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
University of Texas at Arlington

CIKM, October 21, 2015, Melbourne, Australia

Humans' Obsession: Comparing Things

Who is the better footballer?

http://visual.ly/messi-vs-ronaldo

Which is the better company to work for?

Facebook vs. Google

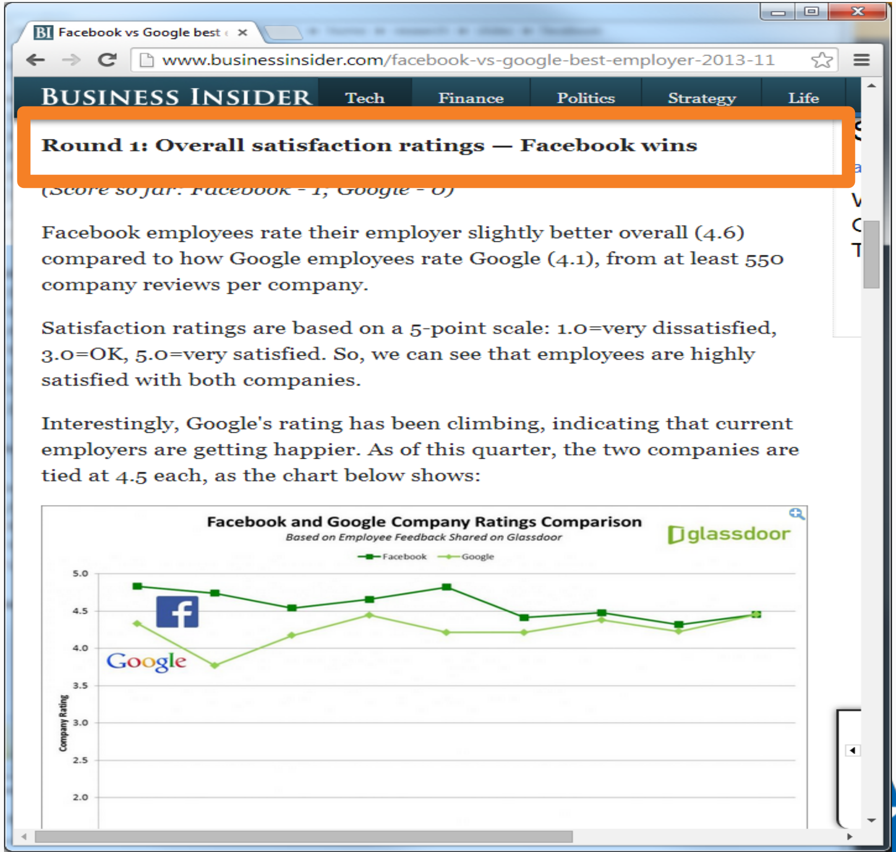

http://www.businessinsider.com/facebook-vs-google-best-employer-2013-11

Facebook vs. Google

Multiple criteria

Overall satisfaction	CEO approval	Employee confidence in the future	Perks and salaries	Interview difficulty
f ≻ g	f ≻ g	f ≻ g	g ≻ f	f ≻ g

Which one is better?

Overview

1. Ask crowd to compare objects on individual criteria
2. Derive partial knowledge about preference relations based on responses from crowd
3. Find all Pareto-optimal objects without exhausting all possible comparison questions

General framework and its instantiations (algorithms), with the goal of minimizing number of questions

Preference (Better-Than) Relation P_c

$(x, y) \in P_c$ (or $x \succ_c y$)

“x is better than (preferred over) y with regard to criterion c”

Assumptions on data model:

- P_c is a strict partial order (as opposed to a total order)
  Irreflexivity, Transitivity, Asymmetry
- No explicit attribute representation,
  thus no equivalence on a criterion

Notations:

- (indifferent) $x \sim_c y \Leftrightarrow (x, y) \notin P_c \wedge (y, x) \notin P_c$
- (not better than) $x \nsucc_c y \Leftrightarrow (x, y) \notin P_c$

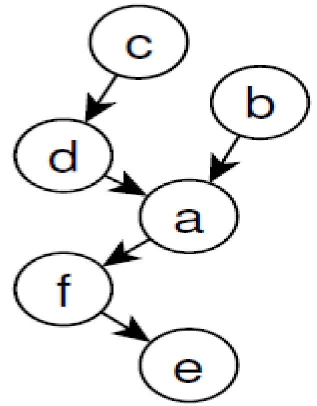

$b \succ_s a$
$c \succ_s a$
$d \succ_s a$
$a \succ_s e$
$a \succ_s f$
$b \sim_s c$
$b \sim_s d$
$b \succ_s e$
$b \succ_s f$
$c \succ_s d$
$c \succ_s e$
$c \succ_s f$
$d \succ_s e$
$d \succ_s f$
$f \succ_s e$

Pareto-optimal Objects

Object dominance: consider objects O, criteria C

$y \succ x \iff \forall c \in C : x \nsucc_c y$ and $\exists c \in C$ such that $y \succ_c x$ (i.e., x is not better than y by any criterion and y is better than x by at least one criterion)

$x \in O$ is Pareto-optimal $\iff$ x is not dominated by any other object

Example:

- $c \succ d \Leftarrow c \succ_{story} d,\ c \succ_{music} d,\ c \sim_{acting} d$
- Only one Pareto-optimal object: b

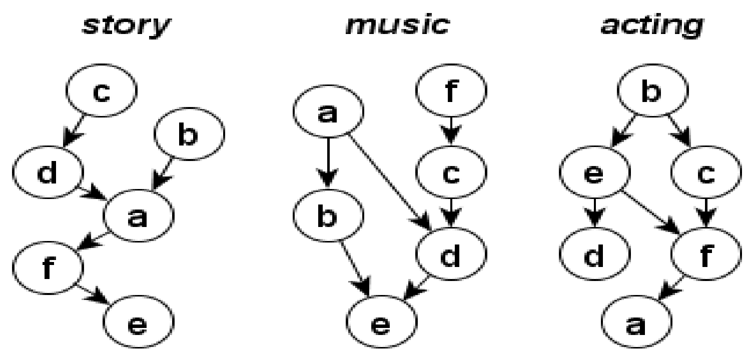

Deriving Preference Relations by Aggregating Crowd's Responses to Pairwise Comparisons

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

In each following pair(each row), which picture has more interesting colors?

Left Picture	Right Picture	Preference
		Left Picture Right Picture No Preference

Pareto-Optimal Object Finding

Problem statement: Given objects O and criteria C, find all Pareto-optimal objects, using pairwise comparisons by individual criteria

Cost metric

- Goal: as few pairwise comparison questions as possible
- Simple, but reflect real-world monetary cost and time delay
- Brute-force approach : $|C| \times |O| \times (|O|-1)/2$ questions

Assumptions on execution model

- Sequential execution: get $rlt(q_i)$ before asking q_{i+1}
- No consideration of worker quality

Applications

Collecting Public Opinion

- Best companies to work for, best cities to live in

Group Decision Making

- Where for lunch, which product to use, which candidate to hire

Information Exploration

- Compare photos by color, sharpness, and landscape

Back to the "which one is better"?

- After finding Pareto-optimal objects, further actions (ranking, filtering, visualization) to find desirable objects

Related Work

	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

[12] Chen et al. WSDM13 || [15] Davidson et al. ICDT13 || [17] Grozet al. PODS15 || [23] Lofi et al. EDBT13 || [25] Polychronopoulos et al. WebDB13 || [4] technical report of this paper

Other related work: collaborative filtering, learning to rank, …

Related Work

Explicit attribute representation

- Total order on ordinal attributes (sizes of houses, ratings of restaurants), Partial-order on categorical attributes (genres of movies)
- Not always easy to model and/or for users to provide missing values (e.g., story of movies)

Pairwise comparisons

- Known to be easier, faster, and less error-prone. Widely used in social choice and welfare, preferences, and voting

Partial order vs. Total order

- A direct effect of using pairwise comparisons
- Not always natural to enforce a total order

General, Iterative Algorithm Framework

4-steps in each iteration

(1) Question selection
(2) Outcome derivation
(3) Contradiction resolution
(4) Termination test

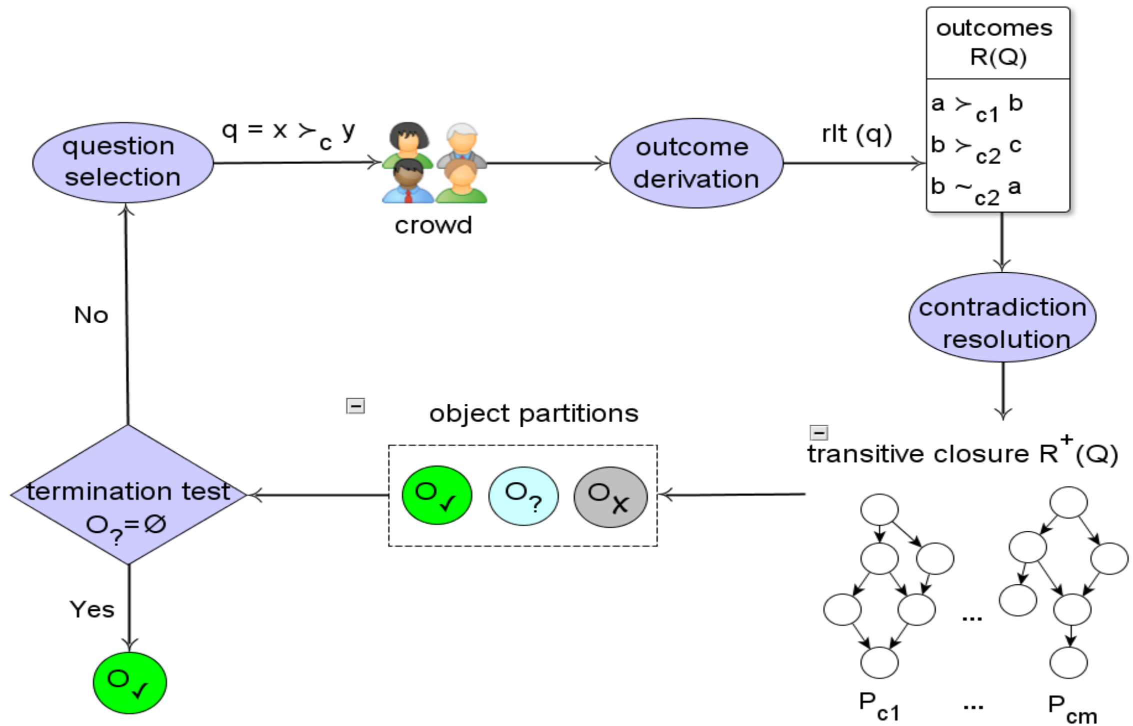

(2) Outcome Derivation

The framework is agnostic to the outcome derivation method.

Other conceivable method can be plugged in.

Question

- $q = x \, ?_c \, y$ (compare x and y by criterion c)

Three possible outcomes based on voting

$$\mathrm{rlt}(q) = \begin{cases} x \succ_c y \\ y \succ_c x \\ x \sim_c y \end{cases}$$

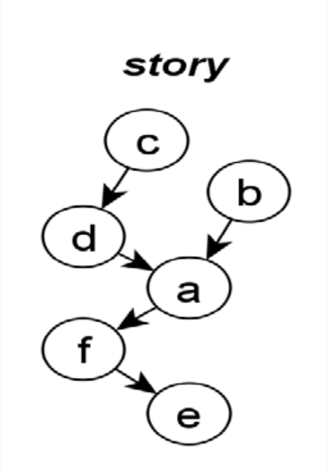

QUESTION	ANSWER			OUTCOME
	$\succ$	$\sim$	$\prec$	
$a ?_s b$	1	0	4	$b \succ_s a$
$a ?_s c$	0	0	5	$c \succ_s a$
$a ?_s d$	0	2	3	$d \succ_s a$
$a ?_s e$	4	0	1	$a \succ_s e$
$a ?_s f$	3	1	1	$a \succ_s f$
$b ?_s c$	1	2	2	$b \sim_s c$
$b ?_s d$	1	3	1	$b \sim_s d$
$b ?_s e$	5	0	0	$b \succ_s e$
$b ?_s f$	4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

(3) Contradiction Resolution

Assume transitivity in preference relation, and enforce it.

case 1

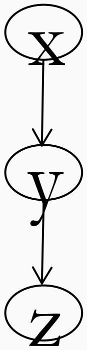

criterion c

(3) Contradiction Resolution

Assume transitivity in preference relation, and enforce it.

case 1

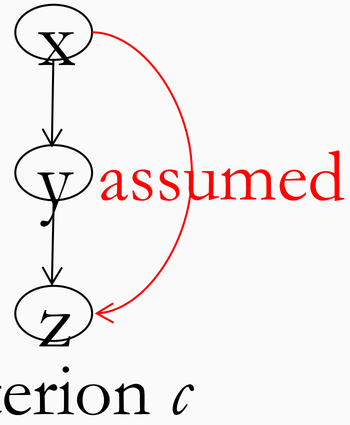

(3) Contradiction Resolution

Assume transitivity in preference relation, and enforce it.

case 1

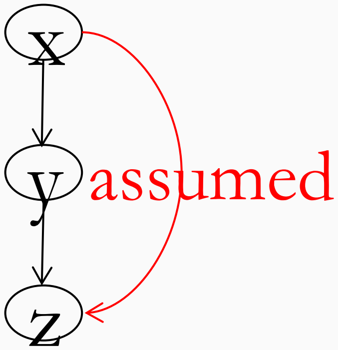

criterion *c*

case 2

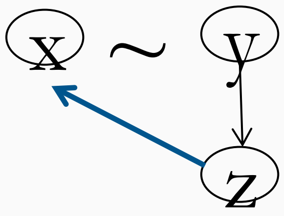

criterion *c*

(3) Contradiction Resolution

Assume transitivity in preference relation, and enforce it.

case 1

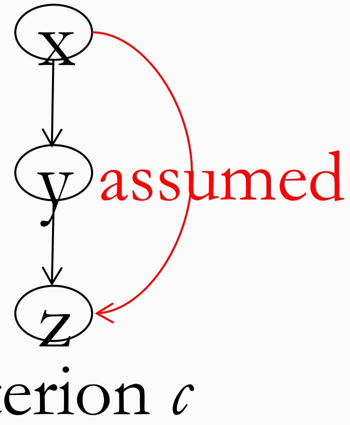

case 2

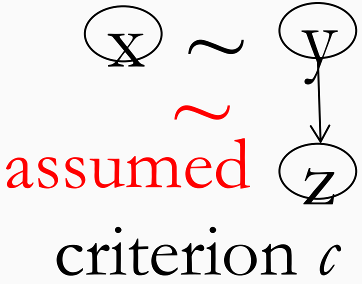

(4) Termination Test

At the end of each iteration, objects are partitioned into 3 sets, based on incomplete preference relations $R^+(Q)$ so far.

$$O_{\surd} = \{x \in O \mid \forall y \in O : (\exists c \in C : x \succ_c y \in R^+(Q)) \vee (\forall c \in C : x \sim_c y \in R^+(Q))\};$$

$$O_{\times} = \{x \in O \mid \exists y \in O : (\forall c \in C : y \succ_c x \in R^+(Q) \vee x \sim_c y \in R^+(Q)) \wedge (\exists c \in C : y \succ_c x \in R^+(Q))\};$$

$$O_{?} = O \backslash (O_{\surd} \cup O_{\times}).$$

- $O\surd$: Pareto-optimal objects
- $O\times$: Non Pareto-optimal objects
- $O?$: $R^+(Q)$ is insufficient for discerning these objects' Pareto-optimality
- $O? = \varnothing \Rightarrow$ terminate

(1) Question Selection

The process of executing a question sequence Q = ⟨q1, ..., qn⟩

- Q is a terminal sequence if O? = ∅ based on $R^{+}(Q)$.
- Goal: among many terminal sequences, execute a short sequence

Lower bound

- Theorem 2: At ldeast (|O|−k)×|C|+(k−1)×2 pairwise comparison questions are necessary, where k is the number of Pareto-optimal objects.

Bad news

- Worst-case : |C| x |O| x (|O|-1)/2 questions; cannot do better than brute-force
- E.g., suppose all objects are indifferent by every criterion. If any comparison $x\ ?_c\ y$ is skipped, we will not be able to determine if x and y are indifferent or if one dominates another.

Transitivity of Object Dominance: Doesn't Hold

A cost-saving property for skyline queries

- Object dominance transitivity: $x \succ y, y \succ z \Rightarrow x \succ z$
- Immediately prune a dominated object from further comparison. (Any object dominated by y is also dominated by x.)

Transitivity of Object Dominance: Doesn't Hold

Fundamental reason: lack of explicit attribute representation

- In skyline/preference queries:
  on any attribute, $x >= y$, $y>=z \Rightarrow x>=z$.
- In Pareto-optimal object finding:
  $x \succ_c$ or $\sim_c y$, $y \succ_c$ or $\sim_c z \Rightarrow x \succ_c$ or $\sim_c z$ (not true)

Even possible that an object is dominated by only one non-Pareto optimal object.

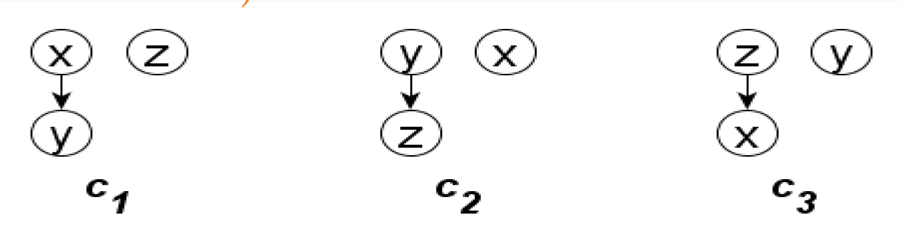

Can Still Benefit From A Similar Idea

For a non-Pareto optimal object, we only need to know at least one object dominates it. We don't care about which other objects also dominate it.

Overriding principle of the framework:

- Identify non-Pareto objects as early as possible
- Postpone their comparisons with other objects as much as possible

Candidate Questions

Given asked questions $Q = \langle q1, ..., qn \rangle$, $x \, ?_c \, y$ is a candidate question iff it satisfies 3 conditions:

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

(ii) $x \in O?$

(iii) Based on R+(Q), $y \succ x$ is not ruled out yet.

How to rule out $y \succ x$?

$\exists c \in C$ such that $x \succ_c y \in R^+(Q) \Rightarrow y \nsucc x$

Only Choosing from Candidate Questions

Sufficient

Property 2: $Qcan = \varnothing \Leftrightarrow O? = \varnothing$

Efficient

Theorem 1: If Q contains non-candidate questions, there exists a shorter or equally long sequence Q' without non-candidate questions such that Q' finds at least all dominated objects found by Q.

Macro-ordering, Micro-ordering

Both guided by the overriding principle

Macro-ordering: When available, we choose a candidate question $x ?_c y$ such that $y \notin O_x$

Micro-ordering: Several question ordering heuristics:

- Random Question (RandomQ): randomly choose a candidate question $x ?_c y$
- Random Pair (RandomP): randomly choose a candidate question $x ?_c y$, continue to finish all remaining candidate questions between x and y.
- Fewest Remaining Questions (FRQ): Choose a pair with the fewest remaining questions. Ties are broken based on how many objects are better/worse than x and y on the criterion.

Experiments by Simulation

- Used an 10000-tuple NBA dataset that records players' per-season performance on 10 criteria (points-per-game, …)

- Simulated Partial orders based on players' performance comparison, with some perturbations.

Effectiveness of Candidate Questions and Macro-Ordering

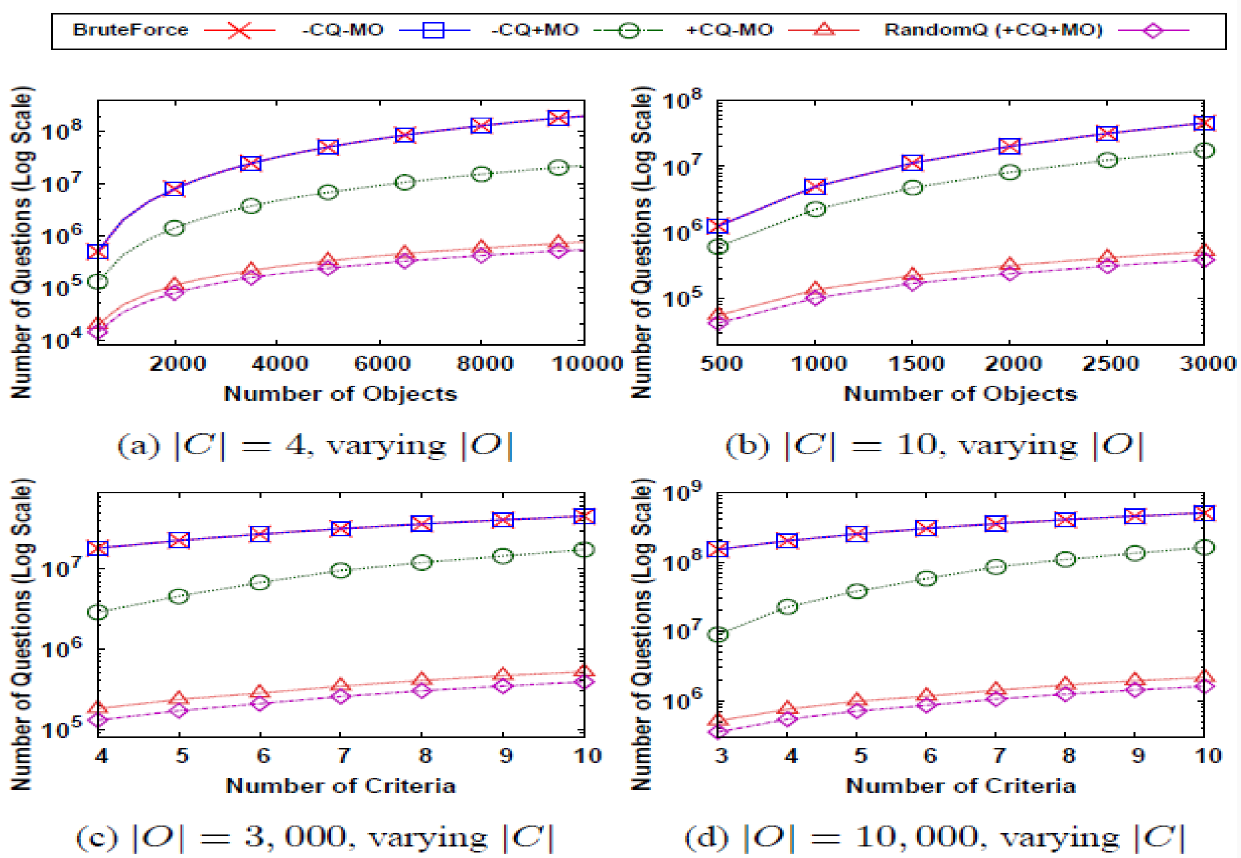

(a) $|C| = 4$, varying $|O|$

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

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

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

Effectiveness of Micro-Ordering

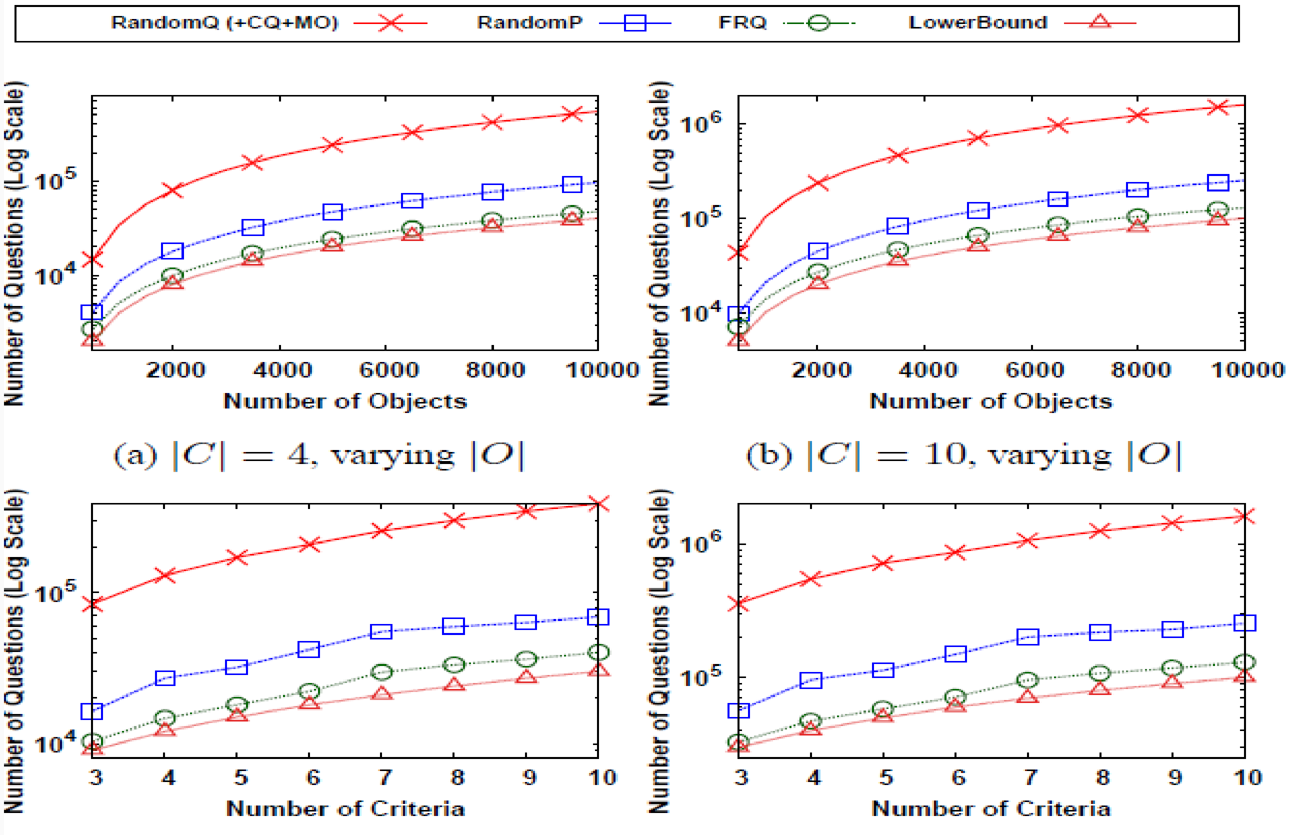

(a) $|C| = 4$, varying $|O|$

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

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

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

Experiments using Amazon Mechanical Turk

- Compare 100 photos of UT-Arlington campus, by color, sharpness, landscape.

- All 14, 850 possible pairwise questions were partitioned into 1, 650 tasks, each containing 9 questions on a criterion.

- Worker qualification:
  - ❑ responded to at least 100 HITs before with at least 90% approval rate
  - ❑ 2 additional validation questions mixed in each task

Experiments using Amazon Mechanical Turk

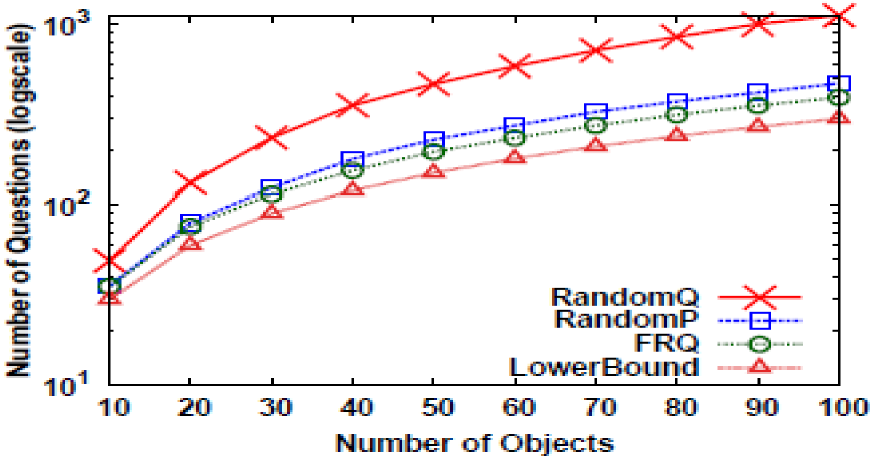

Figure 8: No. of questions by different micro-ordering heuristics. $|C| = 3$, varying $|O|$.

Limitations and Future Work

No performance guarantee: as bad as brute-force in worst-case

Non-deterministic results: due to contradiction resolution

Possibly empty result: due to lack of object dominance transitivity

No consideration of different levels of confidence on question outcomes or crowdsourcers' quality.

Future work: Pareto-optimal objects in probabilistic sense?

No consideration of parallel/batch-execution scheme

Future work: Parallel scheme

Thank You! Questions?

Chengkai Li
http://ranger.uta.edu/~cli
cli@uta.edu

(3) Contradiction Resolution

How often does contradiction happen?

- Depends on data itself, k, and θ
- It may not happen a lot.
  - Intuitively: as long as the underlying relation is transitive, collective wisdom of crowd should reflect it.
  - Preference judgments of relevance in document retrieval are transitive [27, 11].

Brute-Force on the Toy Example

6 objects, 3 criteria: 45 comparisons

RandomQ on the Toy Example

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

RandomP on the Toy Example

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

FRQ on the Toy Example

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	$b \succ_s \underline{a}$		(a,b), $\{m, a\}$			
2	$\underline{a} \succ_m b$	$a \sim b$	(a,c), $\{s, a, m\}$			
3	$c \succ_s \underline{a}$		(a,c), $\{a, m\}$			
4	$c \sim_a \underline{a}$		(a,c), $\{m\}$			
5	$c \succ_m \underline{a}$	$c \succ a$	(b,c), $\{a, s, m\}$	$\varnothing$	{b,c,d,e,f}	{a}
6	$\underline{b} \sim_a c$		(b,c), $\{s, m\}$			
7	$\underline{b} \sim_s c$		(b,c), $\{m\}$			
8	$\underline{b} \succ_m c$	$b \succ_m a$, $b \succ c$	(d,b), $\{a, s, m\}$	$\varnothing$	{b,d,e,f}	{a,c}
9	$b \sim_a \underline{d}$		(d,b), $\{s, m\}$			
10	$b \sim_s \underline{d}$		(d,b), $\{m\}$			
11	$b \succ_m \underline{d}$	$b \succ d$	(e,b), $\{a, s, m\}$	$\varnothing$	{b,e,f}	{a,c,d}
12	$b \succ_a \underline{e}$		(e,b), $\{s, m\}$			
13	$b \succ_s \underline{e}$		(e,b), $\{m\}$			
14	$b \succ_m \underline{e}$	$a \succ_m e$, $b \succ e$	(f,b), $\{a, s, m\}$	$\varnothing$	{b,f}	{a,c,d,e}
15	$b \succ_a \underline{f}$		(f,b), $\{s, m\}$			
16	$b \succ_s \underline{f}$		(f,b), $\{m\}$			
17	$b \succ_m \underline{f}$	$b \succ f$		{b}	$\varnothing$	{a,c,d,e,f}