# A Partial-Order-Based Framework for Cost-Effective Crowdsourced Entity Resolution 

CHENGLIANG CHAI • GUOLIANG LI • JIAN LI • DONG DENG • JIANHUA FENG


#### Abstract

Crowdsourced entity resolution has recently attracted significant attentions because it can harness the wisdom of crowd to improve the quality of entity resolution. However existing techniques either cannot achieve high quality or incur huge monetary costs. To address these problems, we propose a cost-effective crowdsourced entity resolution framework, which significantly reduces the monetary cost while keeping high quality. We first define a partial order on the pairs of records. Then we select a pair as a question and ask the crowd to check whether the records in the pair refer to the same entity. After getting the answer of this pair, we infer the answers of other pairs based on the partial order. Next we iteratively select pairs without answers to ask until we get the answers of all pairs. We devise effective algorithms to judiciously select the pairs to ask in order to minimize the number of asked pairs. To further reduce the cost, we propose a grouping technique to group the pairs and we only ask one pair instead of all pairs in each group. We develop error-tolerant techniques to tolerate the errors introduced by the partial order and the crowd. We also study the budget-aware entity resolution, which, given a budget, finds the maximum number of matching pairs within the budget, and propose effective optimization techniques. Experimental results show that our method reduces the cost to $1.25 \%$ of existing approaches (or existing approaches take $80 \times$ monetary cost of our method) while not sacrificing the quality.


## 1 Introduction

Entity resolution aims to find records that refer to the same entity from a collection of records. For example, consider the 11 records in Table $1 r_{1}, r_{2}$ and $r_{3}$ refer to the same entity. $r_{4}, r_{5}, r_{6}$ and $r_{7}$ refer to the same entity. Entity resolu-

[^0]tion has many real-world applications, particularly in health data integration, knowledge-base construction, web search, comparison shopping, and law enforcement.

However existing machine-based methods are still far from perfect [44, 46], because the same entity may have many unpredictable representations. Crowdsourced entity resolution that leverages the crowd's ability to solve this problem has attracted significant attentions[27, 43, 45, 46, 47]. A brute-force method enumerates every pair of records and asks the crowd to check whether they refer to the same entity. This method involves huge monetary costs, especially for large datasets. To address this problem, several algorithms have been proposed to reduce the cost by pruning some pairs that do not need to be asked. Wang et al. 45] utilized the transitivity to reduce the cost, but were unable to provide quality guarantees. This is because the transitivity may not hold for some records, which leads to incorrect deduction and uncontrollable error propagation. Wang et al. [46] proposed a correlation-clustering method, which adaptively assigned the records referring to the same entity into the same cluster. This method improves the quality at the expense of asking more questions and thus involves high monetary costs. In summary, existing methods either cannot achieve high quality or involve huge monetary costs.

To address these problems, we propose Power, a partialorder based crowdsourced entity resolution framework, which significantly reduces the monetary cost while keeping high quality. The basic idea is that we define a partial order on all pairs of records based on the similarity of each pair and prune many pairs that do not need to be asked based on the partial order. Specifically, we first define a partial order: (1) If a pair of records refer to the same entity, then the pairs preceding this pair also refer to the same entity; (2) If a pair of records refer to different entities, then the pairs succeeding this pair refer to different entities. Then we select a pair as a question and ask the crowd to check whether the records in the pair refer to the same entity. Based on the answer of
this pair, we infer the answers of other pairs based on the partial order. Thus our goal is to judiciously select the pairs to ask in order to minimize the number of asked pairs. To this end, we devise effective algorithms to iteratively select pairs without answers to ask until we get the answers of all the pairs. To further reduce the cost, we propose a grouping technique to group the pairs such that we only need to ask one pair instead of all pairs in each group. Since asking only one pair in each iteration leads to a high latency, we propose effective techniques to select multiple pairs in each iteration. As both the partial order and the crowd may introduce errors, we develop error-tolerant techniques to tolerate the errors. We also study the budget-aware entity resolution, which, given a budget, finds the maximum number of matching pairs within the budget, and propose effective algorithms to address this problem.

To summarize, we make the following contributions.
(1) We propose a partial-order based crowdsourced entity resolution framework. We define a partial order on record pairs and utilize the partial order to infer the answers of some unasked pairs so as to reduce the monetary cost.
(2) We construct a graph based on the partial order and utilize the graph to ask questions and infer answers. We devise efficient algorithms to construct the graph. We develop a grouping technique to group the record pairs, which can further reduce the cost. We prove that the optimal grouping is NP-hard and propose approximation algorithms.
(3) We judiciously select pairs to ask in order to minimize the number of asked pairs. We propose a path-based algorithm that asks one question in each iteration and prove that the algorithm is optimal in general. To reduce the latency, we devise a topological-sorting-based algorithm that asks multiple questions in parallel in each iteration.
(4) We develop a probability-based method to tolerate the errors introduced by the crowd and the partial order.
(5) We propose budget-aware algorithms to maximize the number of matching pairs within a given budget.
(6) We conduct experiments using real-world datasets on a real crowdsourcing platform. Experimental results show that our method reduces the cost to $1.25 \%$ of existing approaches (or existing approaches take more than 80 times money of our method) while not sacrificing the quality.

The rest of this paper is structured as follows. We first define the problem and review related work in Section 2 and then propose our framework in Section 3 The grouping strategy, question selection, and error-tolerant techniques are discussed in Sections 4,5,6respectively. We study the budgetaware problem in Section 7 . We report experimental results in Section 8 and conclude in Section 9

## 2 Preliminaries

### 2.1 Problem Definition

Definition 1 (Crowdsourced Entity Resolution) Consider a table $\mathcal{T}$ with $m$ attributes $\left\{\mathcal{A}_{1}, \mathcal{A}_{2}, \ldots, \mathcal{A}_{m}\right\}$ and $n$ records $\left\{r_{1}, r_{2}, \ldots, r_{n}\right\}$, where each record denotes an entity. The entity resolution aims to identify the records that refer to the same entity. Crowdsourced entity resolution leverages the crowd's ability to address this problem.

In most cases, requesters have a budget and aim to identify as many as possible matching pairs within the given budget. Next we formulate the problem of budget-aware crowdsourced entity resolution.

Definition 2 (Budget-Aware Crowdsourced Entity Resolution) Consider a table $\mathcal{T}$ and a budget $B$. The budgetaware crowdsourced entity resolution aims to identify the maximum number of pairs that refer to the same entity within the budget, i.e., asking at most $B$ questions to the crowd.

For example, Table 1 shows a table with 4 attributes and 11 records. $r_{1}, r_{2}$, and $r_{3}$ refer to the same entity. $r_{4}, r_{5}, r_{6}$, and $r_{7}$ refer to the same entity. Each of $r_{8}, r_{9}, r_{10}, r_{11}$ represents a different entity. Crowdsourced entity resolution asks questions to the crowd (or workers) for identifying the records referring to the same entity. As we need to pay the workers for answering a question, the objective is to reduce the number of questions while keeping high quality. Given a budget of asking 2 questions, the budget-aware crowdsourced entity resolution aims to find the maximum number of matching pairs with 2 questions.

We summarize our notations used in our paper in Table 2. We first focus on the first problem and then extend our techniques to support the second problem in Section 7 .

### 2.2 Related Work

### 2.2.1 Machine-based Entity Resolution

There are a bunch of works studying entity resolution with machine-based approaches. For example, Konda at. el [21] propose an end-to-end entity matching system to using machine learning approach. Also, Done et. al [7] identified matching attributes corresponding to the same real-world entity in a database using a dependency graph. .

### 2.2.2 Crowdsourced Entity Resolution

Generating Questions for Workers. An important problem in crowdsourced entity resolution is to design questions for workers. A straightforward method is to generate pair-comparison-based questions, where each question is a pair of two records and asks workers to check whether the two records refer to the same entity. This method may generate a large number of questions. To address this problem, clustering-based questions are proposed [27,44], where each

Table 1: Eleven Records In A Real Restaurant Dataset.

|  | Name $\left(\mathcal{A}_{1}\right)$ | Address $\left(\mathcal{A}_{2}\right)$ | City $\left(\mathcal{A}_{3}\right)$ | Flavor $\left(\mathcal{A}_{4}\right)$ |
| :---: | :---: | :---: | :---: | :---: |
| $r_{1}$ | ritz-carlton restaurant (atlanta) | 181 w. peachtree st. | atlanta | european french |
| $r_{2}$ | ritz-carlton restaurant | 181 peachtree dr | atlanta | european(french) |
| $r_{3}$ | ritz-carlton restaurant Georgia | 181 peachtree st. | city of atlanta | european France |
| $r_{4}$ | cafe ritz-carlton buckhead | 3434 peachtree rd. | city of atlanta | american |
| $r_{5}$ | cafe ritz-carlton (buckhead) | 3434 peachtree rd. | city of atlanta | american |
| $r_{6}$ | dining room ritz-carlton buckhead | 3434 peachtree ave. | atlanta | international |
| $r_{7}$ | dining room ritz-carlton (buckhead) | 3434 peachtree ave. | atlanta | international |
| $r_{8}$ | cafe claude | 20183 rd st. | new york | cafe |
| $r_{9}$ | cafe bizou (american) | 1354 th st. | new york | american food |
| $r_{10}$ | gotham bar \& grill | 12 th rd. | new york | american(new) |
| $r_{11}$ | mesa grill | 102 5th rd. | new york | southwestern |

Table 2: Notations Used In This Paper.

| Notation | Description |
| :---: | :---: |
| $\mathcal{T}=\left\{r_{1}, r_{2}, \ldots, r_{n}\right\}$ | a set or records |
| $\mathcal{A}=\left\{\mathcal{A}_{1}, \mathcal{A}_{2}, \ldots, \mathcal{A}_{m}\right\}$ | a set of attributes |
| $r_{i}[k]$ | value of attribute $\mathcal{A}_{k}$ in record $r_{i}$ |
| $p_{i j}$ | $\left(r_{i}, r_{j}\right)$ |
| $s_{i j}^{k}$ | similarity between $r_{i}[k]$ and $r_{j}[k]$ |
| $\mathcal{G}=(\mathcal{V}, \mathcal{E})$ | a DAG of pairs in $\mathcal{T}$ |
| $\mathcal{G}^{\prime}=\left(\mathcal{V}^{\prime}, \mathcal{E}^{\prime}\right)$ | a grouped DAG of $\mathcal{G}=(\mathcal{V}, \mathcal{E})$ |
| $\succ$ | partial order |
| $\mathcal{C}\left(p_{i j}\right)$ | the child vertex set of $p_{i j}$ |
| $\mathcal{P}\left(p_{i j}\right)$ | the parent vertex set of $p_{i j}$ |
| $\theta$ | the number of incomparable vertices |

question is a group of records and asks workers to classify the records into different clusters such that records in the same cluster refer to the same entity and records in different clusters refer to different entities. As the clustering-based method does not need to enumerate every pair, it can reduce the monetary cost. However, it has no inference power and our method based on the partial order graph can infer much more pairs compared with clustering method.

Pruning Dissimilar Pairs. Intuitively, we do not need to ask the dissimilar pairs that have low probabilities referring to the same entity. Wang et al. [44] proposed a similarity-based method, which computed the similarity of record pairs and pruned the pairs with small similarities. As this method can prune many dissimilar pairs without sacrificing the quality of final answers, most of existing studies used this technique to reduce the cost. We also utilize the method to do pruning in the first step. Afterwards, we utilize the similarity score to do more pruning, which can result in much lower cost.
Leveraging Transitivity to Reduce The Cost. Transitivity can be used to reduce the cost: Given three records, $r_{1}, r_{2}, r_{3}$, if $r_{1}=r_{2}$ ( $r_{1}$ and $r_{2}$ refer to the same entity) and $r_{2}=r_{3}$, we can deduce that $r_{1}=r_{3}$ and do not need to ask whether $r_{1}=r_{3}$. Wang et al. [45] and Vesdapunt et al. [43] studied how to utilize the transitivity to reduce the number of questions. Although this method can reduce the cost, the quality may be reduced. For example, suppose $r_{1}=r_{2}$ and $r_{2} \neq r_{3}$, but the crowd returns $r_{1}=r_{2}$ and $r_{2}=r_{3}$. Then it introduces an incorrect deduction $r_{1}=r_{3}$. Though this
method can reduce the cost, it uses pure crowdsourcing to do that. However, we use partial-order which has higher pruning power.

Improving The Quality. Wang et al. [46] proposed a correlation clustering method, which includes three steps. It first prunes dissimilar pairs with small similarities. Then, it selects some pairs to ask and divides the records into a set of clusters based on the workers' results of these asked pairs. Finally, it refines the clusters by selecting more pairs to ask, checking whether their answers are consistent with the initial clusters, and adjusting the clusters based on the inconsistencies. This method improves the accuracy at the expense of huge monetary costs. We use the confidence value of answers given by different workers for the same question to control the quality. We can see from the evaluation that we also have a good performance in quality control.

Question Selection. A natural problem is how to select next questions to ask in order to improve the quality. Whang et al. [47] proposed a probabilistic model to select high-quality questions. Verroios et al. [42] improved the model by tolerating workers' errors. Gokhale et al. [13] studied the crowdsourced record linkage problem, which linked two records from two tables, which is different from ours as we focus on linking multiple records in the same table.In our framework, we select question based on our graph. We select the optimal questions every time and the result can show that we can save a lot under our question selection strategy.

Compared with existing techniques, our model can significantly reduce the cost while not sacrificing the quality. Compared with the conference version [4], we make the following new contributions. Firstly, we define the budget-based entity resolution problem and propose a budget-aware algorithm. Secondly, we conduct extensive experiments to evaluate our budget-aware algorithms. Experiment results show that our method signi ficantly outperforms existing algorithms, leading by $60 \%-80 \%$ on recall. Thirdly, we add the proofs of all theorems and lemmas. Fourthly, we discuss more related works in this manuscript.

### 2.2.3 Other Related Work

Crowdsourced Operators. There are many studies on leveraging crowd's ability to improve database operators, e.g., crowdsourced selection [32, 31, 38, 3, 53, 23], crowdsourced sort[5, 36, 8, 14, 6, 49, 20], crowdsourced max/top- $k$ [14,41], crowdsourced collection [39, 35], crowdsourced join [45, 43, 46, 13]. They focus on trading-off monetary cost, quality and latency.
Crowdsourced Systems. Several crowdsourced databases, e.g. Deco[33, 34], Quak [28], CrowdDB[11], CDB [24], were proposed, aiming to implement and optimize crowdsourced operators. Li et al. [25] give a survey on crowdsourced data management.
Crowdsourced Quality Control. Many methods are proposed to improve the quality [16, 30, ,9, 26, 56, 55, 54]. Most of these studies focus on devising a worker model to capture worker's quality, computing the worker's model, eliminating bad workers, assigning questions to appropriate workers, and aggregating the results from multiple workers [14, [50, 3, 19, 2, 18, 48, 17, 1, 48, 40, 37, $9,52,51,15]$. They design majority voting or EM algorithm to improve the quality.

## 3 Partial-Order-Based Framework

We first define a partial order (Section 3.1) and then propose a partial-order-based algorithm (Section 3.2).

### 3.1 Partial Order

Record Similarity. Given two records $r_{i}$ and $r_{j}$, we use $p_{i j}$ to denote the pair $\left(r_{i}, r_{j}\right)$ and use $s_{i j}^{k}$ to denote the similarity of $p_{i j}$ on attribute $\mathcal{A}_{k}$. We can utilize any similarity function to compute the similarity, e.g., edit distance, Jaccard, Euclidean distance. Here we take Jaccard and edit similarity as examples. Let $r_{i}[k]$ denote the value of $r_{i}$ on attribute $\mathcal{A}_{k}$. For Jaccard, we tokenize $r_{i}[k]$ into a set of tokens and compute Jaccard on token sets as below.
$s_{i j}^{k}=\mathrm{JAC}\left(r_{i}[k], r_{j}[k]\right)=\frac{\left|r_{i}[k] \cap r_{j}[k]\right|}{\left|r_{i}[k] \cup r_{j}[k]\right|}$,
where $\left|r_{i}[k]\right|$ is the token-set size of $r_{i}[k]$.
For edit similarity, we first compute their edit distance, which is the minimum number of edit operations (insertion, deletion, substitution) required to transform one string to the other, and then compute the edit similarity as below.

$$
\begin{equation*}
\left.s_{i j}^{k}=\operatorname{EDS}\left(r_{i}[k], r_{j}[k]\right)\right)=1-\frac{\operatorname{ED}\left(r_{i}[k], r_{j}[k]\right)}{\max \left(\left|r_{i}[k]\right|,\left|r_{j}[k]\right|\right)}, \tag{2}
\end{equation*}
$$

where $\operatorname{EDS}(E D)$ is the edit similarity (distance) function.
For example, we use the edit similarity on attributes $\mathcal{A}_{1}$ and $\mathcal{A}_{4}$, and Jaccard on attributes $\mathcal{A}_{2}$ and $\mathcal{A}_{3}$. For instance, $s_{12}^{1}=1-\frac{9}{33}=0.72$, and $s_{12}^{2}=\frac{2}{5}=0.4$. As discussed in

Table 3: Record Similarity.

| $p_{i j}$ | $s_{i j}^{1}$ | $s_{i j}^{2}$ | $s_{i j}^{3}$ | $s_{i j}^{4}$ | $p_{i j}$ | $s_{i j}^{1}$ | $s_{i j}^{2}$ | $s_{i j}^{3}$ | $s_{i j}^{4}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $p_{12}$ | 0.72 | 0.4 | 1 | 0.88 | $p_{37}$ | 0.28 | 0.2 | 0.33 | 0 |
| $p_{13}$ | 0.75 | 0.75 | 0.33 | 0.8 | $p_{45}$ | 0.92 | 1 | 1 | 1 |
| $p_{23}$ | 0.77 | 0.5 | 0.33 | 0.69 | $p_{46}$ | 0.69 | 0.5 | 0.33 | 0 |
| $p_{24}$ | 0.51 | 0.2 | 0.33 | 0 | $p_{47}$ | 0.65 | 0.5 | 0.33 | 0 |
| $p_{25}$ | 0.53 | 0.2 | 0.33 | 0 | $p_{56}$ | 0.63 | 0.5 | 0.33 | 0 |
| $p_{26}$ | 0.42 | 0.2 | 1 | 0 | $p_{57}$ | 0.71 | 0.5 | 0.33 | 0 |
| $p_{27}$ | 0.45 | 0.2 | 1 | 0 | $p_{67}$ | 0.94 | 1 | 1 | 1 |
| $p_{34}$ | 0.39 | 0.2 | 1 | 0 | $p_{89}$ | 0.33 | 0.2 | 1 | 0 |
| $p_{35}$ | 0.39 | 0.2 | 1 | 0 | $p_{10,11}$ | 0.5 | 0.25 | 1 | 0 |

Section 2.2, we do not need to consider pairs whose similarities are smaller than a similarity bound $\tau$, as they have small probabilities to be the same entity. Formally, we only consider the similar pair $p_{i j}$ such that (1) $s_{i j}=\mathrm{JAC}\left(r_{i}, r_{j}\right) \geq \tau$ for Jaccard, where $\operatorname{JAC}\left(r_{i}, r_{j}\right)$ is the Jaccard similarity on the token sets of $r_{i}$ and $r_{j}$; or (2) $s_{i j}=\operatorname{EDS}\left(r_{i}, r_{j}\right) \geq \tau$ for edit similarity, where $\operatorname{EDS}\left(r_{i}, r_{j}\right)$ is the edit similarity on records $r_{i}$ and $r_{j}$. The similar record pairs with $\tau=0.2$ are shown in Table 3 If $s_{i j}^{k}<\tau$, we set $s_{i j}^{k}=0$ for simplicity.
Partial Order. We define a partial order on record pairs. Given two pairs $p_{i j}=\left(r_{i}, r_{j}\right), p_{i^{\prime} j^{\prime}}=\left(r_{i^{\prime}}, r_{j^{\prime}}\right), p_{i j} \succeq p_{i^{\prime} j^{\prime}}$, if $\left(r_{i}, r_{j}\right)$ has no smaller similarities than $\left(r_{i^{\prime}}, r_{j^{\prime}}\right)$ on every attribute. $p_{i j} \succ p_{i^{\prime} j^{\prime}}$, if $p_{i j} \succeq p_{i^{\prime} j^{\prime}}$ and $\left(r_{i}, r_{j}\right)$ has larger similarities on at least one attribute than $\left(r_{i^{\prime}}, r_{j^{\prime}}\right)$. Formally,

$$
\begin{gather*}
p_{i j} \succeq p_{i^{\prime} j^{\prime}} \quad \text { if } s_{i j}^{k} \geq s_{i^{\prime} j^{\prime}}^{k} \text { for } 1 \leq k \leq m  \tag{3}\\
p_{i j} \succ p_{i^{\prime} j^{\prime}} \quad \text { if } p_{i j} \succeq p_{i^{\prime} j^{\prime}} \text { and } \exists k, s_{i j}^{k}>s_{i^{\prime} j^{\prime}}^{k} \tag{4}
\end{gather*}
$$

For example, in Table 3 $p_{34} \succeq p_{35}, p_{27} \succ p_{34}, p_{27} \succ p_{35}$.

### 3.2 Graph-Based Algorithm

We model the pairs as a graph based on the partial order.
Definition 3 (Graph Model) Given a table $\mathcal{T}$, we build a directed acyclic graph $\mathcal{G}=(\mathcal{V}, \mathcal{E})$, where each vertex in $\mathcal{V}$ is a similar record pair. Given two pairs $p_{i j}$ and $p_{i^{\prime} j^{\prime}}$, if $p_{i j} \succ p_{i^{\prime} j^{\prime}}$, there is a directed edge in $\mathcal{E}$ from $p_{i j}$ to $p_{i^{\prime} j^{\prime}}$.

Figure 1 shows the graph for the pairs in Table 1 In the figure, we do not show all the edges for illustration purpose: given two vertices, if there is already a path between them, we do not show the direct edge between them. For example, there should be an edge between $p_{67}$ and $p_{12}$, but we omit it as there is already a path from $p_{67}$ to $p_{12}$.
Graph Coloring. Each vertex in $\mathcal{G}$ has two possibilities: (1) they refer to the same entity and we color it Green; (2) they refer to different entities and we color it RED. Initially each vertex is uncolored. Our goal is to utilize the crowd to color all vertices. A straightforward method is to take the record pair on each vertex as a question and ask workers to answer the question, i.e. whether the two records in the pair refer to the same entity. If a worker thinks that the two records on the vertex refer to the same entity, the worker returns Yes; No


Fig. 1: Partial Order and Graph Model.
otherwise. For each pair, to tolerate the noisy results from workers, we assign it to multiple workers, say 5 . Based on the workers' results, we get a voted answer on each vertex. If majority workers vote Yes, we color it Green; otherwise we color it RED. Next, we interchangeably use vertex, pair and question if the context is clear.

Obviously this method is rather expensive as there are many vertices on the graph. To address this issue, we propose an effective coloring framework to reduce the number of questions. Algorithm 1 shows the pseudo code. It first computes the partial orders between pairs and constructs a graph (line 1). Then it selects an uncolored vertex $p_{i j}$ (line 3 ) and asks workers to answer Yes or No on the vertex,
(1) If majority workers vote Yes, we not only color $p_{i j}$ Green, but also color all of its ancestors Green (line 5]. In other words, for $p_{i^{\prime} j^{\prime}} \succ p_{i j}$, we also take $r_{i^{\prime}}$ and $r_{j^{\prime}}$ as the same entity. This is because $p_{i^{\prime} j^{\prime}}$ has larger similarity on every attribute than $p_{i j}$, and since $r_{i}$ and $r_{j}$ refer to the same entity (denoted by $r_{i}=r_{j}$ ), we deduce that $r_{i^{\prime}}=r_{j^{\prime}}$.
(2) If majority workers vote No, we not only color $p_{i j}$ RED, but also color all of its descendants RED (line 7). In other words, for $p_{i j} \succ p_{i^{\prime} j^{\prime}}$, we also take $r_{i^{\prime}}$ and $r_{j^{\prime}}$ as different entities. This is because $p_{i^{\prime} j^{\prime}}$ has smaller similarity on every attribute than $p_{i j}$, and since $r_{i}$ and $r_{j}$ refer to different entities (denoted by $r_{i} \neq r_{j}$ ), we deduce that $r_{i^{\prime}} \neq r_{j^{\prime}}$.

If all the vertices have been colored, the algorithm terminates (line 4); otherwise, it selects an uncolored vertex and repeats the above steps (lines 2/7).

Obviously, this method can reduce the cost as we can avoid asking many unnecessary vertices. For example, consider the constructed graph in Figure 1. A naive method is to ask all eighteen pairs. However, if we first ask $p_{10,11}$, as majority workers vote No, we can color $p_{10,11}$ and its descendants $p_{27}, p_{26}, p_{34}, p_{35}, p_{89}$ and $p_{37}$ RED without needing to ask these descendants. Then if we select $p_{56}$, as majority workers vote Yes, we color $p_{56}$ and its ancestors $p_{46}, p_{47}$, $p_{57}, p_{23}, p_{45}, p_{67}$ and $p_{13}$ GREEN without needing to ask them. In Section 5, we will show that we need to ask at least 4 questions (e.g., $p_{12}, p_{10,11}, p_{25}, p_{56}$ ) to color all vertices.

There are several challenges in this algorithm.

```
Algorithm 1: A Partial-Order-Based Framework
    Input: \(\mathcal{T}=\left\{r_{1}, r_{2} \cdots, r_{n}\right\}\)
    Output: All vertices are colored as Green or Red
    Construct \(\mathcal{G}=(\mathcal{V}, \mathcal{E})\) based on partial orders;
    while there exist uncolored vertices in \(\mathcal{V}\) do
        Select an uncolored vertex \(p_{i j}\) to ask workers;
        if majority workers vote Yes then
            color \(p_{i j}\) and \(p_{i^{\prime} j^{\prime}}\left(p_{i^{\prime} j^{\prime}} \succ p_{i j}\right)\) Green;
        else
            color \(p_{i j}\) and \(p_{i^{\prime} j^{\prime}}\left(p_{i j} \succ p_{i^{\prime} j^{\prime}}\right)\) RED;
    8 return colored \(\mathcal{V}\);
```

(1) Graph Construction. As there are large numbers of pairs, how to efficiently construct the graph? Can we reduce the graph size so as to reduce the number of questions?
(2) Question Selection. How to select the minimum number of vertices to ask in order to color all vertices?
(3) Error Tolerant. The coloring strategy and the workers may introduce errors. So how to tolerate the errors?

We address these challenges in the following sections.

## 4 Graph Construction

We first propose efficient graph-construction algorithms (Section 4.1) and then present grouping methods (Section 4.2).

### 4.1 Graph Construction Algorithms

Brute-Force Method. It enumerates every pair of vertices and checks whether they satisfy the partial order. If so, the algorithm adds an edge between them. The complexity of this method is $\mathcal{O}\left(|\mathcal{V}|^{2}\right)$. Obviously this method is rather expensive, especially if there are a large number of vertices.
Quicksort-Based Method. Quicksort is an efficient algorithm for the sorting problem and it can be extended to construct the graph. We first randomly select a vertex $p_{i j}$ as pivot, and then split other vertices into three disjoint parts by comparing them with $p_{i j}$ :
(1) Parent Vertex Set: $\mathcal{P}\left(p_{i j}\right)=\left\{p_{i^{\prime} j^{\prime}} \mid p_{i^{\prime} j^{\prime}} \succ p_{i j}\right\}$. For each $p_{i^{\prime} j^{\prime}}$ in $\mathcal{P}\left(p_{i j}\right)$, we add an edge from $p_{i^{\prime} j^{\prime}}$ to $p_{i j}$;
(2) Child Vertex Set: $\mathcal{C}\left(p_{i j}\right)=\left\{p_{i^{\prime} j^{\prime}} \mid p_{i j} \succ p_{i^{\prime} j^{\prime}}\right\}$. For each $p_{i^{\prime} j^{\prime}}$ in $\mathcal{C}\left(p_{i j}\right)$, we add an edge from $p_{i j}$ to $p_{i^{\prime} j^{\prime}}$;
(3) Incomparable Vertex Set: $\mathcal{U}\left(p_{i j}\right)=\mathcal{V}-\mathcal{P}\left(p_{i j}\right)-\mathcal{C}\left(p_{i j}\right)=$ $\left\{p_{i^{\prime} j^{\prime}} \mid p_{i j} \nsucc p_{i^{\prime} j^{\prime}} \& p_{i^{\prime} j^{\prime}} \nsucc p_{i j}\right\}$. For each $p_{i^{\prime} j^{\prime}}$, there is no edge between $p_{i j}$ and $p_{i^{\prime} j^{\prime}}$, as they are incomparable.

Obviously, $\forall p \in \mathcal{P}\left(p_{i j}\right), p^{\prime} \in \mathcal{C}\left(p_{i j}\right), p \succ p^{\prime}$, and thus we do not need to compare the pairs in $\mathcal{P}\left(p_{i j}\right) \times \mathcal{C}\left(p_{i j}\right)$. Then, we consider the pairs in $\left(\mathcal{P}\left(p_{i j}\right) \cup \mathcal{U}\left(p_{i j}\right)\right) \times\left(\mathcal{P}\left(p_{i j}\right) \cup\right.$ $\left.\mathcal{U}\left(p_{i j}\right)\right)$ and $\left(\mathcal{C}\left(p_{i j}\right) \cup \mathcal{U}\left(p_{i j}\right)\right) \times\left(\mathcal{C}\left(p_{i j}\right) \cup \mathcal{U}\left(p_{i j}\right)\right)$. To add edges between these pairs, we can recursively utilize the above method ${ }^{\top}$ The worst-case complexity of this method

[^1]is also $\mathcal{O}\left(|\mathcal{V}|^{2}\right)$ if all the vertices are incomparable. However, this method has better performance than brute-force in practice, because it can prune many unnecessary pairs (e.g., $\left.\mathcal{P}\left(p_{i j}\right) \times \mathcal{C}\left(p_{i j}\right)\right)$.
Index-Based Method. The quicksort-based method still has poor performance for large datasets. To address this issue, we propose an index-based method. As the similarity $s_{i j}^{k}$ is a numerical value, we can utilize geometric relationship to compare two pairs. For simplicity, we first assume there are two attributes $(m=2)$. So the similarity of $p_{i j}$ has two components $s_{i j}^{1}$ and $s_{i j}^{2}$. Therefore, we can map each vertex to a point in a two-dimensional coordinate as shown in Figure 2(a)

If we want to find the child set of $p_{i j}, \mathcal{C}\left(p_{i j}\right)=\left\{p_{i^{\prime} j^{\prime}} \mid p_{i j} \succ\right.$ $\left.p_{i^{\prime} j^{\prime}}\right\}$, we report the left-bottom vertices (i.e., vertices in the rectangle). Similarity, if we compute $\mathcal{P}\left(p_{i j}\right)=\left\{p_{i^{\prime} j^{\prime}} \mid p_{i^{\prime} j^{\prime}} \succ\right.$ $\left.p_{i j}\right\}$, we report the top-right vertices. We can utilize the 2dimensional range trees to achieve this goal [22].
Range Search Tree Construction. We first construct a firstlevel balanced binary tree based on $s_{i j}^{1}$ for all vertices as shown in Figure 2(b), where leaves are vertices in $\mathcal{V}$ and the internal nodes are guided search values. (There are multiple pairs in a node because they have the same similarity. For example, $p_{34}, p_{35}$ are in the same node, because $s_{34}^{1}=s_{35}^{1}=0.39$.) The value of a node is the largest $s_{i j}^{1}$ for all vertices in its left subtree, and thus the $s_{i j}^{1}$ values of vertices under the left subtree are not larger than the value of this node; while the $s_{i j}^{1}$ values of vertices under its right subtree are larger than the value. We can build the binary tree in a bottom-up way. For each internal node, we construct the second-level balanced binary tree based on $s_{i j}^{2}$ for vertices under this node.
Reporting $\mathcal{C}\left(p_{i j}\right)$ with Range Search Tree. Given a vertex $p_{i j}$, we use the range search tree to report $\mathcal{C}\left(p_{i j}\right)$. We first find the tree nodes whose descendants' similarities on $\mathcal{A}_{1}$ are not larger than $s_{i j}^{1}$ using the first-level tree. For each of such qualified nodes on $\mathcal{A}_{1}$, we visit its second-level tree and find the nodes whose descendants' similarities on $\mathcal{A}_{2}$ are not larger than $s_{i j}^{2}$. Then the vertices under these nodes are added into $\mathcal{C}\left(p_{i j}\right)$. Next we discuss how to find such qualified nodes in the first-level tree and the same techniques can be used to search the second-level tree.

To find the qualified nodes on $\mathcal{A}_{1}$, we search the firstlevel tree from the root. For each node, (1) If its value is not larger than $s_{i j}^{1}$, (1.1) if it is a leaf, it is a qualified node; (1.2) if it is not a leaf, the similarities of all the vertices under its left child on attribute $\mathcal{A}_{1}$ are not larger than $s_{i j}^{1}$, and its left child is a qualified node. Next we recursively process its right child; (2) If its value is larger than $s_{i j}^{1}$, (2.1) if it is a leaf, we prune it; (2.2) if it is not a leaf, we prune its right subtree as the similarities of all the vertices under its right child on attribute $\mathcal{A}_{1}$ must be larger than $s_{i j}^{1}$. Next we recursively process its left child. Iteratively, we can iden-
tify all qualified nodes on $\mathcal{A}_{1}$. This method accesses at most $\log (|\mathcal{V}|)$ nodes in the first-level tree.

For example, suppose we want to compute $\mathcal{C}\left(p_{12}\right)$ where $s_{12}^{1}=0.72$ and $s_{12}^{2}=0.4$. We first compare $s_{12}^{1}$ with the root $s_{56}^{1}=0.63$. As $s_{12}^{1}>s_{56}^{1}$, its left child (i.e., $p_{26}$ ) is a qualified node. Next we go to the right child $p_{12}$. As $s_{12}^{1}=$ $s_{12}^{1}$, we visit its left child $p_{46}$. As $s_{12}^{1}>s_{46}^{1}$, its left child $p_{47}$ is a qualified node. Next we go to the right child $p_{57}$. As $s_{12}^{1}>s_{57}^{1}$, its left child $p_{57}$ is a qualified node and we go to its right node $p_{12}$. As $p_{12}$ is a leaf, it is a qualified node. Next for each qualified node ( $p_{26}, p_{47}, p_{57}, p_{12}$ ), we check it on the second attribute. Take $p_{26}$ as an example. As $s_{12}^{2}=0.4$ is larger than the root's value, its left child $p_{37}$ is a qualified node. We then visit its right child $p_{56}$. As $s_{12}^{2} \leq s_{56}^{2}$, we go to its left child which is a leaf. As the value is larger than $s_{12}^{2}$, we prune it. Thus the pairs under node $p_{37}$ are added into $\mathcal{C}\left(p_{12}\right)$.
Building The Graph with Range Search Tree. For each vertex $p_{i j}$, we use the range search tree to find $\mathcal{C}\left(p_{i j}\right)$ and add vertices in $\mathcal{C}\left(p_{i j}\right)$ as the children of $p_{i j}$. Then we can build the graph. It is straightforward to generalize 2-dimensional range trees to $m$-dimensional range trees.
Complexity. Both the time and space complexities of constructing the tree is $\mathcal{O}\left(|\mathcal{V}| \log ^{m-1}|\mathcal{V}|\right)$. The time complexity of computing $\mathcal{C}\left(p_{i j}\right)$ is $\mathcal{O}\left(\log ^{m}|\mathcal{V}|+\left|\mathcal{C}\left(p_{i j}\right)\right|\right)$, where $\left|\mathcal{C}\left(p_{i j}\right)\right|$ is the size of $\mathcal{C}\left(p_{i j}\right)$. After using the fractional cascading technique [22], the complexity is reduced to $\mathcal{O}\left(\log ^{m-1}\right.$ $\left.|\mathcal{V}|+\left|\mathcal{C}\left(p_{i j}\right)\right|\right)$. Thus the overall time complexity of constructing the graph is $\mathcal{O}\left(|\mathcal{V}| \log ^{m-1}|\mathcal{V}|+|\mathcal{E}|\right)$.

### 4.2 Vertex Grouping

Note that some vertices have very close similarities and we can combine them to reduce the graph size, which not only reduces the cost but also saves the graph construction cost. For example, $p_{67}$ and $p_{45}$ have close similarities on the four attributes, i.e., $p_{67}:(0.94,1,1,1)$ and $p_{45}:(0.92,1,1,1)$ as shown in Table 3. Thus we can combine them as a single vertex. Next we formulate the problem.

Definition 4 (Vertex Group) Given a threshold $\varepsilon$, a subset $g \subseteq \mathcal{V}$ is called a vertex group, if for any pairs $p_{i j}$ and $p_{i^{\prime} j^{\prime}}$ in $g,\left|s_{i j}^{k}-s_{i^{\prime} j^{\prime}}^{k}\right| \leq \varepsilon$ for $1 \leq k \leq m$.

As the similarities between different pairs in a group should not have large gap, we use $\varepsilon$ to set a constraint. For example, suppose $\varepsilon=0.1$. $\left\{p_{26}, p_{34}, p_{35}\right\}$ is a group as the difference of their similarities on every attribute is smaller than 0.1 ( $\left.p_{26}:(0.42,0.2,1,0), p_{34}:(0.39,0.2,1,0), p_{35}:(0.39,0.2,1,0)\right)$.

Next we partition the vertices into different groups.
Definition 5 (Grouping Strategy) Given a set of vertices $\mathcal{V}$, a grouping strategy is a partition of $\mathcal{V}$ to generate a set of groups $g_{1}, g_{2}, \ldots, g_{x}$, which satisfies,

(a) $2 d$ coordinate
(1) Complete: For any $p_{i j} \in \mathcal{V}, \exists g_{t}, p_{i j} \in g_{t}$; and
(2) Disjoint: For any two groups $g_{i}, g_{j}, g_{i} \cap g_{j}=\phi$.

For example, consider the eighteen pairs in Table 1 Given threshold $\varepsilon=0.1$, the groups $\left\{p_{67}, p_{45}\right\},\left\{p_{12}\right\},\left\{p_{13}\right\},\left\{p_{23}\right\}$, $\left\{p_{10,11}, p_{27}\right\},\left\{p_{57}, p_{47}, p_{46}, p_{56}\right\},\left\{p_{24}, p_{25}\right\},\left\{p_{26}, p_{34}, p_{89}, p_{35}\right\}$ $\left\{p_{37}\right\}$ satisfy the two constraints.
Partial Order on Groups. We can define the partial order on the groups. For any two groups $g_{i}$ and $g_{j}$,

$$
\begin{align*}
g_{i} \succeq g_{j} & \text { if } \forall p \in g_{i}, p^{\prime} \in g_{j}, p \succeq p^{\prime}  \tag{5}\\
g_{i} \succ g_{j} & \text { if } \forall p \in g_{i}, p^{\prime} \in g_{j}, p \succ p^{\prime} \tag{6}
\end{align*}
$$

Let $g^{k} . l / g^{k} . u$ denote the smallest/largest similarity of pairs in $g$ on $\mathcal{A}_{k}$, i.e., $g^{k} . l=\min _{p_{i j} \in g} s_{i j}^{k}$ and $g^{k} . u=$ $\max _{p_{i j} \in g} s_{i j}^{k}$. We can prove that if $g_{i}^{k} \cdot l \geq g_{j}^{k}$.u for $1 \leq k \leq$ $m, g_{i} \succeq g_{j}$; if $g_{i}^{k} . l \geq g_{j}^{k} . u$ and $\exists k g_{i}^{k} \cdot l>g_{j}^{k} . u, g_{i} \succ g_{j}$. Thus we can use $g_{i}^{k} . l$ and $g_{j}^{k}$.u to easily determine the partial orders of two groups. Given a set of groups, if $g_{i} \succ g_{j}$, we add an edge from $g_{i}$ to $g_{j}$. Then we can construct a grouped graph.

Definition 6 (Grouped Graph) Given a set of vertices $\mathcal{V}$ and a set of groups $g_{1}, g_{2}, \ldots, g_{x}$ generated using the grouping strategy, we construct a grouped graph $\mathcal{G}^{\prime}=\left(\mathcal{V}^{\prime}, \mathcal{E}^{\prime}\right)$, where each vertex in $\mathcal{V}^{\prime}$ is a group, and there is an edge in $\mathcal{E}^{\prime}$ from $g_{i}$ to $g_{j}$ if $g_{i} \succ g_{j}$.

Coloring The Grouped Graph. We ask workers to color the grouped graph. If a group is selected to ask, we randomly select a pair in the group and take the answer of this pair as the answer of the group. Then we utilize our coloring algorithm (Section 3.2) to color the grouped graph.
Optimal Group Generation. There are multiple grouping strategies. We quantify how good a grouping strategy is. Obviously, the smaller the number of vertices in the grouped graph is, the lower the cost is. Thus we aim to generate the minimum number of groups.

Definition 7 (Optimal Group Generation). Given a set of vertices $\mathcal{V}$ and a threshold $\varepsilon$, we aim to generate the minimum number of groups.


Fig. 4: The Group Tree.
We can prove that the optimal group generation problem is NP-hard as proved in Theorem 1 .

Theorem 1 The optimal group generation problem is NPHard.

Proof We prove the problem is NP-Hard even $m=2$ by a reduction from the following rectangle cover problem. In a rectangle cover instance, we are given a set of points in the Euclidean plane $\mathbb{R}^{2}$. Our goal is to use the minimum number of unit squares to cover all points. The problem is known to be NP-Hard [29]. In our problem, it is easy to see a vertex group can be covered by a square of side length $\epsilon$. We can partition the set of vertices into $k$ groups, if and only if all vertices can be covered by $k$ squares of side length $\epsilon$. Therefore, our problem is equivalent to the rectangle cover problem, thus is NP-Hard as well.

We propose a greedy algorithm and a heuristic algorithm.
Greedy Algorithm. The basic idea is that we first generate all the maximal groups, which are defined as below.

Definition 8 (Maximal Group) A group $g$ is called a maximal group if $\forall p_{i j} \in \mathcal{V}-g, g \cup\left\{p_{i j}\right\}$ is not a group (i.e., it does not satisfy the $\varepsilon$-constraint in Definition 4 ).

```
Algorithm 2: Vertex Grouping: Greedy
    Input: \(\mathcal{G}=(\mathcal{V}, \mathcal{E})\)
    Output: A set of groups \(g_{1}, g_{2}, \ldots, g_{x}\)
    Generate maximal groups \(\mathcal{M}\);
    while \(\mathcal{M}\) is not empty do
        Pick the largest group \(g\) from \(\mathcal{M}\);
        for each \(g_{i}\) in \(\mathcal{M}\) do
            \(g_{i}=g_{i}-g\)
    return \(g_{1}, g_{2}, \ldots, g_{x}\)
```

For example, $\left\{p_{26}, p_{34}, p_{35}\right\}$ is a group, but it is not a maximal group, because if we add $p_{89},\left\{p_{26}, p_{34}, p_{35}, p_{89}\right\}$ is still a group satisfying Definition 4, which contradicts with Definition $8\left\{p_{26}, p_{34}, p_{35}, p_{89}\right\}$ is a maximal group, as we cannot add any pair to form a new group.

Next we introduce a greedy algorithm. At the beginning of the algorithm, we need to compute the set of all maximal groups. Next we show how to solve this problem.
Generating Maximal Groups. We first consider the onedimensional case, i.e., $m=1$. We generate all the maximal groups based on $s_{i j}^{1}$. We first sort $p_{i j}$ based on $s_{i j}^{1}$ in a descending order, denoted by $p_{1}, p_{2}, \ldots, p_{n}$. For the first pair $p_{1}$, we generate a longest group $\left\{p_{1}, p_{2}, \ldots, p_{t}\right\}$ where $p_{1}-p_{t} \leq \varepsilon$ and $p_{1}-p_{t+1}>\varepsilon$. Obviously this longest group is a maximal group. Next we generate the longest group for $p_{2}$. If the longest group of $p_{2}$ is not contained by that of $p_{1}$, it is a maximal group. Iteratively we can generate all the maximal groups. The complexity is $\mathcal{O}\left(|\mathcal{V}|^{2}\right)$.

For the $m$-dimensional case, we first generate the maximal groups $\mathcal{M}^{i}$ on every attribute $\mathcal{A}_{i}$. Then we join them to generate the maximal groups, i.e., $\mathcal{M}^{1} \bowtie \mathcal{M}^{2} \bowtie \cdots \bowtie$ $\mathcal{M}^{m}=\left\{\mathcal{M}_{i_{1}}^{1} \cap \mathcal{M}_{i_{2}}^{2} \cap \cdots \cap \mathcal{M}_{i_{m}}^{m}\right\}$ where $1 \leq i_{j} \leq$ $\left|\mathcal{M}^{j}\right|$. We prove that the generated groups contain all maximal groups. Then we utilize these groups to run the greedy algorithm.
Theorem $2 \mathcal{M}^{1} \bowtie \mathcal{M}^{2} \bowtie \cdots \bowtie \mathcal{M}^{m}=\left\{\mathcal{M}_{i_{1}}^{1} \cap \mathcal{M}_{i_{2}}^{2} \cap\right.$ $\left.\cdots \cap \mathcal{M}_{i_{m}}^{m}\right\}$ contains all maximal groups.
Proof We prove that for any maximal group, $g$, there exist $\mathcal{M}_{i_{1}}^{1}, \mathcal{M}_{i_{2}}^{2}, \cdots, \mathcal{M}_{i_{m}}^{m}, g=\mathcal{M}_{i_{1}}^{1} \cap \mathcal{M}_{i_{2}}^{2} \cap \cdots \cap \mathcal{M}_{i_{m}}^{m}$. As $g$ is a maximal group, $g_{i}^{k} \cdot u-g_{i}^{k} \cdot l \leq \varepsilon$ for any attribute $\mathcal{A}_{k}$. Let $s_{i j}^{k}=g_{i}^{k}$.l. We generate the maximal group $\mathcal{M}_{i_{k}}^{k}$ on attribute $\mathcal{A}_{k}$ based on $s_{i j}^{k}$. Obviously $g \subseteq \mathcal{M}_{i_{k}}^{k}$. Thus $g \subseteq \mathcal{M}_{i_{1}}^{1} \cap \mathcal{M}_{i_{2}}^{2} \cap \cdots \cap \mathcal{M}_{i_{m}}^{m}$. As $g$ is a maximal group, $g=\mathcal{M}_{i_{1}}^{1} \cap \mathcal{M}_{i_{2}}^{2} \cap \cdots \cap \mathcal{M}_{i_{m}}^{m}$.

Algorithm 2 shows the pseudo code. It first generates all the maximal groups (line 1), and then greedily picks the largest group (line 3). Finally it updates other groups by removing the vertices in the largest group (line 5).

We first generate the set of all maximal groups, denoted by $\mathcal{M}$. Then we greedily pick the largest group $g$ in $\mathcal{M}$ with the maximum number of vertices. For each $g_{i}$ in $\mathcal{M}$, we remove the vertices in $g$ from $g_{i}$ and update $g_{i}$ to $g_{i}-g$. (If $g_{i}-g$ is empty, we remove it from $\mathcal{M}$.) Next we iteratively pick
the largest group from $\mathcal{M}$ until $\mathcal{M}$ is empty. This greedy algorithm has a $\ln (|\mathcal{V}|)$ approximation ratio. However it is expensive to generate the maximal groups and the complexity of this greedy algorithm is $\mathcal{O}\left(|\mathcal{V}|^{m}\right)$.

For example, we want to group the vertices in Figure 1. Firstly, we generate all the maximal groups $\mathcal{M}=\left\{\left\{p_{67}\right.\right.$, $\left.p_{45}\right\},\left\{p_{12}\right\},\left\{p_{13}\right\},\left\{p_{23}\right\},\left\{p_{10,11}, p_{27}, p_{26}\right\},\left\{p_{27}, p_{26}, p_{34}\right.$, $\left.p_{35}\right\},\left\{p_{26}, p_{34}, p_{35}, p_{89}\right\},\left\{p_{47}, p_{57}, p_{46}, p_{56}\right\},\left\{p_{24}, p_{25}\right\}$, $\left.\left\{p_{37}\right\}\right\}$. Then we select the largest group $\left\{p_{27}, p_{26}, p_{34}, p_{35}\right\}$ from the maximal group set as a group. Next we remove vertices in it from other maximal groups. Now $\mathcal{M}=\left\{\left\{p_{67}\right.\right.$, $\left.p_{45}\right\},\left\{p_{12}\right\},\left\{p_{13}\right\},\left\{p_{23}\right\},\left\{p_{10,11}\right\},\left\{p_{89}\right\},\left\{p_{47}, p_{57}, p_{46}\right.$, $\left.\left.p_{56}\right\},\left\{p_{24}, p_{25}\right\},\left\{p_{37}\right\}\right\}$. Then we select the largest group. Finally the groups are $\mathcal{M}=\left\{\left\{p_{67}, p_{45}\right\},\left\{p_{12}\right\},\left\{p_{13}\right\},\left\{p_{23}\right\}\right.$, $\left\{p_{10,11}\right\},\left\{p_{27}, p_{26}, p_{34}, p_{35}\right\},\left\{p_{89}\right\},\left\{p_{47}, p_{57}, p_{46}, p_{56}\right\}$, $\left.\left\{p_{24}, p_{25}\right\},\left\{p_{37}\right\}\right\}$.
Split-Based Algorithm. As the greedy algorithm is expensive, we propose an efficient algorithm. The basic idea is that we first take all the pairs as a group, and if any attribute does not satisfy the threshold constraint, we partition the group based on this attribute. The pseudo code is shown in Algorithm 3 . Formally, we build a tree structure and the root is $\mathcal{N}_{1}=\mathcal{V}$. Let $\mathcal{N}_{1}^{i} . l / \mathcal{N}_{1}^{i} . u$ denote the minimal/maximal similarity of pairs in $\mathcal{N}_{1}$ on attribute $\mathcal{A}_{i}$. If $\mathcal{N}_{1}^{i} \cdot u-\mathcal{N}_{1}^{i} . l>\varepsilon$, we split $\mathcal{N}_{1}$ based on $\mathcal{A}_{i}$ and generate two ranges $\left[\mathcal{N}_{1}^{i} \cdot l, \frac{\mathcal{N}_{1}^{i} . l+\mathcal{N}_{1}^{i} \cdot u}{2}\right],\left(\frac{\mathcal{N}_{1}^{i} . l+\mathcal{N}_{1}^{i} \cdot u}{2}, \mathcal{N}_{1}^{i} . u\right]$; otherwise, we do not split $\mathcal{N}_{1}$ based on this attribute. Suppose we split $\mathcal{N}_{1}$ based on $\mathcal{A}_{i_{1}}, \mathcal{A}_{i_{1}}, \ldots, \mathcal{A}_{i_{t}}$. We generate $2^{t}$ children of $\mathcal{N}_{1}$ by enumerating the two ranges of these attributes. For each node, we add the pairs that fall in the corresponding ranges into the node. If a node cannot be split on any attribute, it is a leaf. Finally the groups on leaves are the result.

For example, we walk through our algorithm on the records in Table 33. Suppose $\varepsilon=0.1$. Figure 4 shows the group tree. Firstly, the root $\mathcal{N}_{1}\left(\left[\mathcal{N}_{1}^{1} . l, \mathcal{N}_{1}^{1} \cdot u\right],\left[\mathcal{N}_{1}^{2} . l, \mathcal{N}_{1}^{2} . u\right],\left[\mathcal{N}_{1}^{3} . l, \mathcal{N}_{1}^{3} . u\right]\right.$, $\left.\left[\mathcal{N}_{1}^{4} . l, \mathcal{N}_{1}^{4} . u\right]\right)$ is denoted as $([0.28,0.94],[0.2,1],[0.33,1]$, $[0,1]$ ) in Figure 4 As $\mathcal{N}_{1}^{i} \cdot u-\mathcal{N}_{1}^{i} . l>\varepsilon$ for $i \in[1,4]$, we split $\left[\mathcal{N}_{1}^{1} \cdot l, \mathcal{N}_{1}^{1} \cdot u\right],\left[\mathcal{N}_{1}^{2} \cdot l, \mathcal{N}_{1}^{2} \cdot u\right],\left[\mathcal{N}_{1}^{3} \cdot l, \mathcal{N}_{1}^{3} \cdot u\right]$ and $\left[\mathcal{N}_{1}^{4} \cdot l, \mathcal{N}_{1}^{4} \cdot u\right]$ into $\langle[0.28,0.61],(0.61,0.94]\rangle,\langle[0.2,0.6],(0.6,1]\rangle,\langle[0.33,0.67]$, $(0.67,1]\rangle$ and $\langle[0,0.5](0.5,1]\rangle$ respectively. Then we move each pair in $\mathcal{N}_{1}$ into the $2^{4}$ children (empty children are removed). For $i \in[1,4], s_{45}^{i}$ and $s_{67}^{i}$ are in the range of $(0.61,0.94],(0.6,1],(0.67,1],(0.5,1]$, and $p_{45}$ and $p_{67}$ are added into $\mathcal{N}_{4}$. Then we calculate $\mathcal{N}_{4}^{i} . l, \mathcal{N}_{4}^{i} . u$ and get $([0.92,0.94]$, $[1,1],[1,1],[1,1])$. As each range is smaller than $\varepsilon, \mathcal{N}_{4}=$ $\left\{p_{45}, p_{67}\right\}$ is a leaf. Next, we move $\left\{p_{24}, p_{25}, p_{37}\right\}$ into $\mathcal{N}_{5}$ ([0.28, 0.53], [0.2, 0.2], [0.33, 0.33], [0, 0]). It is not a group and split again. As $\left[\mathcal{N}_{5}^{2} \cdot u-\mathcal{N}_{5}^{2} \cdot l\right]<\varepsilon,\left[\mathcal{N}_{5}^{3} \cdot u-\mathcal{N}_{5}^{3} \cdot l\right]<\varepsilon$ and $\left[\mathcal{N}_{5}^{4} \cdot u-\mathcal{N}_{5}^{4} \cdot l\right]<\varepsilon$, we split $\mathcal{N}_{5}^{1}$ and get two leaves $\mathcal{N}_{9}$ and $\mathcal{N}_{10}$. At last, we get 9 groups (as shown in Figure 3).
Complexity. At the root of the tree, the maximum similarity interval of each attribute is 1 . Each time we partition the at-

```
Algorithm 3: Vertex Grouping: Split
    Input: \(\mathcal{G}=(\mathcal{V}, \mathcal{E})\)
    Output: A set of groups \(g_{1}, g_{2}, \ldots, g_{x}\)
    \(\mathcal{N}_{1} \leftarrow \mathcal{V}\); Priority queue \(Q=\left\{\mathcal{N}_{1}\right\} ;\)
    while \(Q\) is not empty do
        Pop node \(\mathcal{N}_{i}\) from \(Q\)
        for \(k \in[1, m]\) do
            if \(\mathcal{N}_{i}^{k} \cdot u-\mathcal{N}_{i}^{k} . l>\varepsilon\) then
                Split \(\mathcal{N}_{i}\) based on \(\mathcal{A}_{k}\);
        if \(\mathcal{N}_{i}\) is split by \(\mathcal{A}_{i_{1}}, \mathcal{A}_{i_{2}}, \ldots, \mathcal{A}_{i_{+}}\)then
            Generate \(2^{t}\) children of \(\mathcal{N}_{i}\);
            Move pairs in \(\mathcal{N}_{i}\) into corresponding children;
            Add these children into \(Q\);
        else
            \(\mathcal{N}_{i}\) is a leaf and taken as a group \(g ;\)
    return the groups on the leaves;
```

tribute, the interval is reduced by half. We stop the algorithm until $\frac{1}{2^{H-1}} \leq \epsilon$, where $H$ is the number of levels of the tree. Therefore, the tree has at most $\left(\log \frac{1}{\varepsilon}+1\right)$ levels. Thus the time complexity of constructing the tree is $\mathcal{O}\left(|\mathcal{V}| \log \frac{1}{\varepsilon}\right)$.
Remark. In the experiment, we evaluate the quality, cost and latency by varying the vertex grouping threshold $\epsilon$. Our method Power is not sensitive to the threshold much on quality. For cost, our algorithm with small threshold only took slightly more cost than the algorithm with larger $\epsilon$ but it was still much less than state-of-the-art algorithms. Therefore, if one cares about the cost, one can choose a large threshold like 0.3 , which may sacrifice the quality a little. Otherwise, one can choose a small threshold like 0.1 , which can achieve a high recall.

## 5 Question Selection

An important problem is to select the minimum number of vertices as questions to color all vertices. We first formulate the question-selection problem (Section 5.1.) and then propose a serial algorithm that selects one vertex in each iteration (Section 5.2 ) and parallel algorithms that select multiple vertices in each iteration (Section 5.3).

### 5.1 Optimal Vertex Selection

We first assume that (1) if a vertex is Green, then all of its ancestors are Green; and (2) if a vertex is Red, then all of its descendants are RED. We will discuss how to support the case that the two conditions do not hold in Section 6
Definition 9 (Optimal Graph Coloring) Given a graph, the optimal graph coloring problem aims to select the minimum number of vertices as questions to color all the vertices using the coloring strategy.

For example, in Figure 3, if we sequentially select vertices $g_{8}, g_{7}, g_{5}, g_{2}, g_{3}, g_{4}$ and $g_{6}$, we ask 7 questions. The
optimal crowdsourced vertices are $g_{2}, g_{5}, g_{6}$ and $g_{8}$ (highlighted by bold circles), because the colors of these vertices cannot be deduced based on the colors of other vertices. Next we study how to identify the optimal vertices. We first introduce a notation for ease of presentation.

Definition 10 (Boundary Vertex) A vertex is a boundary vertex if its color cannot be deduced based on other vertices' colors. There are four cases: (1) all of its parents have different colors with the vertex; (2) all of its children have different colors with the vertex; (3) it has no child and its color is Green; or (4) it has no parent and its color is RED.

For example, $g_{6}$ is a boundary vertex as its child $g_{8}$ has different color with $g_{6} . g_{4}$ is not a boundary vertex as its child $g_{6}$ has the same color and $g_{4}$ 's color can be deduced based on $g_{6}$ 's color.

We prove that all the boundary vertices must be asked, as their colors cannot be deduced. Thus the number of asked vertices using any algorithm is not smaller than the number of boundary vertices. However, as we do not know the ground truth, we cannot identify the boundaries in advance. To address this problem, we propose effective algorithms to identify the boundary vertices with theoretical guarantees.

### 5.2 Serial Algorithm

Comparable Vertices. Given any two vertices $p_{i j}, p_{i^{\prime} j^{\prime}}$, if they are comparable, i.e., $p_{i j} \succ p_{i^{\prime} j^{\prime}}$ or $p_{i^{\prime} j^{\prime}} \succ p_{i j}$, we may deduce $p_{i j}$ 's color based on $p_{i^{\prime} j}{ }^{\prime}$ 's color, and vice versa. Obviously, two comparable vertices must be on a (directed) path in the graph, and the vertices on a path are totally ordered (i.e., any two vertices are comparable). Given a path, we can use a binary-search method to select the boundary vertices. Formally given a path, we first ask the mid-vertex on the path. (1) If the vertex is colored GREEN, its ancestors' colors can be deduced but its descendants' colors cannot be deduced, and thus we ask the mid-vertex between this vertex and the destination vertex of the path; (2) If the vertex is colored RED, its descendants' colors can be deduced but its ancestors' colors cannot be deduced, and thus we ask the mid-vertex between this vertex and the source vertex of the path. Iteratively, we can find the boundary vertices. For the path $P$ with $|P|$ vertices, the number of asked vertices is $\mathcal{O}(\log |P|)$. This is optimal and cannot be improved in general. For example, $g_{1} \rightsquigarrow g_{4} \rightsquigarrow g_{6} \rightsquigarrow g_{8} \rightsquigarrow g_{9}$ is a path. We first ask the mid-vertex $g_{6}$. As $g_{6}$ is Green, we ask the mid-vertex between $g_{6}$ and $g_{9}$, i.e., $g_{8}$. As $g_{8}$ is RED, all the vertices are colored in the path.
Incomparable Vertices. If two vertices are incomparable, we cannot deduce one vertex's color based on the other vertex's color. Suppose there are $\theta$ incomparable vertices (any two vertices are incomparable). We can divide the graph into


Fig. 5: Single-Path Method.

(a) Ask $g_{5}, g_{3}, g_{6}$

(b) Ask $g_{2}$

(c) Ask $g_{8}$

(d) All Colored

Fig. 6: Multi-Path Method.
$\theta$ disjoint paths (i.e., any two paths have no common vertices). Then we can ask each path using the binary search method. As the maximum length of a path is $|\mathcal{V}|$, the number of asked vertices is $\mathcal{O}(\theta \log |\mathcal{V}|)$. This is optimal in the worst case and cannot be improved in general. This is because if $\theta=1$, we need to ask $\log |\mathcal{V}|$ vertices. For example, in Figure 5. we have 3 disjoint paths $g_{1} \rightsquigarrow g_{4} \rightsquigarrow g_{6} \rightsquigarrow g_{8} \rightsquigarrow g_{9}$ , $g_{2} \rightsquigarrow g_{5} \rightsquigarrow g_{7}$, and $g_{3}$. We need to ask these paths using the binary-search algorithm.

Finding $\theta$ Disjoint Paths. We transform the graph $\mathcal{G}$ into a bipartite graph $\mathcal{G}^{b}=\left(\left(\mathcal{V}_{1}^{b}, \mathcal{V}_{2}^{b}\right), \mathcal{E}^{b}\right)$, where $\mathcal{V}_{1}^{b}=\mathcal{V}_{2}^{b}=\mathcal{V}$ and there is an edge between $v_{1} \in \mathcal{V}_{1}^{b}$ and $v_{2} \in \mathcal{V}_{2}^{b}$ if there is an edge $\left(v_{1}, v_{2}\right) \in \mathcal{V}$. We find a maximal matching in $\mathcal{G}^{b}=\left(\left(\mathcal{V}_{1}^{b}, \mathcal{V}_{2}^{b}\right), \mathcal{E}^{b}\right)$, which is a maximal set of edges in $\mathcal{G}^{b}=\left(\left(\mathcal{V}_{1}^{b}, \mathcal{V}_{2}^{b}\right), \mathcal{E}^{b}\right)$ where any two edges do not share a common vertex in $\mathcal{V}_{1}^{b}$ and $\mathcal{V}_{2}^{b}$, i.e., for any two edges $\left(v, v^{\prime}\right)$, $\left(u, u^{\prime}\right)$ in the matching, $v \neq u$ and $v^{\prime} \neq u^{\prime}$. Obviously any two edges in the matching sharing the same vertex in $\mathcal{V}$ must be on the same path, i.e., for any two edges $\left(v, v^{\prime}\right),\left(u, u^{\prime}\right)$ in the matching, if $v^{\prime}=u$, then $v \rightsquigarrow v^{\prime}=u \rightsquigarrow u^{\prime}$ must be on the same path based on the partial order. Note that the maximal matching can be computed in $\mathcal{O}\left(\theta|\mathcal{V}|^{2}\right)$ [10]. Based on this idea, we utilize the maximal matching to find the $\theta$ disjoint paths as follows.

Let $\mathcal{Y}$ denote the maximal matching, $\mathcal{Y}_{1}$ denote the set of the first vertices in $\mathcal{Y}$ and $\mathcal{Y}_{2}$ denote the set of the second vertices in $\mathcal{Y}$. Then $\mathcal{V}_{2}^{b}-\mathcal{Y}$ is the set of vertices that have no in-edges, and we can take them as the first vertex of a path. For each such vertex $v$, if it has an edge $\left(v, v^{\prime}\right)$, we take $v^{\prime}$ as the second vertex in the path. Then we check whether $v^{\prime}$ has an edge $\left(v^{\prime}, v^{\prime \prime}\right)$. Iteratively, we can find the path starting at $v$. The paths computed in our method satisfy: disjoint, complete and minimal, and the correctness is guaranteed by the following theorem.

Theorem 3 The set of paths found by the maximal matching of $\mathcal{G}^{b}$ satisfy:
(1) Disjoint: any two paths do not share a vertex;
(2) Complete: the paths contain all the vertices;
(3) Minimal: the size is exactly $\theta$ and is not larger than the size of any other set of paths satisfying (1) and (2).

Proof The proof essentially follows the Fulkerson's proof of Dilworth theorem [12].
(1) Disjoint: If there exist two paths with common vertices, this vertex has at least two edges in the maximal matching, which contradicts with the definition of maximal matching. (2) Complete: Consider any vertex $v$. If its in-degree is 0 , it must be covered by a path. If its in-degree is not 0 , it has an

```
Algorithm 4: Question Selection: SinglePath
    Input: \(\mathcal{G}=(\mathcal{V}, \mathcal{E})\)
    Output: All vertices in \(\mathcal{V}\) are colored as Green or RED
    while there exist uncolored vertices in \(\mathcal{V}\) do
        Compute disjoint paths using maximal matching;
        Color the longest path using binary search;
        Remove the colored vertices;
    return colored \(\mathcal{V}\)
```

in-edge $\left(v^{\prime}, v\right)$. We call $v^{\prime}$ the parent of $v$. If the in-degree of $v^{\prime}$ is $0, v^{\prime}$ and $v$ will be covered by the same path starting at $v^{\prime}$; otherwise we find the parent of $v^{\prime}$. Iteratively we find an ancestor of $v$ whose in-degree is 0 , and then $v$ is covered by the path starting at this ancestor.
(3) Minimal: Let $J$ denote the number of edges in a matching and $D$ denote the number of disjoint paths in the graph. Fulkerson et al. [12] proved that $J+D=|\mathcal{V}|$. As $|\mathcal{V}|$ is fixed, if we find the maximal matching, then $D$ is minimal.

For example, consider the graph in Figure 3. We construct a bipartite graph as shown in Figure 5(a). As there is an edge from $g_{1}$ to $g_{3}$ in $\mathcal{G}$, there is an edge from $g_{1}$ in $\mathcal{V}_{1}^{b}$ to $g_{3}$ in $\mathcal{V}_{2}^{b}$. Thus $\mathcal{G}$ and $\mathcal{G}^{b}$ have the same number of edges. Then we find a maximal matching which is the set of the colored edges. The vertices $g_{1}, g_{2}$ and $g_{3}$ in $\mathcal{V}_{2}^{b}$ have no in-edges in the maximal matching. We compute the disjoint paths starting from them. From $g_{1}$ we get path $g_{1} \rightsquigarrow g_{4} \rightsquigarrow g_{6} \rightsquigarrow g_{8} \rightsquigarrow g_{9}$; from $g_{2}$ we get $g_{2} \rightsquigarrow g_{5} \rightsquigarrow g_{7}$; and $g_{3}$ itself is a path. Thus we get 3 disjoint paths.
SinglePath Algorithm. Then we propose a path-based question selection algorithm. The pseudo code is shown in Algorithm 4 It first computes the $\theta$ disjoint paths. Then it asks the longest path using the binary-search method, colors the graphs, and then removes the colored vertices. Next it recomputes the disjoint paths and asks the next longest path. Iteratively it can color all vertices. The complexity of this algorithm is $\mathcal{O}\left(\theta|\mathcal{V}|^{2}\right)$.

For example, in Figure 55 we first identify the minimal disjoint paths as shown in Figure 5(a). Then we select the longest path (Figure 5(b)), ask the path using binary search. We first ask $g_{6}$ and color the graph based on the answers of asked vertices (Figure 5(c)). Next we ask $g_{8}$ and get Figure 5(d) Then we recompute the disjoint paths, ask midvertex of the longest path $g_{2} \rightsquigarrow g_{5} \rightsquigarrow g_{7}$ (Figure 5(d)p, and color the graph (Figure 5(e). Next as there is only one vertex left, we ask it and get the final result (Figure 5(f). This method totally asks 4 vertices and involves 4 iterations.

### 5.3 Parallel Algorithm

If users do not care about the latency, the single-path algorithm is a good choice. However if the latency is very crucial, the single-path algorithm is not acceptable as it needs to post one question at a time on crowdsourcing platforms,
which would result in a long time latency. To address this issue, we design parallel algorithms, which select multiple vertices and ask them together in each iteration.

### 5.3.1 Multi-Path Algorithm

We extend the path-based algorithm to support the parallel setting. We first identify the $\theta$ disjoint paths and then ask their mid-vertices in parallel. Based on the answers on these vertices, we color the graph. Next we remove the colored vertices and repeat the above step until all the vertices are colored. Figure 6 shows an example. Note that the parallel algorithm may generate conflicts. For example, if $g_{i}$ is colored GREEN and $g_{j}$ is colored RED, then there is a conflict on $g$ where $g \succ g_{i}$ and $g_{j} \succ g$, because $g$ is deduced as Green based on $g_{i}$ and deduced as Red based on $g_{j}$. To address this issue, we can use majority voting to vote $g$ 's color. Algorithm 5 shows the pseudo code. It first finds the minimal disjoint paths (line 2) and then asks their mid-vertices in parallel (lines 4.5. Next it colors the graph based on the answers and removes the colored vertices (line 6). Finally, it repeats the above step if there exist uncolored vertices in $\mathcal{V}$. For example, we first compute the three disjoint paths and asks their mid-vertices $g_{5}, g_{3}$ and $g_{6}$ together in Figure 6 We get the answers: $g_{5}$ is RED, and $g_{3}$ and $g_{6}$ are Green. We color the graph based on these three answers(Figure 6(b). Next we generate a path: $g_{2} \rightsquigarrow g_{8}$ and we ask $g_{2}$. The answer is: $g_{2}$ is Green, and we color the graph (Figure 6(c). Iteratively we color all vertices (Figure 6(d)). This method asks 5 vertices and involves 3 iterations.

### 5.3.2 Topological-Sorting-Based Algorithm

In the multi-path algorithm, the asked vertices may have ancestor-descendent relationships, and thus it may ask unnecessary questions. For example, in Figure 6(a) we do not need to ask $g_{3}$ and $g_{6}$ together, as the color of $g_{3}$ can be deduced based on the color of $g_{6}$. To address this issue, we aim to ask independent vertices in each iteration.

To this end, we perform a topological sorting on the vertices. We first identify the set of vertices with zero in-degree, denoted by $\mathcal{L}_{1}$. Then we delete them from the graph and find another set of vertices whose in-degrees are zero, denoted by $\mathcal{L}_{2}$. We repeat this step until all vertices are deleted. Suppose there are $|\mathcal{L}|$ sets, $\mathcal{L}_{1}, \mathcal{L}_{2}, \cdots, \mathcal{L}_{|\mathcal{L}|}$. Obviously vertices in each $\mathcal{L}_{i}$ have no in-edges (as their in-degrees are 0 ) and thus can be taken as an independent set. Moreover, the vertices in the sets with small subscripts (e.g., $\mathcal{L}_{1}, \mathcal{L}_{2}$ ) are more likely to be colored GREEN and the vertices in the sets with large subscripts (e.g., $\mathcal{L}_{|\mathcal{L}|}$ ) are more likely to be colored RED, and thus we cannot deduce the colors of many uncolored vertices based on them. In other words, the boundary


Fig. 7: Topological-Sorting Based Method.

```
Algorithm 5: Question Selection: Multi-Path
    Input: G}=(\mathcal{V},\mathcal{E}
    Output: All vertices in \mathcal{V}}\mathrm{ are colored as Green or RED
    while there exist uncolored vertices in \mathcal{V}}\mathrm{ do
        Compute 0 disjoint paths;
        for each path of these disjoint paths do
            N}\leftarrow\mathrm{ mid-vertex of the path;
        Ask \mathcal{N}}\mathrm{ to workers in parallel and color }\mathcal{G}\mathrm{ ;
        Removed colored vertices;
    return colored V
```

```
Algorithm 6: Topological Sorting
    Input: \(\mathcal{G}=(\mathcal{V}, \mathcal{E})\)
    Output: All vertices in \(\mathcal{V}\) are colored as Green or Red
    while there exist uncolored vertices in \(\mathcal{V}\) do
        Do a topological sorting on the uncolored vertices in \(\mathcal{G}\) and
        obtain \(|\mathcal{L}|\) sets, \(\mathcal{L}_{1}, \mathcal{L}_{2}, \cdots, \mathcal{L}_{|\mathcal{L}|}\);
        Ask workers to color vertices in \(\mathcal{L}_{\frac{|\mathcal{L}|+1}{2}}\);
    return colored \(\mathcal{V}\);
```

vertices are more likely to be in the middle sets. To this end, we first ask vertices in $\mathcal{L}_{\frac{|\mathcal{L}|+1}{2}}$.

Next we design a topological-sorting-based algorithm and Algorithm 6 illustrates the pseudo code. It first computes topological-sorted sets $\mathcal{L}_{1}, \mathcal{L}_{2}, \cdots, \mathcal{L}_{|\mathcal{L}|}$. Then it asks vertices in $\mathcal{L}_{\frac{|\mathcal{L}|+1}{2}}$ in parallel. Based on the results of these vertices, it colors the graph, removes the colored vertices, and repeats the above step. Iteratively it colors all vertices.

For example, we construct the topological structure as shown in Figure 7(a) $\mathcal{L}_{1}=\left\{g_{1}\right\}, \mathcal{L}_{2}=\left\{g_{2}, g_{3}, g_{4}\right\}, \mathcal{L}_{3}=$ $\left\{g_{5}, g_{6}\right\}, \mathcal{L}_{4}=\left\{g_{7}, g_{8}\right\}, \mathcal{L}_{5}=\left\{g_{9}\right\}$ and $|\mathcal{L}|=5$. So we select $\mathcal{L}_{3}=\left\{g_{5}, g_{6}\right\}$ and ask the vertices. After getting their answers, we obtain Figure 7(b) Then we compute the topological sorting on the graph of the uncolored vertices. Next, $\mathcal{L}_{1}=\left\{g_{2}\right\}, \mathcal{L}_{2}=\left\{g_{8}\right\}$. We ask $g_{2}$. After this iteration, only $g_{8}$ is uncolored. We ask it and get the final result (Figure 7(d). This method asks 4 vertices and has 3 iterations.

## 6 Tolerating Errors

There are two types of possible errors in our framework. The first is caused by workers' errors and the second is introduced by our coloring strategy. For example, suppose a vertex $p_{i j}$ is actually RED. However the workers wrongly color it Green. This error is caused by workers. Consider $p_{i j}$ 's ancestor, $p_{i^{\prime} j^{\prime}}$, whose color is Red. Our coloring strategy will wrongly color it Green based on partial order. This error is caused by our coloring strategy. Next we discuss how to address these errors.
Confidence of Workers' Answers. To tolerate workers' errors, we assign each vertex to multiple workers and aggregate their answers. There are many methods to compute the confidence of workers' answers, and we take majority voting as an example and any other techniques can be integrated into our method. Suppose each vertex is assigned to $z$ workers and $y>\frac{z}{2}$ workers vote a consensus answer (e.g, Yes) and $z-y$ workers vote the other answer (e.g., No). The confidence of the voted answer is $c=\frac{y}{z}$.
Error-Tolerant Coloring Strategy. For each crowdsourced vertex, if the confidence of workers on this vertex is larger than a confidence threshold, e.g, $\geq 0.8$, we use our coloring strategy to color its ancestors or descendants; otherwise, we color it BLUE and do not color its ancestors or descendants. For the Green and Red pairs, we take them as ground truth as their answers have large confidences. Next we utilize them to color BLUE pairs.

We first need to compute the weights of different attributes which reflect the importance in determining the colors of each pair. Let $P^{g}$ denote the set of Green pairs. For every $p_{i j} \in P^{g}$, if $s_{i j}^{k}$ is large, then attribute $\mathcal{A}_{k}$ plays an important role to determine the color of $p_{i j}$, and we should assign it with a large weight; otherwise it is insignificant to determine the color of $p_{i j}$. To this end, we assign a weight $\omega_{k}$ for each attribute $\mathcal{A}_{k}$ as below
$\omega_{k}=\frac{\sum_{p_{i j} \in P^{g}} s_{i j}^{k}}{\sum_{p_{i j} \in P^{g}} \sum_{1 \leq t \leq m} s_{i j}^{t}}$.

```
Algorithm 7: Error-Tolerant
    Input: \(\mathcal{G}=(\mathcal{V}, \mathcal{E})\)
    Output: All vertices in \(\mathcal{V}\) are colored as Green or RED
    while there exist uncolored vertices in \(\mathcal{V}\) do
        Select a set of uncolored vertices to ask workers;
        for each asked \(p_{i j}\) with an answer do
            if confidence \(\geq 0.8\) then
                    color \(p_{i j}\) and its ancestors or descendents;
            else color \(p_{i j}\) BLUE;
    Generate histogram \(h_{i}\) and compute \(\mathrm{Pr}_{i}\);
    for each \(p_{i^{\prime} j^{\prime}}\) colored BLUE in \(h_{i}\) do
        if \(P r_{i}>0.5\) then color \(p_{i^{\prime} j^{\prime}}\) GREEN;
        else color \(p_{i^{\prime} j^{\prime}}\) RED;
    return colored \(\mathcal{V}\),
```

Then we compute a weighted similarity of $p_{i j}$,
$\hat{s}_{i j}=\sum_{k \in[1, m]} \omega_{k} \cdot s_{i j}^{k}$.
Coloring The Pairs in Low-Confidence Groups. We use a histogram based method to color pairs in BLUE vertices [44, 47]. We first generate equi-width histograms based on the weighted similarities of pairs in GREEN and RED vertices. Each histogram $h_{i}$ contains a set of pairs within a similarity range. We count the number of GREEN pairs in $h_{i}$ and compute the probability $\mathrm{Pr}_{i}$ that pairs in $h_{i}$ should be colored Green, i.e., the number of Green pairs to the total number of pairs in $h_{i}$. Then we assign the pairs in Blue vertices into the histograms and color them based on probability $\mathrm{Pr}_{i}$. For example, if a pair falls in a histogram with high probability of Green, the vertex is colored Green; otherwise RED. Algorithm 7 shows the pseudo code. It uses the coloring strategy only for the vertices with high-confidence answers (line 5) and utilizes the histograms to color the vertices with low-confidence answers (lines $7 \mid 10$.

Recall the topological-sorting method in Figure 7(b). The workers return the answer of $g_{2}$ with a low confidence, and we color it BLUE and do another topological sorting among the rest groups, i.e., $g_{8} . g_{8}$ is colored BLUE as workers give a low confidence answer. We get Figure 8 . Then we need to color pairs in $g_{2}$ and $g_{8}$ (i.e., $p_{12}, p_{24}, p_{25}$ ) based on the colored pairs. First, we calculate the attribute weight $\omega$ based on the pairs $P^{g}=\left\{p_{12}, p_{67}, p_{45}, p_{23}, p_{46}, p_{56}, p_{47}, p_{57}\right\}$ in the colored groups. Using Equation 7, we obtain $\omega=\{0.32$, $0.28,0.21,0.19\}$. Then we build 5 histograms with width 0.2 . We compute $\hat{s}_{i j}$ of each colored pair by Equation 8 and assign it into the corresponding histogram. Figure 10 shows the histograms and table 9 shows the estimated similarities. $\left\{p_{67}, p_{45}\right\}$ are assigned into $h_{5}$ ( $[0.8,1]$ ). As all of them are colored GREEN, $\mathrm{Pr}_{5}=1 .\left\{p_{23}, p_{13}\right\}$ are assigned into $h_{4}$ ( $\left[0.6,0.8\right.$ )), and $\mathrm{Pr}_{4}=1 .\left\{p_{46}, p_{57}, p_{47}, p_{56}, p_{10,11}, p_{26}\right.$, $\left.p_{27}\right\}$ are assigned into $h_{3}([0.4,0.6))$, and $\mathrm{Pr}_{3}=\frac{4}{7}=0.57$. $\left\{p_{37}, p_{89}, p_{34}, p_{35}\right\}$ are assigned into $h_{2}$ ([0.2,0.4)), and $\mathrm{Pr}_{2}=0$. Next we compute $\hat{s}_{i j}$ of $p_{12}, p_{24}$ and $p_{25}$. For

| $p_{i j}$ | $\hat{s}_{i j}$ | $p_{i j}$ | $\hat{s}_{i j}$ |
| :---: | :---: | :---: | :---: |
| $p_{12}$ | 0.72 | $p_{37}$ | 0.21 |
| $p_{13}$ | 0.68 | $p_{45}$ | 0.97 |
| $p_{23}$ | 0.60 | $p_{46}$ | 0.43 |
| $p_{24}$ | 0.28 | $p_{47}$ | 0.42 |
| $p_{25}$ | 0.29 | $p_{56}$ | 0.41 |
| $p_{26}$ | 0.40 | $p_{57}$ | 0.44 |
| $p_{27}$ | 0.41 | $p_{67}$ | 0.98 |
| $p_{34}$ | 0.39 | $p_{89}$ | 0.37 |
| $p_{35}$ | 0.39 | $p_{10,11}$ | 0.44 |

Fig. 9: Estimated Similarity $\hat{s}_{i j}$.


Fig. 10: Equi-width
Histograms.
instance, $\hat{s}_{12}=0.32 \times 0.72+0.28 \times 0.4+0.21 \times 1+$ $0.19 \times 0.88=0.72$, so we assign it into $h_{4}$ and color it Green as $\mathrm{Pr}_{4}>0.5$. Similarly, we color $p_{24}$ and $p_{25}$ Red.
Remark For the confidence threshold, if we select a too low confidence, we are likely to prune less pairs and increase the cost. If we select a too high confidence, we may have a low quality. Obviously, the most suitable confidence depends on the difficulty of the datasets. In our framework, we first ask the pairs in middle layer, which are likely to be the most difficult pairs, and thus reflect the difficulty of the dataset. Then in the first iteration, we can collect the answers per pair and then compute the numbers of pairs corresponding to different confidence values. Then we can select the confidence value with the most number of pairs. For example, if there are 50 pairs with confidence 0.8 and 20 with confidence 0.6 which means that most workers have confidence 0.8 . Thus we choose 0.8 as the final confidence threshold.

## 7 Budget-Aware Methods

In this section, we first formulate the budget-based entity resolution problem (Section 7.1), and then propose a serial algorithm and a parallel algorithm to solve the problem (Section 7.2).

### 7.1 Budget-based Question Selection

Crowdsourced entity resolution aims to minimize the number of questions to color the whole graph. Thus it regards the Red vertices with the same importance as the Green ones. The budget-aware problem, however, aims to identify as many Green vertices (i.e., the matching pairs that refer to the same entity) as possible within the budget, which is formalized as below.

## Definition 11 (Optimal Budget-based Question Selection)

Given a budget $B$ and a graph $\mathcal{G}=(\mathcal{V}, \mathcal{E})$, color the maximum number of Green vertices by asking at most $B$ questions.

Intuitively, it should first ask the GREEN boundary vertices. For example, in Figure 1, if $B$ is equal to or more than

2 , the vertices $p_{12}$ and $p_{56}$ are the optimal choices. If $B$ is less than 2 , i.e., only one question can be asked, then $p_{56}$ is the optimal vertex because it can deduce more vertices than $p_{12}$.

We can prove that if the budget $B$ exceeds the number of Green boundary vertices, the optimal solution is to ask all Green boundary vertices. Otherwise, the budget-aware problem is NP-hard as proved in Theorem4.

Theorem 4 If budget $B$ exceeds the number of Green boundary vertices, the optimal solution asks the Green boundary vertices. Otherwise, the budget-aware problem is NP-hard.

Proof We first consider the case that the budget exceeds the number of Green boundary vertices. Apparently, the optimal solution is that we must ask all Green boundary vertices because any matching pair can be deduced by a Green boundary vertex.

We then consider the case that the budget is smaller than the number of GREEN boundary vertices. We prove that the problem is NP-hard, by a reduction from the set cover problem. Set cover problem is defined as follows: given a collection $C$ of subsets of a finite set $S$, a positive integer $K \leq|C|$, it asks whether there exists a subset $C^{\prime} \subseteq C$ with $\left|C^{\prime}\right| \leq K$ such that every element of $S$ belongs to at least one member of $C^{\prime}$ ? To reduce this problem to the budget-based selection problem, we take the finite set $S$ as the set of Green vertices and take each element of collection $C$ as each Green boundary node and its ancestors. Therefore, the set cover problem can be reduced to budget-based optimal selection problem.

For the first case, we aim to ask all the Green boundary vertices. For the second case, we aim to first ask the Green boundary vertices which can be used to deduce a large number of vertices, i.e., those having many ancestors. To this end, we propose a unified framework to address the two cases together.

### 7.2 Budget-Aware Algorithm

### 7.2.1 Budget-Aware Serial Algorithm

We need to consider two factors to select a vertex to ask. Firstly, the vertex $v$ should have a large probability $\operatorname{Pr}(v)$ to be a Green boundary vertex. Secondly, the vertex $v$ should have a large number $|\mathrm{F}(v)|$ of ancestors. Thus we combine them together and propose a benefit function $\operatorname{Exp}(v)$, which denotes the benefit to select a question $p_{i j} \in v$. Obviously we want to select the question with the largest benefit.

Next we define the benefit function as below.
$\operatorname{Exp}(v)=\operatorname{Pr}(v) \times|F(v)|$

Obviously $|F(v)|$ is easy to compute based on the graph structure and next we discuss how to compute $\operatorname{Pr}(v)$. Similar to Section 6, we still compute the weighted similarity of each pair based on Equation 8 and build histograms. Then for each pair, if the pair falls in a histogram with high probability of Green, the pair has large probability to be Green and the probability that the pair is colored Green is set as the probability of this histogram. Then for each vertex, we can compute the probability that the vertex is colored as Green by computing the overall similarity of all pairs in the vertex. Formally, we can compute $\operatorname{Pr}(v)$ as below.
$\operatorname{Pr}(v)=\frac{\sum_{p_{i j} \in v} \operatorname{Pr}_{p_{i j}}}{|v|}$
where $P r_{p_{i j}}$ is the probability that $p_{i j}$ is colored as Green based on the histogram and $|v|$ is the number of pairs in vertex $v$.

Initialization. At the beginning of the algorithm, we have no hints on which vertices are Green and we do not have histograms either. Thus we have to select one vertex to ask without any information. Intuitively, we do not want to select the pair with too large or too small similarities. Instead we want to select the boundary vertices close to the average similarity. To this end, we propose a boundary estimation function to compute the closeness of a pair to the average similarity of all pairs. For each pair $p_{i j}$, we compute the $m$ differences between $s_{i j}^{k}$ and average scores on attribute $\mathcal{A}_{k}$. Then we sum the $m$ differences as the final boundary estimation of this pair, which is defined as below.
$\min _{p_{i j} \in P^{u}} \sum_{k=1}^{m}\left|s_{i j}^{k}-\frac{\sum_{p_{i j} \in P^{u}} s_{i j}^{k}}{\left|P^{u}\right|}\right|$.
where $\left|P^{u}\right|$ is the number of uncolored pairs.
Thus initially we select the pair with the minimal boundary estimation value based on Equation 11 .

The pseudo code is shown in algorithm 8 It first selects the first pair with the minimal boundary estimation value based on Equation 11 . Then it selects pairs with the largest benefits based on Equation 10 . The algorithm terminates until the budget is used up.

For example, given the uncolored graph in Figure 1 and $B=2$, we illustrate how our algorithm works as follows. Since each vertex contains only one pair in this example, we take one pair as a group. Firstly, we compute the average similarity scores $\left(\frac{\sum_{p_{i j} \in P^{u}} s_{i j}^{k}}{\left|P^{u}\right|}\right)$ among all uncolored pairs and we get the average scores $(0.58,0.41,0.66,0.24)$. Then we traverse all the pairs and compute the boundary estimation of all uncolored pairs. The minimum boundary estimation is $p_{56}(|0.63-0.58|+|0.5-041|+|0.33-0.66|+\mid 0-$ $0.24 \mid=0.71)$ and we select $p_{56}$ to ask first. Then $p_{56}, p_{47}$, $p_{46}, p_{57}, p_{13}, p_{23}, p_{45}, p_{67}$ are colored green. Using equation

```
Algorithm 8: Budget-Aware Algorithm:Serial
    Input: \(\mathcal{G}=(\mathcal{V}, \mathcal{E}), B\)
    Output: All Green vertices in \(\mathcal{V}\) after asking \(B\) questions
    Select a pair \(p\) with the minimum offset by equation 11
    Ask workers to color \(p\).
    while \(B>0\) do
        Compute \(\hat{s}\) of all groups, reconstruct the histograms and
        compute the probability corresponding to each histogram.
        Select the group \(g_{i}\) with the most benefit.
        \(B=B-1\)
        Select a pair \(p\) randomly from \(g_{i}\) and ask workers to color
        p.
    return \(\mathcal{V}\);
```

```
Algorithm 9: Budget-Aware Algorithm:Parallel
    Input: \(\mathcal{G}=(\mathcal{V}, \mathcal{E}), B\)
    Output: All Green vertices in \(\mathcal{V}\) after \(B\) uses up
    Select the maximum independent groups by equation 11
    Ask workers to color these pairs.
    while \(B>0\) do
        Compute \(\hat{s}\) of all groups, reconstruct the histograms and
        compute the probability corresponding to each histogram.
        Select the maximum independent groups according to the
        descending order of all uncolored pairs.
        Update the \(B\).
        Ask workers to color pairs randomly selected from these
        groups.
    return \(\mathcal{V}\)
```

7. we obtain $\omega=\{0.33,0.28,0.21,0.18\}$. We next compute $\hat{s}$ of each pair and get that $\hat{s}_{12}=0.72, \hat{s}_{10,11}=0.44$, $\hat{s}_{27}=0.41$. Then we obtain $\operatorname{Exp}\left(p_{12}\right)=0.72 \times 3=2.16$, $\operatorname{Exp}\left(p_{10,11}\right)=0.44 \times 4=1.76, \operatorname{Exp}\left(p_{27}\right)=0.41 \times 5=$ 2.05. So $p_{12}$ is the pair with the most benefit and we select it to ask. Then we can get all RED vertices within the budget.

### 7.2.2 Budget-Aware Parallel Algorithm

However, if we care about the latency, only asking one question in one crowdsourcing iteration is unacceptable. Therefore, we propose a parallel algorithm. Similar to the parallel algorithm proposed in section 5.3 above, we aim to select vertices with no ancestor-descendent relationships in each iteration. Furthermore, we should also select vertices with large benefit. Then we discuss how to select them.

Firstly, at the beginning of the algorithm, we need to select groups without any information in the first iteration. Note that the asked groups in each iteration should not have ancestor-descendent relationships. Therefore, we sort all uncolored groups according to their differences computed by Equation 11 in an ascending order. We select the first group as one of questions greedily. Then we check whether the next group has ancestor-descendent relationship with the first group. If not, we add the pair into the question set. Otherwise we check the next group. We repeat this until there is no vertex can be added to the question set. Then we ask the set of groups in parallel. After workers give answers, we color the

Table 4: Three real-world Datasets

|  | \#Records | \#Attr | \#Pairs | \#Workers/Pair |
| :---: | :---: | :---: | :---: | :---: |
| Restaurant | 858 | 4 | 5010 | 5 |
| Cora | 997 | 8 | 29510 | 5 |
| ACMPub | 66,879 | 4 | 204,000 | 5 |

graph and select the next batch of groups. Then we sort all uncolored groups according to their benefits in a descending order and select independent groups similarly to the method above. Next, we color the graph and select another set of independent groups according to their benefits until the budget is used up. The pseudo code is shown in algorithm 9

For example, assuming that $B=3$, at the beginning, we sort these pairs and get the order $p_{56}, p_{47}, p_{46}, p_{57}$, $p_{10,11}$. Though $p_{47}$ is the second pair in the list, we select $p_{56}$ and $p_{10,11}$ because the next three pairs after $p_{56}$ have ancestor-descendent relationships with it. After the first iteration, there are only $p_{12}, p_{25}$ and $p_{24}$ left and $p_{12}$ is the pair with the most benefit. Since all pairs left have ancestordescendent relationships with $p_{12}$, we can only select $p_{12}$ in this iteration. Then we can see that all green groups have been found out and we use only 2 iterations.

## 8 Experiment

In this section, we evaluate our methods and report experimental results. The goals of our experiments include (1) evaluating our proposed techniques and (2) comparing our method with state-of-the-art approaches.

### 8.1 Experimental Setting

Datasets. We use three real-world datasets which are widely adopted by existing works [13, 44, 45, 46]. (1) Restaurant ${ }^{2}$ is a restaurant dataset consisting of 858 restaurants with 752 different entities. The dataset has four attributes, Name, Address, City and Flavor. (2) Cora ${ }^{3}$ is a dataset of research papers, which contains 997 records with 191 different entities. The dataset has 8 attributes: Author, Title, Journal, Year, Pages, Publisher, Type and Editor. (3) ACMP ub ${ }^{4}$ is a larger publication dataset consisting of 66,879 records with 5347 different entities. It has four attributes: Author, Title, Conference and Year. Table 4 shows the details.
Similarity Functions. We use three similarity functions, Jaccard, edit similarity and bigram Jaccard. For bigram, we generate bigrams for every attribute and compute Jaccard on bigram sets as the similarity, where a bigram is a substring with length 2 and a bigram set contains all the bigrams in an attribute. We use bigram by default.

[^2]

Fig. 11: Graph Construction: Efficiency.
Pruning. As ACMPub has 66,879 records, it will generate $\frac{66879 * 66878}{2}=2,236,366,881$ pairs and it is expensive to consider all of them. Following previous work [46,45], we compute a similarity score for each pair of records by Jaccard and prune pairs whose similarity scores are bellow 0.3. After pruning, there are 5010, 29510 and 204000 pairs left in Restaurant, Cora and ACMP ub datasets respectively.
AMT Setting. We use Amazon Mechanical Turk (AMT). To ensure fair comparison between different algorithms, each question should be answered by the same workers. To this end, we crowdsource all pairs in each dataset to AMT and get their answers. If different algorithms ask the same pair, they will use the same answer. We assign each question to five workers and use the weighted majority voting to integrate the answers. We pack every ten pairs in a HIT and pay 10 cents for each HIT. We vary workers' accuracy which can be specified on AMT, where the worker accuracy is computed based on workers' approval rate in history at AMT.
Comparison. We compare with state-of-the-art methods, including Trans [45], ACD [46] and GCER [47] on the same experimental setting. We get the source codes of ACD and Trans from the authors and implement GCER by ourselves.
Evaluation Metrics. We compare the quality, the number of questions, the number of iterations, and the assignment time. For quality, we use F-measure, which is a combination of precision and recall. Suppose the set of pairs that refer to the same entity is $S_{T}$, and the set of pairs that an algorithm reports as the same entity is $S_{P}$. Then the precision is $p=$ $\frac{\left|S_{T} \cap S_{P}\right|}{\left|S_{P}\right|}$, recall is $r=\frac{\left|S_{T} \cap S_{P}\right|}{\left|S_{T}\right|}$, and F-measure is $\frac{2 p r}{p+r}$.

### 8.2 Evaluation on Graph Construction

### 8.2.1 Evaluation on Graph Construction Algorithms

We compare the efficiency of the three graph construction algorithms (proposed in Section 4.1). (1) Brute-Force: the brute-force method that compares every two pairs. (2) QuickSort: the quicksort-based method. (3) Index: the index-based method ${ }^{5}$ To test the scalability, on ACMPub, we set the bound $\tau$ as 0.18 and generate 500K pairs. Figure 11 shows the results by varying the number of pairs. We

[^3]

Fig. 13: Grouping: Efficiency.
can see that Index significantly outperforms the other two methods, even by 1 order of magnitude. For example, on the Cora dataset with 28 k pairs, Brute-Force takes 20 seconds, QuickSort improves it to 10 seconds, while Index only takes 1 second. On the larger dataset ACMPub, Index still outperforms other methods and achieves higher performance. This is because Index can utilize the range search tree index to efficiently find the children of a pair and can prune many unnecessary pairs (e.g., incomparable pairs). QuickSort outperforms Brute-Force because it can also remove some unnecessary pairs based on the partial order. However the improvement is not signifiant, as many vertices in the graph are not comparable based on the partial order and thus many pairs cannot be pruned. For example, in Restaurant, $70 \%$ pairs of records are not comparable. In Cora, $84 \%$ pairs of records are not comparable. In ACMP ub, $80 \%$ pairs of records are not comparable.

### 8.2.2 Evaluation on Grouping

We first evaluate our two techniques Greedy and Split (proposed in Section 4.2). (1) Greedy: it greedily groups the vertices. (2) Split: it uses the split-based technique. We first compare the number of groups generated by them. Figure 12 shows the number of groups and Figure 13 shows the running time. Note that on the ACMP ub dataset, Greedy cannot report the results within 10 hours and thus we do not show Greedy in the figure. We have several observations on the number of groups. Firstly, compared with the total number of pairs in Restaurant ( 5,010 pairs), Cora ( 29,510 pairs) and ACMPub (204,000 pairs), Split and Greedy only generate less than 150,1300 and 700 groups. Thus the grouping technique can significantly reduce the number of vertices, and thus can reduce the time latency and the crowd cost. Secondly, Split generates a few more groups than Greedy, because Split uses heuristics to generate groups while Greedy utilizes a greedy strategy to generate high-quality groups. For example, on the Cora dataset with grouping threshold $\varepsilon=0.1$, Greedy generates 800 groups and Split generates 1200 groups. Thus





Fi) Restaurant 16 : Question Selection(Serial): Quality.


Fig. 17: Question Selection(Serial): \#Questions.
close. This verifies that we can use Split to generate the groups. (4) The number of questions is not only determined by the number of groups, but also the number of edges. First, fewer groups will lead to fewer questions. Second, fewer edges may lead to more questions, because the answers of many groups cannot be deduced based on the answers of other groups. With the increase of the grouping threshold, the number of groups decreases, and thus the number of questions should decrease intuitively. However, with the increase of the grouping threshold, the groups become larger and it is more difficult to add an edge between two groups. Thus there may be fewer edges in the graph and the number of questions may decrease.

### 8.3 Evaluation on Question Selection

### 8.3.1 Evaluation on Serial Algorithms

We first evaluate the serial question-selection algorithms and compare two algorithms (Section 5.2). (1) Random: which randomly selects a vertex in each iteration. (2) SingleP ath: which selects a vertex from the longest path in each iteration. We compare the two algorithms on the non-grouping graphs. Figure 16 shows the quality and Figure 17 shows the number of questions. We can see that SinglePath outperforms Random and reduces the number of questions. For example, on the Restaurant dataset with 5000 pairs, Random asks 250 pairs while SinglePath only asks 150 pairs. On the ACMPub dataset, Random asks 2500 pairs while SinglePath only asks 1400 pairs. This is because SinglePath can effectively identify the boundary pairs using a binary search strategy. On the other hand, SinglePath achieves similar quality with Random as the question order does not significantly affect the quality. Thus we can utilize the SinglePath to select questions.


Fig. 20: Question Selection(Parallel):\#Iterations.


Fig. 21: Question Selection(Parallel): Time.

### 8.3.2 Evaluation on Parallel Algorithms

We then evaluate the parallel question-selection algorithms (proposed in Section5.3). We compare three algorithms: (1) SinglePath: which selects a vertex from the longest path in each iteration. (2) Multi-Path: which selects multiple vertices from multiple disjoint paths in each iteration. (3) Power: which selects multiple independent vertices based on topological sorting in each iteration. We compare the quality, the number of questions, the number of iterations, and the assignment time in each iteration to select the questions for workers. Figures 18,21 show the results.
(1) For quality, we can see that the three methods achieve similar quality, because different question orders will not affect the quality based on the partial order. (2) For the number of questions, we can see that the two parallel algorithms Multi-Path and Power ask a few more questions than SinglePath. The reason is evident that Multi-Path may ask pairs with ancestor-descendent relationships and Power may ask pairs with the same descendants which can be avoided by serial algorithms based on the partial order. Power outperforms Multi-Path because Power asks independent questions in each iteration while Multi-Path may ask dependent questions. (3) For the number of iterations, the two parallel algorithms Multi-Path and Power significantly outperform SinglePath as they ask ques-

Fig. 24: Error Tolerant: \#Iterations.
tions in parallel. For example, on the Cora dataset, Power and Multi-Path only have 4 iterations while SinglePath involves 200 iterations. On the ACMP ub dataset, Power and Multi-Path have 5 iterations while SinglePath involves 113 iterations. Thus Power and Multi-Path can significantly reduce the latency. In practice, we need to use the parallel algorithms. (4) For assignment time, all the three algorithms can assign tasks within 1 second. Multi-Path and SinglePath take longer time than Power as they are expensive to find multiple independent paths using the graph matching algorithm, which is consistent with the complexity analysis, while Power only needs to compute the topological sorting which is efficient.

### 8.4 Evaluation on Error-Tolerant Techniques

We evaluate the error-tolerant techniques (proposed in Section 6 and compare two algorithms. (1) Power: which does not consider errors. (2) Power+: which extends Power to tolerate errors. We compare quality, the number of questions, and the number of iterations. As they have the same assignment time, we do not compare the assignment time. We build 20 histograms. Figures 22,24 show the results.

Power+ achieves better quality than Power, especially on the Cora dataset, because it can tolerate the errors introduced by workers and the partial order. For example, on the Cora dataset with $\varepsilon=0.1$, Power only has $79 \%$ Fmeasure while Power + improves the quality to $83 \%$. On the ACMP ub dataset, P ower has $87 \%$ F-measure while Power + improves to $90 \%$. On the Restaurant dataset, the improvement is not signifiant because the dataset is easy and Power already achieves $96 \%$ F-measure. On the other hand,


Fig. 25:Evaluation by Varying \#Attributes(Cora).
Power+ asks a little more questions than Power as Power+ does not utilize the partial order for some pairs and thus reduces the number of deduced pairs. The two methods have the same number of iterations, because the only difference is that Power+ does not deduce the answers for some unconfident pairs. Thus we can use the error-tolerant technique to improve the quality.

### 8.5 Evaluation on The Number of Attributes

We vary the number of attributes on the Cora dataset and Figure 25 shows the results. We can see that with the increase of attribute numbers, the number of questions increases, because it is harder to add edges between pairs for more attributes and thus the number of edges decreases. Similar to the number of questions, the number of iterations slightly increases. The quality is not affected as it is determined by the partial order and the crowd error.

### 8.6 Evaluating Worker Accuracy

We compare our methods (Power without error-tolerant techniques and Power+ with error-tolerant techniques) with state-of-the-art approaches ACD [45], Trans [46] and GCER [47]. We compare the number of iterations, the number of questions, and the quality. As GCER requires a parameter to tune the number of asked pairs, we set this parameter the same as $A C D$, i.e., the maximum number of questions among these algorithms. GCER asks 100 questions in each iteration. For our algorithms, we use the split-based grouping algorithm to group the pairs and set the grouping threshold $\varepsilon$ as 0.1 , utilize the index-based method to construct the graph, and employ the topological-sorting algorithm to select questions.

### 8.6.1 Real Exp: Evaluating Worker Accuracy

Existing studies [46 44 13] select the workers with approval rate above $95 \%$ or passing a qualification test to avoid malicious workers. To evaluate the robustness of the algorithms, we vary the workers' accuracy. In the real crowdsourcing platforms AMT, we can specify the worker accuracy by selecting the approval rate. We select three groups of workers, $70 \%-80 \%$ ( $70 \%$ in the figure), $80 \%-90 \%$ ( $80 \%$ in the figure) and above $90 \%$ ( $90 \%$ in the figure) respectively. For each group of workers, we ask them to answer our questions and compare different algorithms. Figures 26,28 show the results. We make the following observations.

Quality. Firstly, Power+ outperforms Power. This because Power+ can tolerate workers' errors. With the increase of worker accuracy, the improvement decreases. This is because for higher worker accuracy, there are fewer errors and Power+ has limited room to further improve the quality. Secondly, Power+ achieves the same quality as state-of-the-art studies and even higher. Even for low-quality workers, our methods still achieve high quality, because (1) Power+ can tolerate errors by not coloring unconfident vertices (and thus avoid enlarging the errors by a wrong coloring vertex); and (2) few pairs invalidate the partial order. Specifically, on the Restaurant dataset, as the tasks are very easy, most workers can correctly compare each pair, and thus all the methods achieve high quality. On the Cora dataset, Power+ and ACD achieve much higher quality than Trans and GCER on all three groups of workers, because this dataset is harder and workers may return noisy results. Trans and GCER cannot tolerate workers' errors and moreover they may expand the error propagation due to the transitivity rules. On the ACMP ub dataset, P ower + and ACD still outperform other methods because both of them consider crowd's errors. Thirdly, with the increase of worker accuracy, the quality of all the algorithms increase, because workers return higher quality answers. Fourthly, even for workers with different accuracy, the algorithms achieve similar quality. This is because the worker accuracy on AMT is computed based on their accuracy on history tasks but not on our tasks. A worker will give higher quality on easy datasets, e.g., Restaurant, and lower quality on hard datasets, e.g., Cora. To address this issue, we conduct a simulation experiment in Section 8.6.2.
\#Questions. Firstly, our two methods Power and Power+ ask fewer questions than state-of-the-art methods, even by 2 orders of magnitude. This is because we can utilize the partial order to prune many pairs that do not need to be asked and use the grouping techniques to reduce the graph size. The partial order can prune the pairs with larger similarities than a Green vertex and the pairs with smaller similarities than a RED vertex, while the grouping technique can prune many pairs with similar similarities close to the asked pairs. Trans can also reduce the number of questions based on transitivity at the expense of lowering down the quality. ACD and GCER achieve high quality at the expense of asking many more questions. For example, on ACMP ub, ACD, GCER, and Trans ask 30,000 questions, and our methods ask 150 questions. Thus our methods can save $200 \times$ monetary cost than ACD, GCER, and Trans. On Restaurant, ACD and GCER ask 4100 questions, Trans asks 3900 questions while Power only asks 51 questions. Thus our methods save $80 \times$ monetary cost than ACD, GCER, and Trans. On Cora, ACD and GCER ask 4800 questions, Trans asks 1020 questions while our methods only ask 354 questions. Trans saves a little cost on the Restaurant dataset because only few pairs satisfy the transitivity rules. Secondly,


Fig. 26: Quality Comparison by Varying Worker Accuracy (Real Experiments).


Fig. 27: \#Question Comparison by Varying Worker


Fig. 28: \#Iteration Comparison by Varying Worker Accuracy (Real Experiments).

Power+ asks a few more questions than Power to tolerate the unconfident vertices and avoid coloring their ancestors and descendants. As there are few unconfident vertices, the gap between Power+ and Power is trivial. Thirdly, the worker accuracy has little effect on the number of questions, because (1) our methods ask few questions and the question number is determined by the graph structure but not worker accuracy and (2) other methods do not consider worker accuracy to select questions.
\#Iterations. Firstly, our methods involve fewer iterations than state-of-the-art approaches. This is because (1) our methods ask smaller number of questions and (2) our methods ask questions in parallel. For example, on the Restaurant dataset, ACD involves 13 iterations, GCER involves 28 iterations, Trans involves 23 iterations, while Power+ only involves 5 iterations. On the Cora dataset, ACD involves 18 iterations, Trans involves 10 iterations, GCER involves 19 iterations, while Power+ only involves 4 iterations. On the ACMP ub dataset, ACD involves 15 iterations, Trans involves 9 iterations, GCER involves 13 iterations, while Power + only involves 5 iterations. Thus our method saves $2-5 \times$ latency cost on the Cora dataset. Secondly, Power and Power + nearly have the same number of iterations, as they have little difference on the number of asked questions. Thirdly, the worker accuracy has little impact on the number of iterations, as existing studies do not consider worker accuracy and our methods ask few questions.

(a) Restaurant


Fig. 29: Quality Comparison by Varying Worker Accuracy (Simulation Experiments).


Fig. 30: \#Question Comparison by Varying Worker Accuracy (Simulation Experiments).


Fig. 31: \#Iteration Comparison by Varying Worker Accuracy (Simulation Experiments).

### 8.6.2 Simulation Exp: Evaluating Worker Accuracy

In real experiments, we select workers based on their history accuracy on AMT. However the approval rates of workers only reflect their history accuracy when they answered other questions in history but not the accuracy on our questions. Workers may have different quality on different datasets and time. For example, on the Restaurant dataset, the problem is easy, and most workers give a correct answer even though they have a low history accuracy. On the Cora dataset, because the dataset is relative dirty and professional, many workers return wrong answers even though they have a high history accuracy. Therefore, the reason why many previous studies [13, 44, 45, 46] set a high approval rate is to filter these malicious workers, but this does not mean that most workers can give right answers under the high approval rate guarantee. To address this issue, we conduct a simulation experiment. We assume the ground truth is known and generate workers with quality in $70 \%-80 \%, 80 \%-90 \%$, and above $90 \%$ respectively. Figures 2931 show the results on the simulation experiments.
Quality. Firstly, Power+ significantly outperforms other methods for low-quality workers, e.g., $70 \%$ and $80 \%$. This is because (1) Power+ can tolerate more errors by postponing coloring the unconfident vertices, and for low-quality workers, there are many more unconfident vertices. For example, for those wrongly answered vertices, we do not color them Green or RED immediately and also do not color their ancestors or descendants, which avoid many errors; (2) Power + can tolerate the malicious answers by first ask-
ing middle-level vertices in the graph and thus has low possibility to wrongly label some pairs (because a high-level vertex will affect many vertices if it is colored RED and a low-level vertex will affect many vertices if it is colored Green). ACD outperforms other baselines, because it tolerates errors based on the clusters refinement (each cluster contains records referring to the same entity). However, on the Restaurant dataset, ACD has lower quality, because there are few records in each cluster and ACD cannot utilize this limited information to infer the answers. Trans and GCER cannot tolerate errors, and thus they have rather low quality for low worker accuracy. Secondly, for high worker accuracy, e.g., $90 \%$, all the algorithms achieve rather high quality as there are few errors in the workers' answers. Thirdly, with the increase of worker accuracy, all the methods achieve higher quality as they can utilize high quality answers. Power + outperforms ACD, which in turns is better than other methods. For example, on the Restaurant dataset, in real experiments, all methods have more than $92 \%$ F-measure whatever the workers' accuracy is, because workers have high quality on this easy dataset. However, in our simulation experiment, for $70 \%$ accuracy, Power+ achieves $92 \%$ F-measure while Power, Trans, ACD, and GCER have $76 \%, 65 \%$, $77 \%$ and $75 \%$ F-measure respectively. For $80 \%$ accuracy, Power+ still outperforms other methods, and all the methods have improvement on quality compared with $70 \%$. For $90 \%$ accuracy, all methods can achieve high quality, e.g., $95 \%$, because the workers return high-quality answers. On the Cora dataset, for 70\% accuracy, Power+ and ACD have high F-measure and reach $91 \%$ due to tolerating crowd's errors while Power has F-measure $86 \%$. Trans and GCER expand the error propagation due to the transitivity rule, whose F-measures are $60 \%$ and $65 \%$. For $80 \%$ accuracy, Power+ and ACD improve to $93 \%$. And Power, Trans and GCER are $88 \%, 80 \%$ and $82 \%$ respectively. For $90 \%$ accuracy, all methods improve the quality to above $90 \%$. Power+ and ACD still outperform others. On the ACMPub dataset, similarly to the Restaurant and Cora datasets, Power+ and ACD outperform other methods when worker accuracy is low. If worker accuracy is high, all methods achieve nearly the same quality.
\#Questions. Our method saves $80 \times$ than Trans, ACD and GCER for Restaurant, $10 \times$ than Trans, ACD and GCER for Cora, and $200 \times$ than Trans, ACD and GCER for ACMPub.
\#Iterations. Since the worker accuracy has little impact on the number of iterations, there is little difference between real experiments and simulation experiments. Our method still saves $2-5 \times$ latency cost on the three datasets.

### 8.7 Evaluating Similarity Functions

We evaluate the effect of different similarity functions. On each dataset, we respectively use Jaccard, edit similarity, and


Fig. 32: Quality Comparison by Varying Similarity Functions (Real Experiments).


Fig. 33: \#Question Comparison by Varying Similarity Functions (Real Experiments).


Fig. 34: \#Iteration Comparison by Varying Similarity Functions (Real Experiments).
bigram on every attribute to generate the graph and compare the results for the three similarity functions. Note that on the Restaurant dataset, Jaccard is not a good similarity function for the Name attribute as there are only 1-2 words in the restaurant name; while on the ACMP ub dataset, edit similarity is not a good choice for the Title attribute as there are many words in the paper title. We want to test whether our methods and state-of-the-art approaches can tolerate the noisy results generated by different similarity functions. We use the worker accuracy of $90 \%$. Figures 3234 show the results. We make the following observations.

Firstly, different similarity functions have little impact on the quality among all the methods, because all of these methods use a property that the pairs with large similarities have large possibility to refer to the same entity. On the real datasets, most of the similar functions satisfy this property. For our methods, if the functions do not significantly invalidate the partial order, they can be used in our methods. In addition, existing methods use record-level similarity while our methods utilize attribute-level similarity to evaluate different pairs. Thus our methods can use more information to determine the partial order. Even if there exist some noisy functions on some attributes (e.g., Jaccard for Name and edit similarity for Title), our methods can utilize other similarity functions to obtain a good partial order and thus can correct the errors caused by the noisy similarity functions. Thus our methods are robust on real datasets, even for some noisy functions. Secondly, the similarity functions have little effect on the number of questions, as the number of questions is determined by the graph structure and the graphs gener-


Fig. 35: Budget-Aware: Simulation Experiment (70\%):


Fig. 36: Budget-Aware: Simulation Experiment (70\%):


Fig. 37: Budget-Aware: Simulation Experiment (70\%): F-Measure.
ated by different functions have no much difference. Thirdly, the similarity functions have little impact on the number of iterations, as the number of iterations is determined by the number of questions but not by the similarity functions.

### 8.8 Evaluation on Budget-Aware Methods

We compare our methods (Budget without error-tolerant techniques and Budget + with error-tolerant techniques) with Power+ and other state-of-the-art methods mentioned above. For Budget proposed in Section 7.2, we easily extend it to tolerate errors, i.e. Budget+, by utilizing the method in section 6 We just need to color the vertex BLUE if its confidence is low. After all vertices in the graph are colored, we use the GREEN vertices to color these BLUE ones.

### 8.8.1 Evaluation on Serial Algorithms

We first evaluate the serial algorithm by varying the quality of different kinds of workers. We evaluate the precision, recall and F-Measure by varying the budget. Figures 35 . 43 show the experimental results on the simulation experiments. We make the following observations.

Firstly, different methods have similar precision except Budget because Budget does not tolerate errors. When the worker quality is $70 \%$, we can see that in Figure 35 on the Restaurant dataset, Budget has a low precision, around $60 \%$, because it does not have an error tolerant strategy, which may deduce the colors of non-asked vertices incorrectly. Budget+ and Power+ achieve the highest pre-


Fig. 38: Budget-Aware: Simulation Experiment (80\%): Precision.


Fig. 39: Budget-Aware: Simulation Experiment (80\%):


Fig. 40: Budget-Aware: Simulation Experiment (80\%): F-Measure.
cision, above $90 \%$ because they can tolerate errors by postprocessing low-confident vertices. Though ACD can also tolerate errors, it is worse than Budget + and Power+ because its technique cannot work well under a small budget. Since Trans and GCER do not consider errors, they have a lower precision (about 90\%) than Budget + and Power+. On Cora and ACMPub datasets, Budget + and Power+ both achieve the highest precision, which are about $95 \%$ and $90 \%$ respectively.

Secondly, Budget and Budget+ achieve much higher recall than other methods. When the worker quality is $70 \%$, we can see that in Figure 36 on the Restaurant dataset, even if we only ask 5 questions, Budget can achieve more than $50 \%$ recall because it can select these large benefit vertices. Trans, GCER and ACD achieve nearly 0 recall because they do not have graph model to deduce more results. Furthermore, with 5 questions, there is no enough transitivities that can be used to reduce the cost. The recall of Budget + is 0 because some vertices may be colored BLUE due to the error-tolerant technique. However, with the budget increasing, Budget + achieves the highest recall because it can tolerant error and select vertices with large benefits. When we ask 25 questions, Budget+ achieves recall of $90 \%$, while Power+ only achieves recall of $78 \%$. When we ask more than 25 questions, the recall of Budget + remains stable because we have almost found out all Green vertices. On Cora, at the beginning, Budget+ and Budget achieve more than $60 \%$ recall while other methods only get a very small recall. With the budget increasing, the recall


Fig. 41: Budget-Aware: Simulation Experiment (90\%):
Precision.




Fig. 42: Budget-Aware: Simulation Experiment (90\%): Recall.


Fig. 43: Budget-Aware: Simulation Experiment (90\%):


Fig. 44: Budget-Aware: Real experiment (70\%): Precision.


Fig. 45: Budget-Aware: Real experiment (70\%): Recall.


Fig. 46: Budget-Aware: Real experiment (70\%):

## F-Measure.

of Budget+ remains stable after only asking 50 questions. On ACMPub, similar to other two datasets, Budget + has the best performance on both recall and costs.

Thirdly, since our methods outperform other approaches on recall and have similar precision, our methods achieve better F-measure. For example, in Figure 37, on Restaurant dataset, Budget+ achieves 94\% F-measure while Trans,


Fig. 47: Budget-Aware: Real experiment (80\%): Precision.


Fig. 48: Budget-Aware: Real experiment (80\%): Recall.


Fig. 49: Budget-Aware: Real experiment (80\%): F-Measure.

GCER and ACD only have F-measure of $50 \%$. On Cora and ACMP ub datasets, Budget + has F-measures $83 \%$ and $82 \%$ respectively while other approaches only have F-measures close to zero, because their recalls are close to zero. The reason why Trans, ACD, GCER have such low recall and F-measure is that they aim to minimize the number of questions to get all the answers. For some selected questions, they have low inference power to infer the answers of other questions and thus they have low recall. For example, on Cora dataset, there are 12793 matching pairs. Power+ can achieve a stable high recall given only a budget of 100 questions, because Power+ can infer more than 10000 matching pairs based on these 100 questions. However, Trans can only infer less than 100 matching pairs. In other words, Trans, ACD, and GCER have low inference power, and with a small budget and the large number of matching pairs in ground truth, they cannot infer the answers of many tasks and thus they have low recall.

Fourthly, with the increase of budgets, the precision decreases and the recall increases, because we can utilize more budgets to get more results and the side effect is that it introduces more errors due to transitivities and crowd errors (for Trans, ACD, GCER) and partial order and crowd errors (for Budget and Power). Besides, the F-measure goes up because the recall increases more than the precision decreases.

Fifthly, with the increase of worker quality, the precision increases but the recall keeps stable, as higher worker quality leads to high result accuracy while it cannot improve the recall which is determined by the number of deduced pairs. When the worker quality is $80 \%$, we can see that in Fig-


Fig. 50: Budget-Aware: Real experiment (90\%): Precision.




Fig. 51: Budget-Aware: Real experiment (90\%): Recall.




Fig. 52: Budget-Aware: Real experiment ( $90 \%$ ): F-Measure.
ure 38, the precision of Budget+ remains the highest because it can tolerate errors. However, the difference between other methods and Budget + is reduced because workers are more accurate. For example, on Cora, the precision of Trans, ACD and GCER are about $94 \%$ and Budget + is about $95 \%$. In Figure 39, Budget + achieves the best performance among all methods, and its recall remains stable after several questions and is much higher than others. When the worker quality is $90 \%$, we can see that in Figure 41 , the precision of all methods is high because workers seldom make mistakes. For recall, Budget + achieves the best result. Since the precision increases and the recall keeps stable, F-measure increases obviously. We can see from Figures 37 and 43 that on Cora dataset, F-measure improves from $92 \%$ to $95 \%$ with the increase of workers' accuracy.

We then evaluate the algorithms on real settings on real crowdsourcing platforms. Figures $44 \mid 52$ show the experimental results on the real experiments. We make the following observations. Firstly, for the real experiment, when the worker quality is $70 \%$ (Figure 44), on Restaurant, the precision of all methods is high, around $96 \%$ because although the worker quality is only $70 \%$, the question is easy for human to answer (the voting result by multiple workers has high quality). On Cora, which is difficult for human to answer, when the budget is around 100, Budget + and Power+ achieve the highest precision, about $80 \%$, while other methods is about $75 \%$. Similarly, on ACMP ub, Budget + and Power+ achieve more precision than other methods because they can tolerate errors. For recall, like the simulation experiment, Budget+ performs the best. Budget+ achieves recall of $90 \%, 83 \%$ and $80 \%$ after asking 25,50 and


Fig. 53: Budget Aware. Machine-based Method Comparison: Precision.


Fig. 54: Budget Aware. Machine-based Method Comparison: Recall.


Fig. 55: Budget Aware. Machine-based Method Comparison: F-measure.

30 questions on Restaurant, Cora and ACMP ub respectively. And afterwards, the recall remains stable, which is much better than any other methods. For example, Power+ is only $70 \%$ when budget is 25 on Restaurant. When the worker quality is $80 \%$ or $90 \%$, similarly, the precision of all methods is high and Budget + performs the best on recall on all datasets. At last, since our methods outperform existing techniques on both precision and recall, obviously achieve high F-measure.

Budget and Budget+ have higher recall, precision and F-measure than other methods with different budgets.

### 8.8.2 Evaluation on Parallel Algorithms

We compare the serial and parallel algorithms by varying worker quality on the three datasets. Tables [5, 6 and 7 show the experimental results on the simulation experiments. We make the following observations.

Firstly, the serial algorithm takes many more (about 10$20 \times$ ) rounds than the parallel algorithm. If one cares about the latency, the serial algorithm is not acceptable. Therefore we design the parallel algorithm in section7.2.2. We can see from the table that the parallel algorithm takes 5 iterations to complete the task, which is significantly less than that of serial algorithm. This is because our parallel algorithm can ask many beneficial and independent questions in each iteration. Secondly, the parallel algorithm takes nearly the same quality with the serial algorithm. In Table 5. on Restaurant, for quality, the recall, precision and f-measure are similar to the serial algorithm. And what's more, different worker

Table 5: Evaluation on Parallel Algorithms (Restaurant, $B=65$ ).

| Worker Quality | \# iterations |  | Precision |  | Recall |  | F-measure |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Parallel | Serial | Parallel | Serial | Parallel | Serial | Parallel | Serial |
| $70 \%$ | 5 | 65 | 0.95 | 0.96 | 0.91 | 0.92 | 0.93 | 0.94 |
| $80 \%$ | 5 | 65 | 0.97 | 0.97 | 0.92 | 0.93 | 0.94 | 0.95 |
| $90 \%$ | 5 | 65 | 0.99 | 0.99 | 0.95 | 0.94 | 0.97 | 0.96 |

Table 6: Evaluation on Parallel Algorithms (Cora, $B=100$ ).

| Worker Quality | \# iterations |  | Precision |  | Recall |  | F-measure |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Parallel | Serial | Parallel | Serial | Parallel | Serial | Parallel | Serial |
| $70 \%$ | 5 | 100 | 0.73 | 0.75 | 0.84 | 0.84 | 0.77 | 0.79 |
| $80 \%$ | 5 | 100 | 0.82 | 0.81 | 0.85 | 0.85 | 0.83 | 0.83 |
| $90 \%$ | 5 | 100 | 0.83 | 0.83 | 0.84 | 0.85 | 0.83 | 0.84 |

Table 7: Evaluation on Parallel Algorithms (AcmPub, $B=70$ ).

| Worker Quality | \# iterations |  | Precision |  | Recall |  | F-measure |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Parallel | Serial | Parallel | Serial | Parallel | Serial | Parallel | Serial |
| $70 \%$ | 5 | 70 | 0.82 | 0.83 | 0.84 | 0.84 | 0.83 | 0.83 |
| $80 \%$ | 5 | 70 | 0.85 | 0.87 | 0.85 | 0.86 | 0.85 | 0.86 |
| $90 \%$ | 5 | 70 | 0.92 | 0.92 | 0.88 | 0.87 | 0.90 | 0.89 |

qualities have little impact on the number of iterations. On Cora, the parallel algorithm takes 5 iterations to achieve a high and stable quality among all kinds of worker quality. On ACMP ub, it takes 4 iterations and $B=70$ to achieve the task with high quality.

### 8.8.3 Comparison with Machine-based Algorithm

We compare our method Budget + with state-of-the-art machinebased entity resolution system Magellan [21]. Magellan first uses blocking rules to prune the pairs with low possibility to be matched. Then it active learning to selec some training data and trains a matching model using random forest. Next it uses the matching model to evaluate each nonpruning pairs. Magellan uses active learning to select the training data. We can also use crowdsourced question selection methods, e.g., ACD, Trans and GCER, to select the training data. We compare with Magellan and Magellan +Trans, Magellan +ACD, Magellan +GCER by varying the budget. We call these hybrid approaches as Tran_ML, ACD_ML and GCER_ML. Figures $53 \sqrt{55}$ show the results.

We can see that all methods achieve nearly the same precision, while Budget + is slightly better than the other two. On Cora, the precision of Budget + is about 10 percent higher than other methods because the dataset is dirtier than others and our methdo can tolerate errors. In Figure 54, we can see that Budget + has a much better recall than others. For example, on Restaurant dataset, Budget+ achieves $95 \%$ recall given the budget of 45 questions, while Magellan has $80 \%$ recall and Tran_ML, GCER_ML and ACD ML have $78 \%, 56 \%$ and $55 \%$ respectively. This because machine-based method cannot perform very well given a small size of budget(training set). Tran_ML, GCER_ML and

ACD ML are worse than Magellan becaues they aim to select some questions to infer all the answers but not maximize the inference power with the given budget. Since Budget + can leverage both human intelligence and machine-based features to select suitable questions and infer a large number of answers according to the partial order graph model, it can achieve a high recall. Since our method performs well on both precision and recall, we can see from 55 that our method also performs better on F-measure.

## 9 Conclusion

We proposed a partial-order based crowdsourced entity resolution framework. We defined a partial order on record pairs based on their similarities on every attribute. We proposed a graph-based coloring strategy to deduce the answer of some pairs based on the answers of asked pairs. We devised two algorithms to construct the graph and proposed two grouping methods to reduce the graph size. We proposed effective algorithms to judiciously select pairs to ask in order to minimize the number of asked pairs. We developed error-tolerant techniques to tolerate the errors. We designed a budget-aware algorithm to maximize the number of matching pairs within a given budget. Experimental results show that our method saves more money than existing approaches while keeping the same quality.

## References

1. A.P.Dawid and A.M.Skene. Maximum likelihood estimation of observer error-rates using em algorithm. Appl.Statist., 28(1):2028, 1979.
2. B. I. Aydin, Y. S. Yilmaz, Y. Li, Q. Li, J. Gao, and M. Demirbas. Crowdsourcing for multiple-choice question answering. In AAAI, pages 2946-2953, 2014.
3. C. C. Cao, J. She, Y. Tong, and L. Chen. Whom to ask? jury selection for decision making tasks on micro-blog services. PVLDB, 5(11):1495-1506, 2012.
4. C. Chai, G. Li, J. Li, D. Deng, and J. Feng. Cost-effective crowdsourced entity resolution: A partial-order approach. In Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016, San Francisco, CA, USA, June 26 - July 01, 2016, pages 969-984, 2016.
5. X. Chen, P. N. Bennett, K. Collins-Thompson, and E. Horvitz. Pairwise ranking aggregation in a crowdsourced setting. In WSDM, pages 193-202, 2013.
6. 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.
7. X. Dong, A. Y. Halevy, and J. Madhavan. Reference reconciliation in complex information spaces. In Proceedings of the ACM SIGMOD International Conference on Management of Data, Baltimore, Maryland, USA, June 14-16, 2005, pages 85-96, 2005.
8. B. Eriksson. Learning to top-k search using pairwise comparisons. In AISTATS, pages 265-273, 2013.
9. J. Fan, G. Li, B. C. Ooi, K. Tan, and J. Feng. icrowd: An adaptive crowdsourcing framework. In SIGMOD, pages 1015-1030, 2015.
10. S. Felsner, V. Raghavan, and J. P. Spinrad. Recognition algorithms for orders of small width and graphs of small dilworth number. Order, 20(4):351-364, 2003.
11. M. J. Franklin, D. Kossmann, T. Kraska, S. Ramesh, and R. Xin. Crowddb: answering queries with crowdsourcing. In SIGMOD, pages 61-72, 2011.
12. D. R. Fulkerson. Note on dilworth's decomposition theorem for partially ordered sets. American Mathematical Society, 7(4):701702, 1956.
13. C. Gokhale, S. Das, A. Doan, J. F. Naughton, N. Rampalli, J. W. Shavlik, and X. Zhu. Corleone: hands-off crowdsourcing for entity matching. In SIGMOD, pages 601-612, 2014.
14. S. Guo, A. G. Parameswaran, and H. Garcia-Molina. So who won?: dynamic max discovery with the crowd. In SIGMOD, pages 385-396, 2012.
15. C.-J. Ho, S. Jabbari, and J. W. Vaughan. Adaptive task assignment for crowdsourced classification. In ICML, pages 534-542, 2013.
16. P. G. Ipeirotis, F. Provost, and J. Wang. Quality management on amazon mechanical turk. In SIGKDD workshop on human computation, pages 64-67. ACM, 2010.
17. M. Joglekar, H. Garcia-Molina, and A. G. Parameswaran. Evaluating the crowd with confidence. In SIGKDD, pages 686-694, 2013.
18. D. R. Karger, S. Oh, and D. Shah. Iterative learning for reliable crowdsourcing systems. In NIPS, pages 1953-1961, 2011.
19. L. Kazemi, C. Shahabi, and L. Chen. Geotrucrowd: trustworthy query answering with spatial crowdsourcing. In SIGSPATIAL, pages 304-313, 2013.
20. A. R. Khan and H. Garcia-Molina. Hybrid strategies for finding the max with the crowd. Technical report, 2014.
21. P. Konda, S. Das, P. S. G. C., A. Doan, A. Ardalan, J. R. Ballard, H. Li, F. Panahi, H. Zhang, J. F. Naughton, S. Prasad, G. Krishnan, R. Deep, and V. Raghavendra. Magellan: Toward building entity matching management systems. PVLDB, 9(12):1197-1208, 2016.
22. M. Kreveld, M. Overmars, O. Schwarzkopf, M. d. Berg, and O. Schwartskopf. Computational geometry: algorithms and applications, 1997.
23. G. Li. Human-in-the-loop data integration. $P V L D B, 10(12): 2006-$ 2017, 2017.
24. G. Li, C. Chai, J. Fan, X. Weng, J. Li, Y. Zheng, Y. Li, X. Yu, X. Zhang, and H. Yuan. CDB: optimizing queries with crowdbased selections and joins. In SIGMOD, pages 1463-1478, 2017.
25. G. Li, J. Wang, Y. Zheng, and M. J. Franklin. Crowdsourced data management: A survey. IEEE Trans. Knowl. Data Eng., 28(9):2296-2319, 2016.
26. X. Liu, M. Lu, B. C. Ooi, Y. Shen, S. Wu, and M. Zhang. CDAS: A crowdsourcing data analytics system. PVLDB, 5(10):1040-1051, 2012.
27. A. Marcus, E. Wu, D. R. Karger, S. Madden, and R. C. Miller. Human-powered sorts and joins. $P V L D B, 5(1): 13-24,2011$.
28. A. Marcus, E. Wu, S. Madden, and R. C. Miller. Crowdsourced databases: Query processing with people. In $C I D R$, pages 211214, 2011.
29. N. Megiddo and K. J. Supowit. On the complexity of some common geometric location problems. SIAM journal on computing, 13(1):182-196, 1984.
30. W. R. Ouyang, L. M. Kaplan, and et al. Debiasing crowdsourced quantitative characteristics in local businesses and services. In IPSN, pages 190-201, 2015.
31. A. G. Parameswaran, S. P. Boyd, H. Garcia-Molina, A. Gupta, N. Polyzotis, and J. Widom. Optimal crowd-powered rating and filtering algorithms. $P V L D B, 7(9): 685-696,2014$.
32. A. G. Parameswaran, H. Garcia-Molina, H. Park, N. Polyzotis, A. Ramesh, and J. Widom. Crowdscreen: algorithms for filtering data with humans. In SIGMOD, pages 361-372, 2012.
33. A. G. Parameswaran, H. Park, H. Garcia-Molina, N. Polyzotis, and J. Widom. Deco: declarative crowdsourcing. In CIKM, pages 1203-1212, 2012.
34. H. Park and J. Widom. Query optimization over crowdsourced data. $P V L D B, 6(10): 781-792,2013$.
35. H. Park and J. Widom. Crowdfill: collecting structured data from the crowd. In SIGMOD, pages 577-588, 2014.
36. T. Pfeiffer, X. A. Gao, Y. Chen, A. Mao, and D. G. Rand. Adaptive polling for information aggregation. In AAAI, 2012.
37. V. C. Raykar, S. Yu, L. H. Zhao, A. K. Jerebko, C. Florin, G. H. Valadez, L. Bogoni, and L. Moy. Supervised learning from multiple experts: whom to trust when everyone lies a bit. In ICML, pages 889-896, 2009.
38. A. D. Sarma, A. G. Parameswaran, H. Garcia-Molina, and A. Y. Halevy. Crowd-powered find algorithms. In ICDE, pages 964975, 2014.
39. B. Trushkowsky, T. Kraska, M. J. Franklin, and P. Sarkar. Crowdsourced enumeration queries. In ICDE, pages 673-684, 2013.
40. M. Venanzi, J. Guiver, G. Kazai, P. Kohli, and M. Shokouhi. Community-based bayesian aggregation models for crowdsourcing. In $W W W$, pages $155-164,2014$.
41. P. Venetis, H. Garcia-Molina, K. Huang, and N. Polyzotis. Max algorithms in crowdsourcing environments. In $W W W$, pages 989 998, 2012.
42. V. Verroios and H. Garcia-Molina. Entity resolution with crowd errors. In $I C D E$, pages 219-230, 2015.
43. N. Vesdapunt, K. Bellare, and N. N. Dalvi. Crowdsourcing algorithms for entity resolution. $P V L D B, 2014$.
44. J. Wang, T. Kraska, M. J. Franklin, and J. Feng. Crowder: Crowdsourcing entity resolution. PVLDB, 2012.
45. J. Wang, G. Li, T. Kraska, M. J. Franklin, and J. Feng. Leveraging transitive relations for crowdsourced joins. In SIGMOD, pages 229-240, 2013.
46. S. Wang, X. Xiao, and C. Lee. Crowd-based deduplication: An adaptive approach. In SIGMOD, pages 1263-1277, 2015.
47. S. E. Whang, P. Lofgren, and H. Garcia-Molina. Question selection for crowd entity resolution. $P V L D B, 2013$.
48. J. Whitehill, P. Ruvolo, T. Wu, J. Bergsma, and J. R. Movellan. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In NIPS, pages 2035-2043, 2009.
49. P. Ye, U. EDU, and D. Doermann. Combining preference and absolute judgements in a crowd-sourced setting. In ICML Workshop, 2013.
50. C. J. Zhang, L. Chen, H. V. Jagadish, and C. C. Cao. Reducing uncertainty of schema matching via crowdsourcing. PVLDB, 6(9):757-768, 2013.
51. Z. Zhao, F. Wei, M. Zhou, W. Chen, and W. Ng. Crowd-selection query processing in crowdsourcing databases: A task-driven approach. In EDBT, pages 397-408, 2015.
52. Z. Zhao, D. Yan, W. Ng, and S. Gao. A transfer learning based framework of crowd-selection on twitter. In $\operatorname{SIGKDD}$, pages 1514-1517, 2013.
53. Y. Zheng, R. Cheng, S. Maniu, and L. Mo. On optimality of jury selection in crowdsourcing. In EDBT, pages 193-204, 2015.
54. Y. Zheng, G. Li, and R. Cheng. DOCS: domain-aware crowdsourcing system. $P V L D B, 10(4): 361-372,2016$.
55. Y. Zheng, G. Li, Y. Li, C. Shan, and R. Cheng. Truth inference in crowdsourcing: Is the problem solved? $P V L D B, 10(5): 541-552$, 2016.
56. Y. Zheng, J. Wang, G. Li, R. Cheng, and J. Feng. QASCA: A quality-aware task assignment system for crowdsourcing applications. In SIGMOD, pages 1031-1046, 2015.

[^0]:    Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China
    E-mail: \{chaicl15,dd11\}@mails.tsinghua.edu.cn
    \{liguoliang,lijian83,fengjh\}@tsinghua.edu.cn

[^1]:    1 Note to avoid duplicately comparing two pairs in $\mathcal{U}\left(p_{i j}\right)$, we can only select pivots from $\mathcal{C}\left(p_{i j}\right)$ and $\mathcal{P}\left(p_{i j}\right)$.

[^2]:    2 http://www.cs.utexas.edu/users/ml/riddle/data/restaurant.tar.gz
    3 https://www.cics.umass.edu/smccallum/data/cora-refs.tar.gz
    $4 \mathrm{http}: / / \mathrm{dbs}$.uni-leipzig.de/en/research/projects/object_matching/ fever/benchmark_datasets_for_entity_resolution

[^3]:    5 The three datasets have 4-8 attributes. As it is too complicated to construct a high dimensional range tree, we use a heuristics: we choose two important attributes in each dataset to construct 2-dimensional indexes. When we search the children of a pair, the pairs reported by the index are a superset as they may not satisfy other attributes. To address this issue, we only need to verify them to remove the false positives based on other non-indexed attributes. In our experiment, we choose attributes Name and Address for Restaurant, Author and Title for Cora, and Author and Title for ACMPub.

