Crowdsourcing Pareto-Optimal Object Finding by Pairwise Comparisons

Abolfazl Asudeh, Gensheng Zhang, Naeemul Hassan, Chengkai Li, Gergely V. Zaruba
Department of Computer Science and Engineering, The University of Texas at Arlington
{ab.asudeh,gensheng.zhang,naeemul.hassan}@mavs.uta.edu
{cli,zaruba}@uta.edu

ABSTRACT

This is the first study of crowdsourcing Pareto-optimal object finding over partial orders and by pairwise comparisons, which has applications in public opinion collection, group decision making, and information exploration. Departing from prior studies on crowdsourcing skyline and ranking queries, it considers the case where objects do not have explicit attributes and preference relations on objects are strict partial orders. The partial orders are derived by aggregating crowdsourcers' responses to pairwise comparison questions. The goal is to find all Pareto-optimal objects by the fewest possible questions. It employs an iterative question-selection framework. Guided by the principle of eagerly identifying non-Pareto optimal objects, the framework only chooses candidate questions which must satisfy three conditions. This design is both sufficient and efficient, as it is proven to find a short terminal question sequence. The framework is further steered by two ideas—macro-ordering and micro-ordering. By different micro-ordering heuristics, the framework is instantiated into several algorithms with varying power in pruning questions. Experiment results using both real crowdsourcing marketplace and simulations exhibited not only orders of magnitude reductions in questions when compared with a brute-force approach, but also close-to-optimal performance from the most efficient instantiation.

1. INTRODUCTION

The growth of user engagement and functionality in crowdsourcing platforms has made computationally challenging tasks unprecedentedly convenient. The subject of our study is one such task—crowdsourcing *Pareto-optimal object finding*. Consider a set of objects O and a set of criteria C for comparing the objects. An object $x \in O$ is *Pareto-optimal* if and only if x is not dominated by any other object. Object y dominates x (denoted $y \succ x$) if and only if x is not better than y by any criterion and y is better than x by at least one criterion, i.e., $\forall c \in C : x \not\succ_c y$ and $\exists c \in C : y \succ_c x$. If x and y do not dominate each other (i.e., $x \not\succ y$ and $y \not\succ x$), we denote it by $x \sim y$. The *preference (better-than) relation P_c* (also denoted $\succ_c$) for each $c \in C$ is a binary relation subsumed by $O \times O$, in which a tuple $(x, y) \in P_c$ (also denoted $x \succ_c y$) is interpreted as "x is better than (preferred over) y with regard to criterion c". Hence,

if $(x, y) \notin P_c$ (also denoted $x \not\succ_c y$), x is not better than y by criterion c. We say x and y are *indifferent* regarding c (denoted $x \sim_c y$), if $(x, y) \notin P_c \wedge (y, x) \notin P_c$. We consider the setting where each P_c is a *strict partial order*, i.e., P_c is irreflexive ($\forall x : (x, x) \notin P_c$) and transitive ($\forall x, y : (x, y) \in P_c \wedge (y, z) \in P_c \Rightarrow (x, z) \in P_c$), which together imply asymmetry ($\forall x, y : (x, y) \in P_c \Rightarrow (y, x) \notin P_c$). We note that such definition of better-than relation has been widely used in modeling preferences (e.g., [20, 13, 28]), and the definition of Pareto-optimal objects follows the concept of *Pareto composition* of preference relations in [13].

Novelty Pareto-optimal queries resembles *skyline* queries [5]. However, except for [23, 4, 17], previous studies on preference and skyline queries do not use the crowd; they focus on query processing on *existing* data. On the contrary, we examine how to obtain sufficient data from the crowd for determining Pareto-optimal objects. Furthermore, our work differs in several radical ways, as summarized in Table 1 and explained below.

- The preference relation for a criterion is *not* governed by explicit scores or values on object attributes (e.g., sizes of houses, prices of hotels), while prior studies (except [4, 17]) assumed explicit attribute representation. For many comparison criteria, it is difficult to model objects by explicit attributes, not to mention asking people to provide such values or scores; people's preferences are rather based on complex, subtle perceptions, as Examples 1 and 2 shall demonstrate. The concept definitions deliberately do not admit notations for attribute or object equivalence.

- Due to the above reason, we request crowdsourcers to perform pairwise comparisons instead of directly providing attribute values or scores. On the contrary, [23] assumes explicit attribute representation and thereby answers skyline queries by asking the crowd to provide missing attribute values. Pairwise comparison is extensively studied in social choice and welfare, preferences, and voting. It is known that people are more comfortable and confident with comparing objects than directly scoring them, since it is easier, faster, and less error-prone [30].

- The crowd's preference relations are modeled as strict partial orders, as opposed to total orders. This is not only a direct effect of using pairwise comparisons instead of numeric scores or explicit attribute values, but also a reflection of the psychological nature of human's preferences [20, 13], since it is not always natural to enforce a total order. Most studies on skyline queries assume total orders, except for [11, 28, 29, 32] which consider partial orders.

Several recent studies used crowdsourcing to compare objects for answering ranking, top-k and group-by queries. Crowd-BT [12] ranks objects by crowdsourcing pairwise object comparisons. Polychronopoulos et al. [25] find top-k items in an itemset by asking human workers to rank small subsets of items. Davidson et al. [15] evaluate top-k and group-by queries by asking the crowd to answer

	Task	Question type	Multiple attributes	Order among objects (on each attribute)	Explicit attribute representation
[12]	full ranking	pairwise comparison	no	total order	no
[25]	top-k ranking	rank subsets of objects	no	total order	no
[15]	top-k ranking and grouping	pairwise comparison	no	total order	no
[23]	skyline/Pareto-optimal queries	missing value inquiry	yes	total order	yes
[17]	skyline/Pareto-optimal queries	pairwise comparison	yes	total order	no
[4], this work	skyline/Pareto-optimal queries	pairwise comparison	yes	strict partial order	no

Table 1: Related work comparison.

(a) Preference relations (i.e., strict partial orders) on three criteria.

QUESTION	ANSWER			OUTCOME
	$\succ$	$\sim$	$\prec$	
a$?_s$b	1	0	4	b$\succ_s$a
a$?_s$c	0	0	5	c$\succ_s$a
a$?_s$d	0	2	3	d$\succ_s$a
a$?_s$e	4	0	1	a$\succ_s$e
a$?_s$f	3	1	1	a$\succ_s$f
b$?_s$c	1	2	2	b$\sim_s$c
b$?_s$d	1	3	1	b$\sim_s$d
b$?_s$e	5	0	0	b$\succ_s$e
b$?_s$f	4	1	0	b$\succ_s$f
c$?_s$d	3	2	0	c$\succ_s$d
c$?_s$e	4	0	1	c$\succ_s$e
c$?_s$f	3	1	1	c$\succ_s$f
d$?_s$e	3	0	2	d$\succ_s$e
d$?_s$f	3	2	0	d$\succ_s$f
e$?_s$f	1	1	3	f$\succ_s$e

(b) Deriving the preference relation for criterion *story* by pairwise comparisons. Each comparison is performed by 5 workers. $\theta = 60\%$.

Figure 1: Finding Pareto-optimal movies by *story*, *music*, *acting*.

type questions (whether two objects belong to the same group) and *value* questions (ordering two objects). They do not consider multiple attributes in modeling objects, while skyline and Pareto-optimal objects are defined in a space of multiple attributes. They assume total orders instead of partial orders.

To recap, ours is so far the only work on crowdsourcing Pareto-optimal object finding over partial orders, and we reported in the extended version of this paper [4] the first study of crowdsourcing Pareto-optimal queries by pairwise comparisons and without explicit attribute representation.

Motivating Applications Pareto-optimal object finding lends itself to applications in several areas, including public opinion collection, group decision making, and information exploration, exemplified by the following motivating examples.

Example 1 (Collecting Public Opinion and Group Decision Making)**.** Consider a set of movies $O=\{a,b,c,d,e,f\}$ and a set of criteria $C=\{story, music, acting\}$ (denoted by s, m, a in the ensuing discussion). Fig.1a shows the individual preference relations (i.e., strict partial orders), one per criterion. Each strict partial order is graphically represented as a directed acyclic graph (DAG), more specifically a Hasse diagram. The existence of a simple path from x to y in the DAG means x is better than (preferred to) y by the corresponding criterion. For example, $(a, e) \in P_m$ (a$\succ_m$e), i.e., a is better than e by *music*. $(b, d) \notin P_s$ and $(d, b) \notin P_s$; hence b$\sim_s$d.

Between **A separation(2011)** and **The big Lebowski(1998)**, which movie is better with regard to *story*?

○ A separation(2011).
○ The big Lebowski(1998).
○ no preference.

[submit] [skip this question]

Figure 2: A question that asks to compare two movies by *story*.

The partial orders define the dominance relation between objects. For instance, movie c dominates d (c$\succ$d), because c is preferred than d on *story* and *music* and they are indifferent on *acting*, i.e., c$\succ_s$d, c$\succ_m$d, and c$\sim_a$d; a and b do not dominate each other (a$\sim$b), since b$\succ_s$a, a$\succ_m$b and b$\succ_a$a. Based on the three partial orders, b is the only Pareto-optimal object, since no other objects dominate it and every other object is dominated by some object.

Note that tasks such as the above one may be used in both understanding the public's preference (i.e., the preference relations are collected from a large, anonymous crowd) and making decisions for a target group (i.e., the preference relations are from a small group of people). ☐

Example 2 (Information Exploration)**.** Consider a photography enthusiast, *Amy*, who is drown in a large number of photos she has taken and wants to select a subset of the better ones. She resorts to crowdsourcing for the task, as it has been exploited by many for similar tasks such as photo tagging, location/face identification, sorting photos by (guessed) date, and so on. Particularly, she would like to choose Pareto-optimal photos with regard to *color*, *sharpness* and *landscape*. ☐

By definition, the crux of finding Pareto-optimal objects lies in obtaining the preference relations, i.e., the strict partial orders on individual criteria. Through crowdsourcing, the preference relations are derived by aggregating the crowd's responses to *pairwise comparison* tasks. Each such comparison between objects x and y by criterion c is a question, denoted x$?_c$y, which has three possible outcomes—x$\succ_c$y, y$\succ_c$x, and x$\sim_c$y, based on the crowd's answers. An example is as follows.

Example 3 (Deriving Preference Relations from Pairwise Comparisons by the Crowd)**.** Fig.1b shows the hypothetical results of all 15 pairwise comparisons between the 6 movies in Example 1, by criterion $s=story$. The outcomes of all comparisons form the crowd's preference relation on *story* (the leftmost DAG in Fig.1a). Fig.2 is the screenshot of a question form designed for one such comparison. A crowdsourcer, when facing this question, would make a choice among the three possible answers or skip a question if they do not have enough confidence or knowledge to answer it. Fig.1b shows how many crowdsourcers have selected each answer. For instance, for question a$?_s$f, three people preferred movie a, one person preferred f, and one person is indifferent. By aggregating these answers, it is derived that a is better than f with regard to *story*, since 60% of the crowdsourcers who responded to the

question chose this answer. For question $b?_s c$, the result is $b \sim_s c$, since neither $b \succ_s c$ nor $b \prec_s c$ received enough votes. (Assuming a threshold $\theta = 60\%$, i.e., either $b \succ_s c$ or $b \prec_s c$ should have at least 60% of votes, in order to not declare $b \sim_s c$.) $\qquad\square$

It is important to point out that, in the aforementioned applications, there is an intrinsic lack of the notion of "ground truth". Since objects are not explicitly represented by attribute values, the partial orders are purely reflections of crowdsourcers' opinions. For this reason, we shall not investigate if crowdsourcers can obtain answers close to "ground truth", which does not exist. This also means that our techniques may get different answers when using different crowdsourcers or even the same crowdsourcers at different times, because people's opinions alter.

Algorithmic Framework Our goal is to find all Pareto-optimal objects with as few questions as possible. A brute-force approach will obtain complete preference relations via pairwise comparisons of all object pairs by every criterion. However, without such exhaustive comparisons, incomplete knowledge collected from a small set of questions may suffice in discerning all Pareto-optimal objects. Toward this end, it may appear that we can take advantage of the transitivity of object dominance—a cost-saving property often exploited in skyline query algorithms (e.g., [5]) to exclude dominated objects from participating in any future comparison once they are detected. But, as also discussed in [13], object dominance in our case is *not* transitive (Property 1), due to the lack of explicit attribute representation. Hence, the aforementioned cost-saving technique is inapplicable.

Aiming at Pareto-optimal object finding by a short sequence of questions, we introduce a general, iterative algorithm framework (Sec.3). Each iteration goes through four steps—*question selection*, *outcome derivation*, *contradiction resolution*, and *termination test*. In the i-th iteration, a question $q_i = x?_c y$ is selected and its outcome is determined based on crowdsourcers' answers. On unusual occasions, if the outcome presents a contradiction to the obtained outcomes of other questions, it is changed to the closest outcome such that the contradiction is resolved. Based on the transitive closure of the outcomes to the questions so far, the objects O are partitioned into three sets—$O_\checkmark$ (objects that must be Pareto-optimal), $O_\times$ (objects that must be non-Pareto optimal), and $O_?$ (objects whose Pareto-optimality cannot be fully discerned by the incomplete knowledge so far). When $O_?$ becomes empty, $O_\checkmark$ contains all Pareto-optimal objects and the algorithm terminates. The question sequence so far is thus a *terminal sequence*.

There are a vast number of terminal sequences. Our goal is to find one that is as short as possible. We observe that, for a non-Pareto optimal object, knowing that it is dominated by at least one object is sufficient, and we do not need to find all its dominating objects. It follows that we do not really care about the dominance relation between non-Pareto optimal objects and we can skip their comparisons. Hence, the overriding principle of our question selection strategy is to identify non-Pareto optimal objects as early as possible. Guided by this principle, the framework only chooses from *candidate questions* which must satisfy three conditions (Sec.3.1). This design is sufficient, as we prove that an empty candidate question set implies a terminal sequence, and vice versa (Proporty 2). The design is also efficient, as we further prove that, if a question sequence contains non-candidate questions, there exists a shorter or equally long sequence with only candidate questions that produces the same $O_\times$, matching the principle of eagerly finding non-Pareto optimal objects (Theorem 1). Moreover, by the aforementioned principle, the framework selects in every iteration such a candidate question $x?_c y$ that x is more likely to be dominated by y. The selection is steered by two ideas—*macro-*

ordering and *micro-ordering*. By using different micro-ordering heuristics, the framework is instantiated into several algorithms with varying power in pruning questions (Sec.4). We also derive a lower bound on the number of questions required for finding all Pareto-optimal objects (Theorem 2).

In summary, this paper makes the following contributions:
- This is the first work on crowdsourcing Pareto-optimal object finding over partial orders and by pairwise comparisons. We define preference relations based on pairwise comparisons and we aim to find all Pareto-optimal objects by fewest possible comparisons.
- We propose a general, iterative algorithm framework (Sec.3) which follows the strategy of choosing only candidate questions that must satisfy three conditions. We prove important properties that establish the advantage of the strategy (Sec.3.1).
- We design macro-ordering and micro-ordering heuristics for finding a short terminal question sequence. Based on the heuristics, the generic framework is instantiated into several algorithms (RandomQ, RandomP, FRQ) with varying efficiency. We also derive a non-trivial lower bound on the number of required pairwise comparison questions. (Sec.4)
- We carried out experiments using both simulations and real crowdsourcing marketplace to compare the amount of comparisons by different instantiations of the framework under varying problem sizes. The results demonstrate the effectiveness of selecting only candidate questions, macro-ordering, and micro-ordering. When these ideas are stacked together, they use orders of magnitude less comparisons than a brute-force approach. The results reveal that FRQ is nearly optimal and the lower bound is practically tight, since FRQ gets very close to the lower bound. (Sec.5)

2. RELATED WORK

Besides [12], there were multiple studies on ranking objects by pairwise comparisons, which date back to decades ago as aggregating the preferences of multiple agents has always been a fundamental problem in social choice and welfare [14, 3]. The more recent studies can be categorized into three types: **1)** Approaches such as [21, 26, 31] predict users' object ranking by completing a user-object scoring matrix. Their predictions take into account users' similarities in pairwise comparisons, resembling *collaborative filtering* [16]. They thus do not consider explicit attribute representation for objects. **2)** Approaches such as [7, 9, 8] infer query-specific (instead of user-specific) ranked results to web search queries. Following the paradigm of *learning-to-rank* [22], they rank a query's result documents according to pairwise result comparisons of other queries. The documents are modeled by explicit ranking features. **3)** Approaches such as [6, 19, 1, 2, 24] are similar to [12] as they use pairwise comparisons to infer a single ranked list that is neither user-specific nor query-specific. Among them, [19] is special in that it also applies learning-to-rank and requires explicit feature representation. Different from our work, none of these studies is about Pareto-optimal objects, since they all assume a total order among objects; those using learning-to-rank require explicit feature representation, while the rest do not consider multiple attributes. Moreover, except [19, 1, 2], they all assume comparison results are already obtained before their algorithms kick in. In contrast, we aim at minimizing the pairwise comparison questions to ask in finding Pareto-optimal objects.

3. GENERAL FRAMEWORK

By the definition of Pareto-optimal objects, the key to finding such objects is to obtain the preference relations, i.e., the strict partial orders on individual criteria. Toward this end, the most

basic operation is to perform *pairwise comparison*—given a pair of objects x and y and a criterion c, determine whether one is better than the other (i.e., $(x,y) \in P_c$ or $(y,x) \in P_c$) or they are indifferent (i.e., $(x,y) \notin P_c \wedge (y,x) \notin P_c$).

The problem of crowdsourcing Pareto-optimal object finding is thus essentially crowdsourcing pairwise comparisons. Each comparison task between x and y by criterion c is presented to the crowd as a question q (denoted $x?_c y$). The outcome to the question (denoted $rlt(q)$) is aggregated from the crowd's answers. Given a set of questions, the outcomes thus contain an (incomplete) knowledge of the crowd's preference relations for various criteria. Fig.2 illustrates the screenshot of one such question (comparing two movies by *story*) used in our empirical evaluation. We note that there are other viable designs of question, e.g., only allowing the first two choices ($x \succ_c y$ and $y \succ_c x$). Our work is agnostic to the specific question design.

Given n objects and r criteria, a brute-force approach will perform pairwise comparisons on all object pairs by every criterion, which leads to $r \cdot n \cdot (n-1)/2$ comparisons. The corresponding question outcomes amount to the complete underlying preference relations. The quadratic nature of the brute-force approach renders it wasteful. The bad news is that, in the worst case, we cannot do better than it. To understand this, consider the scenario where all objects are indifferent by every criterion. If any comparison $x?_c y$ is skipped, we cannot determine if x and y are indifferent or if one dominates another.

In practice, though, the outlook is much brighter. Since we look for only Pareto-optimal objects, it is an overkill to obtain complete preference relations. Specifically, for a Pareto-optimal object, knowing it is not dominated by any object is sufficient, and we do not need to find all the objects dominated by it; for a non-Pareto optimal object, knowing it is dominated by at least one object is sufficient, and we do not need to find all its dominating objects. Hence, without exhausting all possible comparisons, incomplete knowledge on preference relations collected from a set of questions may suffice in fully discerning all Pareto-optimal objects.

Our objective is to find all Pareto-optimal objects with as few questions as possible. By pursuing this goal, we are applying a very simple cost model—the cost of a solution only depends on its number of questions. Although the cost of a task in a crowdsourcing environment may depend on monetary cost, latency and other factors, the number of questions is a generic, platform-independent cost measure and arguably proportionally correlates with the real cost. Therefore, we assume a sequential execution model which asks the crowd an ordered sequence of questions $Q = \langle q_1, ..., q_n \rangle$—it only asks q_{i+1} after $rlt(q_i)$ is obtained. Thereby, we do not consider asking multiple questions concurrently. In practice, such a parallel-execution framework will lead to shorter latency but more questions (which may imply higher monetary cost) than our sequential-execution framework. One interesting direction of future work is to adapt our algorithms to a parallel framework and investigate the tradeoffs between latency and cost. Furthermore, in discussion of our approach, the focus shall be on how to find a short question sequence instead of the algorithms' complexity, since the time taken by crowdsourcers to answer questions will dominate the algorithms' execution time.

To find a short sequence, we design a general algorithm framework, as displayed in Fig.3. Alg.1 shows the framework's pseudocode. Its execution is iterative. Each iteration goes through four steps—*question selection*, *outcome derivation*, *contradiction resolution*, and *termination test*. In the i-th iteration, a question $q_i = x?_c y$ is selected and presented to the crowd. The question outcome $rlt(q_i)$ is derived from the crowd's aggregated answers. On unusu-
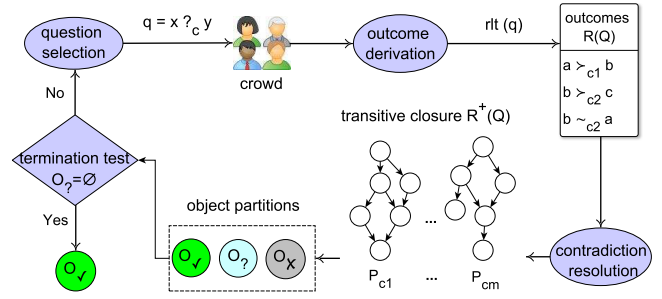

Figure 3: The general framework.

Algorithm 1: The general framework

Input: O : the set of objects
Output: $O_\checkmark$: Pareto-optimal objects of O

1 $R(Q) \leftarrow \varnothing$;   /* question outcomes */
2 **repeat**
3     $x?_c y \leftarrow$ question selection;
4     $rlt(x?_c y) \leftarrow$ outcome derivation;   /* resolve conflict, if any */
5     $R(Q) \leftarrow R(Q) \cup \{rlt(x?_c y)\}$;
6     $(O_\checkmark, O_\times, O_?) \leftarrow$ partitioning objects based on $R^+(Q)$;
        /* $R^+(Q)$ is the transitive closure of $R(Q)$ */
7 **until** $O_? = \{\}$;
8 **return** $O_\checkmark$;

---

al occasions, if the outcome presents a contradiction to the obtained outcomes of other questions so far, it is changed to the closest outcome to resolve contradiction. By computing $R^+(Q_i)$, the *transitive closure* of $R(Q_i)$—the obtained outcomes to questions so far $\langle q_1, \ldots, q_i \rangle$, the outcomes to certain questions are derived and such questions will never be asked. Based on $R^+(Q_i)$, if every object is determined to be either Pareto-optimal or non-Pareto optimal without uncertainty, the algorithm terminates.

Below, we discuss outcome derivation and termination test. Sec.3.1 examines the framework's key step—question selection, and Sec.3.2 discusses contradiction resolution.

Outcome derivation Given a question $x?_c y$, its outcome $rlt(x?_c y)$ must be aggregated from multiple crowdsourcers, in order to reach a reliable result with confidence. Particularly, one of three mutually-exclusive outcomes is determined based on k crowdsourcers' answers to the question:

$$rlt(x?_c y) = \begin{cases} x \succ_c y & \text{if } \frac{\#x}{k} \geq \theta \\ y \succ_c x & \text{if } \frac{\#y}{k} \geq \theta \\ x \sim_c y \ (x \not\succ_c y \wedge y \not\succ_c x) & \text{otherwise} \end{cases} \quad (1)$$

where θ is such a predefined threshold that $\theta > 50\%$, $\#x$ is the number of crowdsourcers (out of k) preferring x over y on criterion c, and $\#y$ is the number of crowdsourcers preferring y over x on c. Fig.1b shows the outcomes of all 15 questions according to Equation (1) for comparing movies by *story* using $k=5$ and $\theta=60\%$. Other conceivable definitions may be used in determining the outcome of $x?_c y$. For example, the outcome may be defined as the choice (out of the three possible choices) that receives the most votes from the crowd. The ensuing discussion is agnostic to the specific definition.

The current framework does not consider different levels of confidence on question outcomes. The confidence on the outcome of a question may be represented as a probability value based on the distribution of crowdsourcers' responses. An interesting direction for future work is to find Pareto-optimal objects in probabilistic sense. The confidence may also reflect the crowdsourcers' quality and credibility [18].

Termination test In each iteration, Alg.1 partitions the objects into three sets by their Pareto-optimality based on the transitive closure of question outcomes so far. If every object's Pareto-optimality has been determined without uncertainty, the algorithm terminates. Details are as follows.

Definition 1 (Transitive Closure of Outcomes). Given a set of questions $Q=\langle q_1,...,q_n\rangle$, the transitive closure of their outcomes $R(Q)=\{rlt(q_1),...,rlt(q_n)\}$ is $R^+(Q)=\{x\sim_c y \mid x\sim_c y \in R(Q)\}\bigcup\{x\succ_c y \mid (x\succ_c y \in R(Q))\vee(\exists w_1,w_2,...,w_m : w_1=x, w_m=y \wedge (\forall 0<i<m : w_i \succ_c w_{i+1} \in R(Q)))\}$. □

In essence, the transitive closure dictates $x\succ_c z$ without asking the question $x?_c z$, if the existing outcomes $R(Q)$ (and recursively the transitive closure $R^+(Q)$) contains both $x\succ_c y$ and $y\succ_c z$. Based on $R^+(Q)$, the objects O can be partitioned into three sets:
$O_{\surd} = \{x \in O \mid \forall y \in O : (\exists c \in C : x\succ_c y \in R^+(Q))\vee(\forall c \in C : x\sim_c y \in R^+(Q))\}$;
$O_{\times} = \{x \in O \mid \exists y \in O : (\forall c \in C : y\succ_c x \in R^+(Q)\vee x\sim_c y \in R^+(Q))\wedge(\exists c \in C : y\succ_c x \in R^+(Q))\}$;
$O_? = O\backslash(O_{\surd}\cup O_{\times})$.
$O_{\surd}$ contains objects that must be Pareto-optimal, $O_{\times}$ contains objects that cannot possibly be Pareto-optimal, and $O_?$ contains objects for which the incomplete knowledge $R^+(Q)$ is insufficient for discerning their Pareto-optimality. The objects in $O_?$ may turn out to be Pareto-optimal after more comparison questions. If the set $O_?$ for a question sequence Q is empty, $O_{\surd}$ contains all Pareto-optimal objects and the algorithm terminates. We call such a Q a *terminal sequence*, defined below.

Definition 2 (Terminal Sequence). A question sequence Q is a terminal sequence if and only if, based on $R^+(Q)$, $O_?=\varnothing$.

3.1 Question Selection

Given objects O and criteria C, there can be a huge number of terminal sequences. Our goal is to find a sequence as short as possible. As Fig.3 and Alg.1 show, the framework is an iterative procedure of object partitioning based on question outcomes. It can also be viewed as the process of moving objects from $O_?$ to $O_{\surd}$ and $O_{\times}$. Once an object is moved to $O_{\surd}$ or $O_{\times}$, it cannot be moved again. With regard to this process, we make two important observations, as follows.

- *In order to declare an object x not Pareto-optimal, it is sufficient to just know x is dominated by another object*. It immediately follows that we do not really care about the dominance relationship between objects in $O_{\times}$ and thus can skip the comparisons between such objects. Once we know $x\in O_?$ is dominated by another object, it cannot be Pareto-optimal and is immediately moved to $O_{\times}$. Quickly moving objects into $O_{\times}$ can allow us skipping many comparisons between objects in $O_{\times}$.
- *In order to declare an object x Pareto-optimal, it is necessary to know that no object can dominate x.* This means we may need to compare x with all other objects including non Pareto-optimal objects. As an extreme example, x may be dominated by only a non-Pareto optimal object y but not by any other object (not even the objects dominating y). This is because object dominance based on preference relations is intransitive, which is formally stated in Property 1. The intransitivity of preference relation formed by Pareto composition of strict partial orders was discussed in [13].

Property 1 (Intransitivity of Object Dominance). Object dominance based on the preference relations over a set of criteria is not transitive. Specifically, if $x\succ y$ and $y\succ z$, it is not necessarily true that $x\succ z$. In other words, it is possible that $x\sim z$ or even $z\succ x$. □

We show the intransitivity of object dominance by an example. Consider objects $O=\{x,y,z\}$, criteria $C=\{c_1, c_2, c_3\}$, and the preference relations in Fig.4. Three dominance relationships violate

Figure 4: Intransitivity of object dominance: $x\succ y$, $y\succ z$, $z\succ x$.

transitivity: (i) $x\succ y$ (based on $x\succ_{c_1}y$, $x\sim_{c_2}y$, $x\sim_{c_3}y$), (ii) $y\succ z$ (based on $y\sim_{c_1}z$, $y\succ_{c_2}z$, $y\sim_{c_3}z$), and (iii) $z\succ x$ (based on $z\sim_{c_1}x$, $z\sim_{c_2}x$, $z\succ_{c_3}x$). As another example, in Fig.1a, $b\succ c$ (since $b\sim_s c$, $b\sim_m c$, $b\succ_a c$, where $s=story$, $m=music$, $a=acting$) and $c\succ a$ (since $c\succ_s a$, $c\sim_m a$, $c\succ_a a$), but $a\sim b$ (since $b\succ_s a$, $a\succ_m b$, $b\succ_a a$) where transitivity does not hold.

Differently, transitivity of object dominance holds in skyline analysis [5]. The contradiction is due to the lack of explicit attribute representation—in our case two objects may be considered equally good on a criterion if they are indifferent, while in skyline analysis they are equally good regarding an attribute if they bear identical values. Skyline query algorithms exploit the transitivity of object dominance to reduce execution cost, because an object can be immediately excluded from further comparison once it is found dominated by any other object. However, due to Property 1, we cannot leverage such pruning anymore.

Based on these observations, the overriding principle of our question selection strategy (Alg.2) is to identify non-Pareto optimal objects as early as possible. At every iteration of the framework (Alg.1), we choose to compare x and y by criterion c (i.e., ask question $x?_c y$) where $x?_c y$ belongs to *candidate questions*. Such candidate questions must satisfy three conditions (Definition 3). There can be many candidate questions. In choosing the next question, by the aforementioned principle, we select such $x?_c y$ that x is more likely to be dominated by y. More specifically, we design two ordering heuristics—*macro-ordering* and *micro-ordering*. Given the three object partitions $O_{\surd}$, $O_{\times}$ and $O_?$, the macro-ordering idea is simply that we choose x from $O_?$ (required by one of the conditions on candidate questions) and y from $O_{\surd}\cup O_?$ (if possible) or $O_{\times}$ (otherwise). The reason is that it is less likely for an object in $O_{\times}$ to dominate x. Micro-ordering further orders all candidate questions satisfying the macro-ordering heuristic. In Sec.4, we instantiate the framework into a variety of solutions with varying power in pruning questions, by using different micro-ordering heuristics.

Definition 3 (Candidate Question). Given Q, the set of asked questions so far, $x?_c y$ is a *candidate question* if and only if it satisfies the following conditions:

(i) The outcome of $x?_c y$ is unknown yet, i.e., $rlt(x?_c y)\notin R^+(Q)$;

(ii) x must belong to $O_?$;

(iii) Based on $R^+(Q)$, the possibility of $y\succ x$ must not be ruled out yet, i.e., $\nexists c' \in C : x \succ_{c'} y \in R^+(Q)$.

Algorithm 2: Question selection

Input: $R^+(Q_i), O_{\surd}(Q_i), O_?(Q_i), O_{\times}(Q_i)$
Output: Q^1_{can} or Q^2_{can}

1   $Q_{can} \leftarrow \{x?_c y \mid rlt(x?_c y)\notin R^+(Q_i) \wedge x \in O_?(Q_i) \wedge (\nexists c' \in C : x \succ_{c'} y \in R^+(Q_i))\}$;
2   $Q^1_{can} \leftarrow \{x?_c y \mid x?_c y \in Q_{can}, y \notin O_{\times}(Q_i)\}$;
3   $Q^2_{can} \leftarrow \{x?_c y \mid x?_c y \in Q_{can}, y \in O_{\times}(Q_i)\}$;
    /* *Macro-ordering: consider Q^1_{can} before Q^2_{can}.*   */
4   **if** $Q^1_{can} \neq \varnothing$ **then**
5   **return** *micro-ordering*(Q^1_{can});
6   **else if** $Q^2_{can} \neq \varnothing$ **then**
7   **return** *micro-ordering*(Q^2_{can});

We denote the set of candidate questions by Q_{can}. Thus, $Q_{can} = \{\mathsf{x}?_c\mathsf{y} \mid rlt(\mathsf{x}?_c\mathsf{y}) \notin R^+(Q) \wedge \mathsf{x} \in O_? \wedge (\nexists c' \in C : \mathsf{x} \succ_{c'} \mathsf{y} \in R^+(Q))\}$. $\square$

If no candidate question exists, the question sequence Q is a terminal sequence. The reverse statement is also true, i.e., upon a terminal sequence, there is no candidate question left. This is formalized in the following property.

Property 2. $Q_{can} = \varnothing$ if and only if $O_? = \varnothing$.

Proof. The proof is omitted due to space limitations and can be found in the technical report [4]. $\square$

Questions violating the three conditions may also lead to terminal sequences. However, choosing only candidate questions matches our objective of quickly identifying non-Pareto optimal objects. Below we justify the conditions.

Condition (i): This is straightforward. If $R(Q)$ or its transitive closure already contains the outcome of $\mathsf{x}?_c\mathsf{y}$, we do not ask the same question again.

Condition (ii): This condition essentially dictates that at least one of the two objects in comparison is from $O_?$. (If only one of them belongs to $O_?$, we make it x.) Given a pair x and y, if neither is from $O_?$, there are three scenarios—(1) $\mathsf{x} \in O_\checkmark, \mathsf{y} \in O_\checkmark$, (2) $\mathsf{x} \in O_\checkmark, \mathsf{y} \in O_\times$ or $\mathsf{x} \in O_\times, \mathsf{y} \in O_\checkmark$, (3) $\mathsf{x} \in O_\times, \mathsf{y} \in O_\times$. Once we know an object is in $O_\checkmark$ or $O_\times$, its membership in such a set will never change. Hence, we are not interested in knowing the dominance relationship between objects from $O_\checkmark$ and $O_\times$ only. In all these three scenarios, comparing x and y is only useful for indirectly determining (by transitive closure) the outcome of comparing other objects. Intuitively speaking, such indirect pruning is not as efficient as direct pruning.

Condition (iii): This condition requires that, when $\mathsf{x}?_c\mathsf{y}$ is chosen, we cannot rule out the possibility of y dominating x. Otherwise, if y cannot possibly dominate x, the outcome of $\mathsf{x}?_c\mathsf{y}$ cannot help prune x. Note that, in such a case, comparing x and y by c may help prune y, if y still belongs to $O_?$ and x may dominate y. Such possibility is not neglected and is covered by a different representation of the same question—$\mathsf{y}?_c\mathsf{x}$, i.e., swapping the positions of x and y in checking the three conditions. If it is determined x and y cannot dominate each other, then their further comparison is only useful for indirectly determining the outcome of comparing other objects. Due to the same reason explained for condition (ii), such indirect pruning is less efficient.

The following simple Property 3 helps to determine whether $\mathsf{y} \succ \mathsf{x}$ is possible: If x is better than y by any criterion, then we can already rule out the possibility of $\mathsf{y} \succ \mathsf{x}$, without knowing the outcome of their comparison by every criterion. This allows us to skip further comparisons between them. Its correctness is straightforward based on the definition of object dominance.

Property 3 (Non-Dominance Property). At any given moment, suppose the set of asked questions is Q. Consider two objects x and y for which the comparison outcome is not known for every criterion, i.e., $\exists c$ such that $rlt(\mathsf{x}?_c\mathsf{y}) \notin R^+(Q)$. It can be determined that $\mathsf{y} \nsucc \mathsf{x}$ if $\exists c \in C$ such that $\mathsf{x} \succ_c \mathsf{y} \in R^+(Q)$. $\square$

In justifying the three conditions in defining candidate questions, we intuitively explained that indirect pruning is less efficient— if it is known that x does not belong to $O_?$ or y cannot possibly dominate x, we will not ask question $\mathsf{x}?_c\mathsf{y}$. We now justify this strategy theoretically and precisely. Consider a question sequence $Q = \langle q_1, \ldots, q_n \rangle$. We use $O_\checkmark(Q)$, $O_?(Q)$, $O_\times(Q)$ to denote object partitions according to $R^+(Q)$. For any question q_i, the subsequence comprised of its preceding questions is denoted $Q_{i-1} = \langle q_1, \ldots, q_{i-1} \rangle$. If q_i was not a candidate question when it was chosen (i.e., after $R(Q_{i-1})$ was obtained), we say it is a

non-candidate. The following Theorem 1 states that, if a question sequence contains non-candidate questions, we can replace it by a shorter or equally long sequence without non-candidate questions that produces the same set of dominated objects $O_\times$. Recall that the key to our framework is to recognize dominated objects and move them into $O_\times$ as early as possible. Hence, the new sequence will likely lead to less cost when the algorithm terminates.

Theorem 1. If Q contains non-candidate questions, there exists a question sequence Q' without non-candidate questions such that $|Q'| \leq |Q|$ and $O_\times(Q') = O_\times(Q)$.

Proof. We prove by demonstrating how to transform Q into such a Q'. Given any non-candidate question $q_i = \mathsf{x}?_c\mathsf{y}$ in Q, we remove it and, when necessary, replace several questions. The decisions and choices are partitioned into three mutually exclusive scenarios, which correspond to violations of the three conditions in Definition 3. The detailed proof is omitted due to space limitations and can be found in the technical report [4]. $\square$

3.2 Resolving Unusual Contradictions in Question Outcomes

A preference relation can be more accurately derived, if more input is collected from the crowd. However, under practical constraints on budget and time, the limited responses from the crowd (k answers per question) may present two types of contradicting preferences.

(i) Suppose $rlt(\mathsf{x}?_c\mathsf{y}) = \mathsf{x} \succ_c \mathsf{y}$ and $rlt(\mathsf{y}?_c\mathsf{z}) = \mathsf{y} \succ_c \mathsf{z}$ have been derived, i.e., they belong to $R(Q)$. They together imply $\mathsf{x} \succ_c \mathsf{z}$, since a preference relation must be transitive. Therefore the question $\mathsf{x}?_c\mathsf{z}$ will not be asked. If the crowd is nevertheless asked to further compare x and z, the result $rlt(\mathsf{x}?_c\mathsf{z})$ might be possibly $\mathsf{z} \succ_c \mathsf{x}$, which presents a contradiction.

(ii) Suppose $rlt(\mathsf{x}?_c\mathsf{y}) = \mathsf{x} \sim_c \mathsf{y}$ and $rlt(\mathsf{y}?_c\mathsf{z}) = \mathsf{y} \succ_c \mathsf{z}$ have been derived from the crowd. If the crowd is asked to further compare x and z, the result $rlt(\mathsf{x}?_c\mathsf{z})$ might be possibly $\mathsf{z} \succ_c \mathsf{x}$. The outcomes $\mathsf{y} \succ_c \mathsf{z}$ and $\mathsf{z} \succ_c \mathsf{x}$ together imply $\mathsf{y} \succ_c \mathsf{x}$, which contradicts with $\mathsf{x} \sim_c \mathsf{y}$. (A symmetric case is $rlt(\mathsf{x}?_c\mathsf{y}) = \mathsf{x} \sim_c \mathsf{y}$, $rlt(\mathsf{y}?_c\mathsf{z}) = \mathsf{z} \succ_c \mathsf{y}$, and the crowd might respond with $rlt(\mathsf{x}?_c\mathsf{z}) = \mathsf{x} \succ_c \mathsf{z}$, which also leads to contradiction with $\mathsf{x} \sim_c \mathsf{y}$. The following discussion applies to this symmetric case, which is thus not mentioned again.)

In practice, such contradictions are uncommon. This is easy to understand intuitively—as long as the underlying preference relation is transitive, collective wisdom of the crowds will reflect it. We can find evidence of it in [27, 10], which confirmed that preference judgments of relevance in document retrieval are transitive.

Nevertheless, contradictions still occur. Type (i) contradictions can be prevented by enforcing the following simple Rule 1 to assume transitivity and thus skip certain questions. They will never get into the derived preference relations. In fact, in calculating transitive closure (Definition 1) and defining candidate questions (Sec.3.1), we already apply this rule.

Rule 1 (Contradiction Prevention by Skipping Questions). Given objects x, y, z and a criterion c, if $rlt(\mathsf{x}?_c\mathsf{y}) = \mathsf{x} \succ_c \mathsf{y}$ and $rlt(\mathsf{y}?_c\mathsf{z}) = \mathsf{y} \succ_c \mathsf{z}$, we assume $rlt(\mathsf{x}?_c\mathsf{z}) = \mathsf{x} \succ_c \mathsf{z}$ and thus will not ask the crowd to further compare x and z by criterion c. $\square$

To resolve type (ii) contradictions, we enforce the following simple Rule 2.

Rule 2 (Contradiction Resolution by Choosing Outcomes). Consider objects x, y, z and a criterion c. Suppose $rlt(\mathsf{x}?_c\mathsf{y}) = \mathsf{x} \sim_c \mathsf{y}$ and $rlt(\mathsf{y}?_c\mathsf{z}) = \mathsf{y} \succ_c \mathsf{z}$ are obtained from the crowd. If $rlt(\mathsf{x}?_c\mathsf{z}) = \mathsf{z} \succ_c \mathsf{x}$ is obtained from the crowd afterwards, we replace the outcome of this question by $\mathsf{x} \sim_c \mathsf{z}$. (Note that we do not replace it by $\mathsf{x} \succ_c \mathsf{z}$, since $\mathsf{z} \succ_c \mathsf{x}$ is closer to $\mathsf{x} \sim_c \mathsf{z}$.) $\square$

i	$rlt(q_i)$	Derived Results	$O_\checkmark$	$O_?$	$O_\times$
1-9	$\underline{b}\succ_m\underline{e}.\underline{c}\sim_a\underline{d}.\underline{a}\sim_m c$ $\underline{c}\succ_s e.\underline{b}\sim_s d.\underline{b}\succ_a \underline{a}$ $\underline{d}\succ_m e.\underline{b}\sim_m d.\underline{b}\succ_s f$		∅	{a,b,c,d,e,f}	∅
10	$\underline{a}\succ_m\underline{d}$	$a\succ_m e$			
11	$\underline{c}\succ_a\underline{a}$				
12	$\underline{b}\sim_s c$				
13	$\underline{c}\succ_m\underline{d}$	$c\succ_m e$			
14-19	$\underline{d}\succ_s\underline{e}.\underline{e}\sim_a c.\underline{d}\sim_a f$ $\underline{a}\sim_a d.f\succ_a\underline{a}.\underline{b}\succ_a e$				
20	$\underline{b}\succ_a\underline{d}$	$b\succ d$	∅	{a,b,c,e,f}	{d}
21-23	$\underline{c}\succ_s f.\underline{a}\succ_s\underline{e}.f\sim_m b$				
24	$\underline{a}\succ_s\underline{f}$	$a\sim f$			
25	$\underline{e}\succ_a f$	$b\succ_a f.b\succ_a a$ $e\succ_a a.b\succ f$	∅	{a,b,c,e}	{d,f}
26	$\underline{b}\succ_a c$				
27	$\underline{a}\succ_m\underline{b}$				
28	$\underline{b}\succ_s\underline{e}$	$b\succ e$	∅	{a,b,c}	{d,e,f}
29	$\underline{c}\succ_s\underline{a}$	$c\succ_s e.c\succ a$	∅	{b,c}	{a,d,e,f}
30	$\underline{b}\sim_m c$	$b\succ c$	{b}	∅	{a,c,d,e,f}

Table 2: RandomQ on Example 1.

i	$rlt(q_i)$	Derived Results	$O_\checkmark$	$O_?$	$O_\times$
1	$\underline{c}\succ_s f$		∅	{a,b,c,d,e,f}	∅
2	$\underline{f}\succ_m c$	$f\sim c$			
3 – 4	$\underline{a}\succ_s\underline{e}.\underline{a}\succ_m e$				
5	$\underline{e}\succ_a a$	$a\sim e$			
6 – 7	$\underline{c}\succ_s\underline{e}.\underline{c}\succ_m\underline{e}$				
8	$\underline{e}\succ_a c$	$c\succ e$	∅	{a,b,c,d,f}	{e}
9	$\underline{b}\succ_s a$	$b\succ_s e$			
10	$\underline{a}\succ_m b$	$a\sim b$			
11	$\underline{d}\succ_s f$				
12	$\underline{f}\succ_m d$	$f\sim d$			
13	$\underline{d}\succ_s a$	$d\succ_s e$			
14	$\underline{a}\succ_m d$	$a\sim d$			
15 – 16	$\underline{b}\sim_s c.\underline{b}\sim_m c$				
17	$\underline{b}\succ_a c$	$b\succ c$	∅	{a,b,d,f}	{c,e}
18 – 19	$\underline{d}\sim_s b.\underline{d}\sim_m b$				
20	$\underline{b}\succ_a d$	$b\succ d$	∅	{a,b,f}	{c,d,e}
21	$\underline{a}\succ_s f$	$b\succ_s f$			
22	$\underline{a}\sim_m f$				
23	$\underline{f}\succ_a a$	$a\sim f$			
24	$\underline{b}\sim_m f$				
25	$\underline{b}\succ_a f$	$b\succ_a a.b\succ f$	{b}	{a}	{c,d,e,f}
26 – 27	$\underline{c}\succ_s\underline{a}.\underline{a}\sim_m c$				
28	$\underline{c}\succ_a\underline{a}$	$c\succ a$	{b}	∅	{a,c,d,e,f}

Table 3: RandomP on Example 1.

4. MICRO-ORDERING IN QUESTION SELECTION

At every iteration of Alg.1, we choose a question $x?_c y$ from the set of candidate questions. By macro-ordering, when available, the question selection strategy (Alg.2) chooses a candidate question in which $y \notin O_\times$, i.e., it chooses from Q_{can}^1. Otherwise, it chooses from Q_{can}^2. The size of Q_{can}^1 and Q_{can}^2 can be large. Micro-ordering is for choosing from the many candidates. As discussed in Sec.3, in order to find a short question sequence, the overriding principle of our question selection strategy is to identify non-Pareto optimal objects as early as possible. Guided by this principle, this section discusses several micro-ordering strategies. Since the strategies are the same for Q_{can}^1 and Q_{can}^2, we will simply use the term "candidate questions" without distinction between Q_{can}^1 and Q_{can}^2.

4.1 Random Question (RandomQ)

RandomQ, as its name suggests, simply selects a random candidate question. Table 2 shows an execution of the general framework under RandomQ for Example 1. For each iteration i, the table shows the question outcome $rlt(q_i)$. Following the question form $x?_c y$ in Definition 3, the object "x" in a question is underlined when we present the question outcome. The column "derived results" displays derived question outcomes by transitive closure (e.g., $a\succ_m e$ based on $rlt(q_7)=d\succ_m e$ and $rlt(q_{10})=a\succ_m d$) and derived object dominance (e.g., $b\succ d$ after q_{20}). The table also shows the object partitions ($O_\checkmark$, $O_?$ and $O_\times$) when the execution starts and when the partitions are changed after an iteration. Multiple iterations may be presented together if other columns are the same for them.

As Table 2 shows, this particular execution under RandomQ requires 30 questions. When the execution terminates, it finds the only Pareto-optimal object b. This simplest micro-ordering strategy already avoids many questions in the brute-force approach. The example clearly demonstrates the benefits of choosing candidate questions only and applying macro-strategy.

4.2 Random Pair (RandomP)

RandomP randomly selects a pair of objects x and y and keeps asking questions to compare them ($x?_c y$ or $y?_c x$) until no such candidate question remains, upon which it randomly picks another pair of objects. This strategy echoes our principle of eagerly identifying non-Pareto optimal objects. To declare an object x non-Pareto optimal, we must identify another object y such that y dominates x. If we directly compare x and y, it requires comparing them by every

criterion in C in order to make sure $y\succ x$. By skipping questions according to transitive closure, we do not need to directly compare them by every criterion. However, Property 4 below states that we still need at least $|C|$ questions involving x—some are direct comparisons with y, others are comparisons with other objects which indirectly lead to outcomes of comparisons with y. When there is a candidate question $x?_c y$, it means y may dominate x. The fewer criteria remain for comparing them, the more likely y will dominate x. Hence, by keeping comparing the same object pair, RandomP aims at finding more non-Pareto objects by less questions.

Property 4. Given a set of criteria C and an object $x\in O$, at least $|C|$ pairwise comparison questions involving x are required in order to find another object y such that $y\succ x$.

Proof. The proof is omitted due to space limitations and can be found in the technical report [4]. □

Table 3 illustrates an execution of RandomP for Example 1. The initial two questions are between c and f. Afterwards, it is concluded that $c\sim f$ by Property 3. Therefore, RandomP moves on to ask 3 questions between a and e. In total, the execution requires 28 questions. Although it is shorter than Table 2 by only 2 questions due to the small size of the example, it clearly moves objects into $O_\times$ more quickly. (In Table 2, $O_\times$ is empty until the 20th question. In Table 3, $O_\times$ already has 3 objects after 20 questions.) The experiment results in Sec.5 exhibit significant performance gain of RandomP over RandomQ on larger data.

4.3 Pair with Fewest Remaining Questions (FRQ)

Similar to RandomP, once a pair of objects x and y are chosen, FRQ keeps asking questions between x and y until there is no such candidate questions. Different from RandomP, instead of randomly picking a pair of objects, FRQ always chooses a pair with the fewest remaining questions. There may be multiple such pairs. To break ties, FRQ chooses such a pair that x has dominated the fewest other objects and y has dominated the most other objects. Furthermore, in comparing x and y, FRQ orders their remaining questions (and thus criteria) by how likely x is worse than y on the criteria. Below we explain this strategy in more detail.

Selecting Object Pair Consider a question sequence Q_i so far and FRQ is to select the next question Q_{i+1}. We use $C_{x,y}$ to

denote the set of criteria c such that $x?_c y$ is a candidate question, i.e., $C_{x,y}=\{c\in C \mid x?_c y\in Q^1_{can}\}$. (We assume Q^1_{can} is not empty. Otherwise, FRQ chooses from Q^2_{can} in the same way; cf. Alg.2.) By Definition 3, the outcomes of these questions are unknown, i.e., $\forall c\in C_{x,y} : rlt(x?_c y)\notin R^+(Q_i)$. Furthermore, if any remaining question (whose outcome is unknown) between x and y is a candidate question, then all remaining questions between them are candidate questions. FRQ chooses a pair with the fewest remaining candidate questions, i.e., a pair belonging to $S_1=\arg\min_{(x,y)} |C_{x,y}|$.

The reason to choose such a pair is intuitive. It requires at least $|C_{x,y}|$ candidate questions to determine $y\succ x$. (The proof would be similar to that of Property 4.) Therefore, $\min_{(x,y)} |C_{x,y}|$ is the minimum number of candidate questions to further ask, in order to determine that an object is dominated, i.e., non-Pareto optimal. Thus, a pair in S_1 may lead to a dominated object by the fewest questions, matching our goal of identifying non-Pareto optimal objects as soon as possible.

We further justify this strategy in a probabilistic sense. For $y\succ x$ to be realized, it is necessary that none of the remaining questions has an outcome $x\succ_c y$, i.e., $\forall c\in C_{x,y} : rlt(x?_c y) \neq x\succ_c y$. Make the simplistic assumption that every question $x?_c y$ has an equal probability p of not having outcome $x\succ_c y$, i.e., $\forall x?_c y\in Q^1_{can}$, $P(rlt(x?_c y)\neq x\succ_c y)=p$. Further assuming independence of question outcomes, the probability of satisfying the aforementioned necessary condition is $p^{|C_{x,y}|}$. By taking a pair belonging to S_1, we have the largest probability of finding a dominated object. We note that, for $y\succ x$ to be realized, in addition to the above necessary condition, another condition must be satisfied—if $\nexists c$ such that $y\succ_c x \in R^+(Q_i)$, the outcome of at least one remaining question should be $y\succ_c x$, i.e., $\exists c\in C_{x,y} : rlt(x?_c y)=y\succ_c x$. Our informal probability-based analysis does not consider this extra requirement.

Breaking Ties There can be multiple object pairs with the fewest remaining questions, i.e., $|S_1|>1$. To break ties, FRQ chooses such an x that has dominated the fewest other objects, since it is more likely to be dominated. If there are still ties, FRQ further chooses such a y that has dominated the most other objects, since it is more likely to dominate x. More formally, FRQ chooses a pair belonging to $S_2=\{(x,y)\in S_1 \mid \nexists(x',y')\in S_1$ such that $d(x')>d(x)\vee (d(x')= d(x) \wedge d(y')>d(y))\}$, where the function $d(\cdot)$ returns the number of objects so far dominated by an object, i.e., $\forall x, d(x) = |\{y|x\succ y$ based on $R^+(Q_i)\}|$. This heuristic follows the principle of detecting non-Pareto optimal objects as early as possible. Note that S_2 may still contain multiple object pairs. In such a case, FRQ chooses an arbitrary pair.

Selecting Comparison Criterion Once a pair (x,y) is chosen, FRQ has to select a criterion for the next question. FRQ orders the remaining criteria $C_{x,y}$ based on the heuristic that the sooner it understands $y\succ x$ will not happen, the lower cost it pays. As discussed before, $|C_{x,y}|$ questions are required in order to conclude that $y\succ x$; on the other hand, only one question (if asked first) can be enough for ruling it out. Consider the case that x is better than y by only one remaining criterion, i.e., $\exists c\in C_{x,y} : rlt(x?_c y)=x\succ_c y$ and $\forall c'\in C_{x,y}, c'\neq c : rlt(x?_{c'} y)=x\not\succ_{c'} y$. If FRQ asks $x?_c y$ after all other remaining questions, it takes $|C_{x,y}|$ questions to understand y does not dominate x; but if $x?_c y$ is asked first, no more questions are necessary, because there will be no more candidate questions in the form of $x?_c y$.

Therefore, FRQ orders the criteria $C_{x,y}$ by a scoring function that reflects the likelihood of x's superiority than y by the corresponding criteria. More specifically, for each $c\in C_{x,y}$, its score is $r_c(x,y) = r_c(y)+r'_c(y)-r''_c(y)-(r_c(x)+r'_c(x)-r''_c(x))$ where $r_c(y) = |\{z \mid z\succ_c y\in R^+(Q_i)\}|$, $r'_c(y)=|\{z \mid y\sim_c z\in R^+(Q_i)\}|$,

i	$rlt(q_i)$	Derived Results	$(x,y), C_{x,y}$	$O_\checkmark$	$O_?$	$O_\times$
			$(a,b),_{\{s,m,a\}}$	$\varnothing$	$\{a,b,c,d,e,f\}$	$\varnothing$
1	$b\succ_s a$		$(a,b),_{\{m,a\}}$			
2	$a\succ_m b$	$a\sim b$	$(a,c),_{\{s,a,m\}}$			
3	$c\succ_s a$		$(a,c),_{\{a,m\}}$			
4	$c\sim_a a$		$(a,c),_{\{m\}}$			
5	$c\succ_m a$	$c\succ a$	$(b,c),_{\{a,s,m\}}$	$\varnothing$	$\{b,c,d,e,f\}$	$\{a\}$
6	$b\sim_a c$		$(b,c),_{\{s,m\}}$			
7	$b\sim_s c$		$(b,c),_{\{m\}}$			
8	$b\succ_m c$	$b\succ_m a. b\succ c$	$(d,b),_{\{a,s,m\}}$	$\varnothing$	$\{b,d,e,f\}$	$\{a,c\}$
9	$b\sim_a d$		$(d,b),_{\{s,m\}}$			
10	$b\sim_s d$		$(d,b),_{\{m\}}$			
11	$b\succ_m d$	$b\succ d$	$(e,b),_{\{a,s,m\}}$	$\varnothing$	$\{b,e,f\}$	$\{a,c,d\}$
12	$b\sim_a e$		$(e,b),_{\{s,m\}}$			
13	$b\sim_s e$		$(e,b),_{\{m\}}$			
14	$b\succ_m e$	$a\succ_m e. b\succ e$	$(f,b),_{\{a,s,m\}}$	$\varnothing$	$\{b,f\}$	$\{a,c,d,e\}$
15	$b\sim_a f$		$(f,b),_{\{s,m\}}$			
16	$b\sim_s f$		$(f,b),_{\{m\}}$			
17	$b\succ_m f$	$b\succ f$		$\{b\}$	$\varnothing$	$\{a,c,d,e,f\}$

Table 4: FRQ on Example 1.

and $r''_c(y)=|\{z \mid y\succ_c z\in R^+(Q_i)\}|$. In this scoring function, $r_c(y)$ is the number of objects preferred over y by criterion c, $r'_c(y)$ is the number of objects equally good (or bad) as y by c, and $r''_c(y)$ is the number of objects to which y is preferred with regard to c. FRQ asks the remaining questions in decreasing order of the corresponding criteria's scores. This way, it may find such a question that $rlt(x?_c y)=x\succ_c y$ earlier than later.

Table 4 presents the framework's execution for Example 1, by applying the FRQ policy. In addition to the same columns in Tables 2 and 3, Table 4 also includes an extra column to show, at each iteration, the chosen object pair for the next question (x,y) and the set of remaining comparison criteria between them $(C_{x,y})$. The criteria in $C_{x,y}$ are ordered by the aforementioned ranking function $r(\cdot)$. At the beginning of the execution, the object pair is arbitrarily chosen and the criteria are arbitrarily ordered. In the example, we assume $a?_s b$ is chosen as the first question. After q_2, FRQ can derive that $a\sim b$. Hence, there is no more candidate question between them and FRQ chooses the next pair (a,c). Three questions are asked for comparing them. At the end of q_5, multiple object pairs have the fewest remaining questions. By breaking ties, (b,c) is chosen as the next pair, since only c has dominated any object so far. The remaining criteria $C_{b,c}$ are ordered as $\{a, s, m\}$, because $r_a(b,c)>r_s(b,c)$ and $r_a(b,c)>r_m(b,c)$. The execution sequence terminates after 17 questions, much shorter than the 30 and 28 questions by RandomQ and RandomP, respectively.

To conclude the discussion on micro-ordering, we derive a lower bound on the number of questions required for finding all Pareto-optimal objects (Theorem 2). The experiment results in Sec.5 reveal that FRQ is nearly optimal and the lower bound is practically tight, since the number of questions used by FRQ is very close to the lower bound.

Theorem 2. Given objects O and criteria C, to find all Pareto-optimal objects in O, at least $(|O|-k)\times|C|+(k-1)\times 2$ pairwise comparison questions are necessary, where k is the number of Pareto-optimal objects in O.

Proof. The proof is omitted due to space limitations and can be found in the technical report [4]. $\square$

5. EXPERIMENTS

We designed and conducted experiments to compare the efficiency of different instantiations of the general framework under varying problem sizes. Our experiments used both a real crowdsourcing marketplace and simulations based on a real dataset.
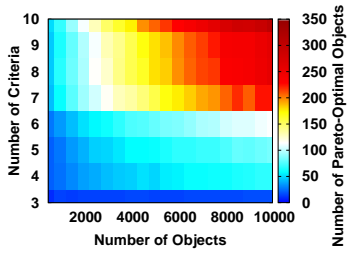
Figure 5: No. of Pareto-optimal objects. Varying $|O|$ and $|C|$.

(a) $|C| = 4$, varying $|O|$

(b) $|C| = 10$, varying $|O|$

(c) $|O| = 3,000$, varying $|C|$

(d) $|O| = 10,000$, varying $|C|$

Figure 6: No. of questions by BruteForce and basic methods.

5.1 Efficiency and Scalability

We studied the efficiency and scalability of various instantiations of the general framework. Given the large number of questions required for such a study, we cannot afford using a real crowdsourcing marketplace. Hence, we performed the following simulation. Each object is an NBA player in a particular year. The objects are compared by 10 criteria, i.e., performance categories such as *points*, *rebounds*, *assists*, etc. We simulated the corresponding 10 preference relations based on the players' real performance in individual years, as follows. Consider a performance category c and two objects x=(player1, year1) and y=(player2, year2). $x.c$ is player1's per-game performance on category c in year1 (similarly for $y.c$). Values in each category c are normalized into the range $[0, 1]$, where 0 and 1 correspond to the minimal and maximal values in c, respectively. Suppose $x.c > y.c$. We generated a uniform random number v in $[0, 1]$. If $v < 1 - e^{-(x.c-y.c)}$, we set $x \succ_c y$, otherwise we set $x \sim_c y$. This way, we introduced a perturbation into the preference relations in order to make sure they are partial orders, as opposed to directly using real performance statistics (which would imply total orders). Fig.5 shows that the number of Pareto-optimal objects increases by the sizes of both object set O (objects are randomly selected) and criteria set C (the first $|C|$ criteria of the aforementioned 10 criteria).

Effectiveness of candidate questions and macro-ordering
To verify the effectiveness of candidate questions and macro-ordering, we compared five methods—BruteForce, −CQ−MO, −CQ+MO, +CQ−MO, and +CQ+MO. The notation +/− before CQ and MO indicates whether a method only selects candidate questions (CQ)

(a) $|C| = 4$, varying $|O|$

(b) $|C| = 10$, varying $|O|$

(c) $|O| = 3,000$, varying $|C|$

(d) $|O| = 10,000$, varying $|C|$

Figure 7: No. of questions by different micro-ordering heuristics.

and whether it applies the macro-ordering strategy (MO), respectively. In all these five methods, qualifying questions are randomly selected, i.e., no particular micro-ordering heuristics are applied. For instance, +CQ+MO selects only candidate questions and applies macro-ordering. Hence, it is equivalent to RandomQ. Fig.6 shows the numbers of required pairwise comparisons (in logarithmic scale) for each method, varying by object set size ($|O|$ from 500 to 10,000 for $|C|$=4 and $|C|$=10) and criterion set size ($|C|$ from 3 to 10 for $|O|$=3,000 and $|O|$=10,000). The figure clearly demonstrates the effectiveness of both CQ and MO, as taking out either feature leads to significantly worse performance than RandomQ. Particularly, the gap between +CQ−MO and −CQ+MO suggests that choosing only candidate questions has more fundamental impact than macro-ordering. If neither is applied (i.e., −CQ−MO), the performance is equally poor as that of BruteForce. (−CQ−MO uses slightly less questions than BruteForce, since it can terminate before exhausting all questions. However, the difference is negligible for practical purpose, as their curves overlap under logarithmic scale.)

Effectiveness of micro-ordering
Fig.7 presents the numbers of pairwise comparisons required by different micro-ordering heuristics (RandomQ, i.e., +CQ+MO, RandomP, FRQ) and LowerBound (cf. Theorem 2) under varying sizes of the object set ($|O|$ from 500 to 10,000 for $|C|$=4 and $|C|$=10) and the criteria set ($|C|$ from 3 to 10 for $|O|$=3,000 and $|O|$=10,000). In all these instantiations of the general framework, CQ and MO are applied. The results are averaged across 30 executions. All these methods outperformed BruteForce by orders of magnitude. (BruteForce is not shown in Fig.7 since it is off scale, but its number can be calculated by equation $|C| \times |O| \times (|O| - 1)/2$.) For instance, for 5,000 objects and 4 criteria, the ratio of pairwise comparisons required by even the naive RandomQ to that used by BruteForce is already as low as 0.0048. This clearly shows the effectiveness of CQ and MO, as discussed for Fig.6. The ratios for RandomP and FRQ are further several times smaller (0.00094 and 0.00048, respectively). The big gain by FRQ justifies the strategy of choosing object pairs with the fewest remaining questions. Especially, FRQ has nearly optimal performance, because it gets very close to LowerBound in Fig.7. The small
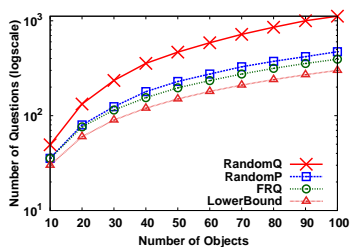
Figure 8: No. of questions by different micro-ordering heuristics. $|C| = 3$, varying $|O|$.

gap between FRQ and LowerBound also indicates that the lower bound is practically tight. The figure further suggests excellent scalability of FRQ as its number of questions grows almost linearly by both $|C|$ and $|O|$.

5.2 Experiments Using a Real Crowdsourcing Marketplace

We also studied the performance of the proposed algorithms using the popular crowdsourcing marketplace Amazon Mechanical Turk (AMT). The task is to compare 100 photos of our institution with regard to *color*, *sharpness* and *landscape*. To obtain the ground-truth data, all $14,850$ possible pairwise questions were partitioned into $1,650$ tasks, each containing 9 questions on a criterion. An AMT crowdsourcer is allowed to perform a task only if they have responded to at least 100 HITs (Human Intelligence Tasks) before with at least 90% approval rate. Furthermore, we implemented basic quality control by including 2 additional validation questions in each task that expect certain answers. For instance, one such question asks the crowd to compare a colorful photo and a dull photo by criterion *color*. A crowdsourcer's responses in a task are discarded if their response to a validation question deviates from our expectation. (236 crowdsourcers failed on this.) The parameters in Equation (1) were set to be $k=5$ and $\theta=0.6$. Hence, in total $(1,650\times 5+236)\times(9+2) = 93,346$ pairwise comparisons were performed by AMT crowdsourcers. We paid 1 cent for each comparison and therefore spent close to $1,000 in total.

The responses to all possible questions provide the ground-truth data. An algorithm execution only needs the responses to a subset of the questions. We randomly selected a subset of photos O ($|O|$ from 10 to 100) and applied various algorithms to find Pareto-optimal photos. Figure 8 shows, for varying $|O|$, the number of questions (in logarithmic scale) required by each micro-ordering strategy. To account for the randomness in RandomP and RandomQ, we repeated these two algorithms, respectively, 30 times, and we reported the average numbers of questions. Confirming the results in Figure 7, FRQ was close to the theoretical lower bound, performing better than the other two methods, and RandomP outperformed RandomQ.

6. CONCLUSIONS

This is the first study on using crowdsourcing to find Pareto-optimal objects when objects do not have explicit attributes and preference relations are strict partial orders. The partial orders are obtained by pairwise comparison questions to the crowd. It introduces an iterative question-selection framework that is instantiated into different methods by exploiting the ideas of candidate questions, macro-ordering and micro-ordering. Experiment were conducted by simulations on large object sets and by using a real crowdsourcing marketplace. The results exhibited not only orders of magnitude reductions in questions against a brute-force approach, but also close-to-optimal performance from the most efficient method.

7. REFERENCES

[1] N. Ailon. Active learning ranking from pairwise preferences with almost optimal query complexity. In *NIPS*, pages 810–818, 2011.

[2] N. Ailon. An active learning algorithm for ranking from pairwise preferences with an almost optimal query complexity. *Journal of Machine Learning Research*, 13(1):137–164, Jan. 2012.

[3] K. J. Arrow. *Social choice and individual values*. Yale University Press, 1951.

[4] A. Asudeh, G. Zhang, N. Hassan, C. Li, and G. V. Zaruba. Crowdsourcing Pareto-Optimal Object Finding by Pairwise Comparisons. *ArXiv e-prints 1409.4161*, Sept. 2014.

[5] S. Borzsony, D. Kossmann, and K. Stocker. The skyline operator. In *ICDE*, pages 421–430, 2001.

[6] M. Braverman and E. Mossel. Noisy sorting without resampling. In *SODA*, pages 268–276, 2008.

[7] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *ICML*, pages 89–96, 2005.

[8] C. J. C. Burges, R. Ragno, and Q. V. Le. Learning to Rank with Nonsmooth Cost Functions. In *NIPS*, pages 193–200, 2006.

[9] Y. Cao, J. Xu, T.-Y. Liu, H. Li, Y. Huang, and H.-W. Hon. Adapting ranking SVM to document retrieval. In *SIGIR*, pages 186–193, 2006.

[10] B. Carterette, P. N. Bennett, D. M. Chickering, and S. T. Dumais. Here or there: Preference judgments for relevance. In *ECIR*, 2008.

[11] C.-Y. Chan, P.-K. Eng, and K.-L. Tan. Stratified computation of skylines with partially-ordered domains. In *SIGMOD*, 2005.

[12] X. Chen, P. N. Bennett, K. Collins-Thompson, and E. Horvitz. Pairwise ranking aggregation in a crowdsourced setting. In *WSDM*, pages 193–202, 2013.

[13] J. Chomicki. Preference formulas in relational queries. *TODS*, 2003.

[14] M. d. Condorcet. Essay on the application of analysis to the probability of majority decisions. *Paris: Imprimerie Royale*, 1785.

[15] 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.

[16] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *CACM*, 1992.

[17] B. Groz and T. Milo. Skyline queries with noisy comparisons. In *PODS*, pages 185–198, 2015.

[18] P. G. Ipeirotis, F. Provost, and J. Wang. Quality management on amazon mechanical turk. In *HCOMP*, pages 64–67, 2010.

[19] K. G. Jamieson and R. D. Nowak. Active ranking using pairwise comparisons. In *NIPS*, pages 2240–2248, 2011.

[20] W. Kießling. Foundations of preferences in database systems. In *VLDB*, pages 311–322, 2002.

[21] N. N. Liu, M. Zhao, and Q. Yang. Probabilistic latent preference analysis for collaborative filtering. In *CIKM*, pages 759–766, 2009.

[22] T.-Y. Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, Mar. 2009.

[23] C. Lofi, K. El Maarry, and W.-T. Balke. Skyline queries in crowd-enabled databases. In *EDBT*, pages 465–476, 2013.

[24] S. Negahban, S. Oh, and D. Shah. Iterative ranking from pair-wise comparisons. In *NIPS*, pages 2483–2491, 2012.

[25] V. Polychronopoulos, L. de Alfaro, J. Davis, H. Garcia-Molina, and N. Polyzotis. Human-powered top-k lists. In *WebDB*, 2013.

[26] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*, pages 452–461, 2009.

[27] M. E. Rorvig. The simple scalability of documents. *Journal of the American Society for Information Science*, 41(8):590–598, 1990.

[28] D. Sacharidis, S. Papadopoulos, and D. Papadias. Topologically sorted skylines for partially ordered domains. In *ICDE*, 2009.

[29] N. Sarkas, G. Das, N. Koudas, and A. K. Tung. Categorical skylines for streaming data. In *SIGMOD*, pages 239–250, 2008.

[30] L. L. Thurstone. A law of comparative judgment. *Psychological Review*, 34:273–286, 1927.

[31] J. Yi, R. Jin, S. Jain, and A. K. Jain. Inferring users' preferences from crowdsourced pairwise comparisons: A matrix completion approach. In *HCOMP*, 2013.

[32] S. Zhang, N. Mamoulis, D. W. Cheung, and B. Kao. Efficient skyline evaluation over partially ordered domains. In *VLDB*, 2010.