrowdGauss	166	133	124	120	120
	HodgeRank	170	140	129	122	119
Rate + Compare	Combine	212	155	136	133	129
	Hybrid	269	225	180	145	122
	HybridMPR	271	235	193	170	136

(b) **PeopleAge** (c-Accuracy = 85.2% r-Accuracy = 100%).

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	399	333	300	282	264
	Copeland	353	305	285	274	264
	Local	346	298	278	269	264
	Indegree	345	303	283	273	264
	ELO	357	302	283	273	265
	BRE	352	304	286	274	264
	URE	394	326	296	279	264
Heuristic (Global)	MPageRank	392	310	288	279	270
	RankCentrality	355	306	286	275	270
	SSRW	705	370	294	276	264
ML	TrueSkill	327	288	275	270	264
	CrowdBT	304	278	270	264	258
	CrowdGauss	307	277	269	264	262
	HodgeRank	311	284	274	269	262
	Combine	362	298	279	267	256
Rate + Compare	Hybrid	426	385	331	266	266
	HybridMPR	424	390	328	269	265

(d) **ImageClarity** (c-Accuracy = 98.6% r-Accuracy = 97%).

Sorting Algorithms		Selection Rate				
		20%	40%	60%	80%	100%
Heuristic (Local)	BordaCount	966	638	456	314	157
	Copeland	757	494	352	244	134
	Local	639	397	278	191	114
	Indegree	742	484	342	236	124
	ELO	732	443	299	206	134
	BRE	757	492	352	244	138
	URE	959	630	449	307	158
Heuristic (Global)	MPageRank	729	413	289	210	136
	RankCentrality	770	512	367	263	212
	SSRW	2735	984	433	263	131
ML	TrueSkill	561	330	227	161	105
	CrowdBT	372	205	139	103	76
	CrowdGauss	325	171	116	90	72
	HodgeRank	482	304	214	151	92
Rate + Compare	Combine	458	250	141	91	62
	Hybrid	860	612	360	265	114
	HybridMPR	878	627	378	275	148

Figure 36: Spearman's Footrule Distance on Real Datasets for All Inference Methods: Varying Selection Rate.




Figure 37: Spearman's Footrule Distance on Simulated Datasets for Selected Inference Methods.




Figure 38: Spearman's Footrule Distance on Real Datasets for Selected Inference Methods.




Figure 39: Spearman's Footrule Distance of Selection Method (Inference Method: CrowdGauss).




Figure 40: Spearman's Footrule Distance of Selection Methods (Inference Method: CrowdBT).




Figure 41: Spearman's Footrule Distance of Selection Methods (Inference Method: HodgeRank).




Figure 42: Spearman's Footrule Distance of Selection Methods (Inference Method: TrueSkill).