Published in final edited form as:

Proceedings VLDB Endowment. 2018 September ; 11(13): 2196–2208. doi:10.14778/3275366.3284965.

Interactive Summarization and Exploration of Top Aggregate Query Answers

Yuhao Wen, Xiaodan Zhu, Sudeepa Roy, and Jun Yang

Department of Computer Science, Duke University

Abstract

We present a system for summarization and interactive exploration of high-valued aggregate query answers to make a large set of possible answers more informative to the user. Our system outputs a set of clusters on the high-valued query answers showing their common properties such that the clusters are diverse as much as possible to avoid repeating information, and cover a certain number of top original answers as indicated by the user. Further, the system facilitates interactive exploration of the query answers by helping the user (i) choose combinations of parameters for clustering, (ii) inspect the clusters as well as the elements they contain, and (iii) visualize how changes in parameters affect clustering. We define optimization problems, study their complexity, explore properties of the solutions investigating the semi-lattice structure on the clusters, and propose efficient algorithms and optimizations to achieve these goals. We evaluate our techniques experimentally and discuss our prototype with a graphical user interface that facilitates this interactive exploration. A user study is conducted to evaluate the usability of our approach.

1. INTRODUCTION

Summarization and diversification of query results have recently drawn significant attention in databases and other applications such as keyword search, recommendation systems, and online shopping. The goal of both result summarization and result diversification is to make a large set of possible answers more informative to the user, since the user is likely not to view results beyond a small number. This brings the need to make the *top-k* results displayed to the user *summarized* (the results should be grouped and summarized to reveal high-level patterns among answers), *relevant* (the results should have high *value* or *score* with respect to user's query or a database query), *diverse* (the results should avoid repeating information), and also providing *coverage* (the results should cover top answers from the original non-summarized result set). In this paper, we present a framework to summarize and explore high valued aggregate query answers to understand their common properties easily and efficiently while meeting the above competing goals simultaneously. We illustrate the challenges and our contributions using the following example:

EXAMPLE 1.1. Suppose an analyst is using the movie ratings data from the MovieLens website [37] to investigate average ratings of different genres of movies by different groups of users over different time periods. So the analyst first joins several relations from this

dataset (information about movies, ratings, users, and their occupations) to one relation R , extracts some additional features from the original attributes (age group, decade, half-decade), and then runs the following SQL aggregate query on R (the join is omitted for simplicity). In this query, $hdec$ denotes disjoint five-year windows of half-decades, e.g., 1990 (=1990–94), 1995 (=1995–99), etc.; $agegrp$ denotes age groups of the users in their teens or 10s (i.e., 10–19), 20s (i.e., 20–29), etc.

```
SELECT hdec, agegrp, gender, occupation, avg(rating) as val FROM R
```

```
GROUP BY hdec, agegrp, gender, occupation
```

```
WHERE genres_adventure = 1
```

```
HAVING count(*) > 50
```

```
ORDER BY val DESC
```

The top 8 and bottom 8 results from this query are shown in Figure 1a. To have a quick summary of these 50 result tuples, The data analyst is interested in seeing the summary in at most four rows to have an idea of the viewers and time periods with a high rating for the adventure genre.

One straightforward option is to output the top 4 result tuples from Figure 1a, but they do not summarize the common properties of the intended viewers/times periods. In addition, despite having high scores, they have attribute values that are close to each other (*e.g.*, male students in their 20s) leading to repetition of information and sub-optimal use of the designated space of $k = 4$ rows. More importantly, *the top k original tuples may give a wrong impression on the common properties of high-valued tuples even if they all share those properties*. For instance, three out of top four tuples share the property (20 s, M), but it is misleading, since a closer look at Figure 1a reveals that many tuples with low values (49th, 46th, 44th) share this property too, suggesting that male viewers in their 20s may or may not give high ratings to the adventure genre. Therefore, we aim to achieve a summarization with the following desiderata: (i) it should be simple and memorable (*e.g.*, male students or (M, Students)), (ii) it should be diverse (*e.g.*, (1975, 20 s, M, Student) and (1980, 20 s, M, Student) might be too similar), and (iii) it should be discriminative (*e.g.*, the properties like (20 s, M) covering both high and low valued tuples should be avoided). Furthermore, it should be achieved at an interactive speed and displayed using a user-friendly interface.

In recent years, work has been done to *diversify* a set of result tuples by selecting a subset of them (discussed further in Section 2), *e.g.*, *diversified top- k* [31] takes account of diversity and relevance while selecting top result tuples; *DisC diversity* [8] takes into account similarity with the tuples that have not been selected, and diversity and relevance in the selected ones. In contrast, we intend to output summarized information on the result tuples by displaying the common attribute values in each cluster to give the user a holistic view of the result tuples with high value. In this direction, the *smart drill-down* [24] framework helps the user explore summarized “interesting” tuples in a database, but it does not focus on

aggregate answers, or helping the user choose input parameters and understand consecutive solutions, which are two key features of our framework. As discussed in Section 2 and observed in experiments in our initial exploration, standard clustering or classification approaches do not give a meaningful summary of high-valued results as well. In particular, we support summarization and interactive exploration of aggregate answers in the following ways each posing its own technical challenges.

(1) Summarizing Aggregate Answers with Relevance, Diversity, and Coverage. To meet the desiderata of a good summarization, the basic operation of our framework involves generating a set of *clusters* summarizing the common properties or common attribute values of high-valued answers (Section 3). If all elements in a cluster do not share the same value for an attribute, then the value of that attribute is replaced with a ‘*’¹. The clusters can be expanded to show the elements contained in them to the user. To compute the clusters, our framework can take (up to) three parameters as input: (i) *size constraint* k denotes the number of rows or clusters to be displayed ($k = 4$ in Example 1.1), (ii) *coverage parameter* L , requiring that the top- L tuples in the original ranking must be covered by the k chosen clusters, and (iii) *distance parameter* D , requiring that the summaries should be at least distance D from each other to avoid repeating similar information.

EXAMPLE 1.2. Suppose we run our framework for the query in Example 1.1 with parameters $k = 4$, $L = 8$, and $D = 2$, i.e., the user would like to see at most 4 clusters, these clusters should cover top 8 tuples from Figure 1a, and any two clusters should not have identical values for more than two attributes. Our framework first displays the four clusters shown in Figure 1b along with the average scores of result tuples contained in them.

The user may choose to investigate any of these clusters by expanding the cluster on our framework (clicking ▼). If all four clusters are expanded by the user, the second-layer will reveal all original result tuples they cover, as shown in Figure 1c. In this particular example, no other tuples outside top 8 have been chosen by our algorithm (which is also the optimal solution), but in general, the selected clusters may contain other tuples (high-valued but not necessarily in top L). The above example illustrates several advantages and features of our framework in providing a meaningful and holistic summary of high-valued aggregate query answers. *First*, the original top 8 result tuples are not lost thanks to the second layer, whereas the properties that combine multiple top result tuples are clearly highlighted in the clusters in the first layer. *Second*, the chosen clusters are diverse, each contributing some extra novelty to the answer. *Third*, the clustering captures the properties of the top result tuples that distinguish them from those with low values. For instance, the cluster for (20 s, M) does not appear in the solution, since this property is prevalent in both high-valued and low-valued tuples as discussed before. Clearly, this could not be achieved by simply clustering top L tuples by k clusters. This is ensured by our objective function that aims to maximize the *average* value of the tuples covered by all clusters (instead of maximizing the sum).

¹Our framework and algorithms can be extended to more fine-grained generalizations of values beyond * (by introducing a concept hierarchy over the domain). Details in Appendix A.6.

To achieve the solution as described above, we make the following technical contributions in the paper:

- To ensure that the chosen clusters cover answer tuples with high values, we formulate an optimization problem that takes k , L , D as input, and outputs clusters such that the *average value* of the tuples covered by these clusters is maximized. We study the complexity of the above problem (both decision and optimization versions) and show NP-hardness results (Section 4).
- We design efficient heuristics satisfying all constraints using properties of the *semi-lattice structure* on the clusters imposed by the attributes (Section 5).
- We perform extensive experimental evaluation using the Movie-Lens [37] and TPC-DS benchmark [28] datasets (Section 7).

(2) Interactive Clustering and Parameter Selection. The intended application of our framework is an interactive exploration of query results where the user may keep updating k , L , or D to understand the key properties of the high-valued aggregate answers. One challenge in this exploration is to select values of k , L , D while ensuring interactive speed, since straightforward implementations of our algorithms would not be fast enough. To support parameter selection, we provide the user with a holistic view of how the objective varies with different choices of parameters. This view helps users identify “flat regions” (uninteresting for parameter changes) vs. “knee points” (possibly interesting for parameter changes) in the parameter space. One example is shown in Figure 2, where given selected values of $L = 15$ as in Section 6.1, how the average value of the solutions (y -axis) varies with different k (x -axis) is shown. This figure illustrates that if k is changed from $k = 11$ to $k = 7$, there will be a drop in the overall value. The user can select different legends for different lines and check the value in detail by hovering over a point. This visualization also helps the user validate the choice of parameters, *e.g.*, if a smaller value of k can give a similar quality result, the user may want to reduce the value of k to have a more compact solution. This feature not only helps in guiding the user select parameter values², it also serves as a *precomputation* step to retrieve the actual solutions for different combinations of input parameters k , L , and D at an interactive speed. (Section 6).

- We develop techniques for *incremental computation and efficient storage* for solutions for multiple combinations of input parameters using an *interval tree* data structure.
- We implement multiple optimization techniques to further speed up computation of these solutions. We evaluate the effect of these optimizations experimentally. Eventually we achieved $30\times$ - $1000\times$ speed up from these optimizations, which helped us achieve our goal of interactive speed.

Roadmap. We discuss the related work in Section 2 and define some preliminary concepts in Section 3. The above sets of results are discussed in Sections 4, 5 and 6. The experimental

²To further assist in parameter selection, our system also allows visual comparison of two successive solutions showing how the clusters are redistributed. We formulated an optimization problem to enable clean visualization and provided optimal solutions. The details are in Appendix A.7 and in our demonstration paper in SIGMOD 2018 [43].

results are presented in Section 7. A user study is conducted on Section 8. We conclude in Section 9 with scope of future work. Some details are deferred to Appendix.

2. RELATED WORK

First we discuss three recent papers relevant to our work that consider *result diversification* or *result summarization*: smart drilldown [24], diversified top- k [31], and DisC diversity [8]. We explored using or adapting the approaches proposed in these papers for our problem, but since they focus on different problems, as expected, the optimization, objective, and the setting studied in [24, 31, 8] do not suffice to meet the goals in our work; There are several other related work in the literature that we briefly mention below. Qualitative comparison results and details are discussed in Appendix A.11.

Smart drill-down [24]: In a recent work, Joglekar et al. [24] proposed the *smart drill-down* operator for interactive exploration and summarizing interesting tuples in a given table. The outputs show top- k rules (clusters) with don't-care $*$ -values. The goal is to find an ordered set of rules with maximum score, which is given by the sum of product of the *marginal coverage* (elements in a rules that are not covered by the previous rules) and *weight* of the rules (a “goodness” factor, *e.g.*, a rule with fewer $*$ is better as it is more specific). In Appendix A.5, we show with examples that this approach is not suitable for summarizing aggregate query answers, since it will prefer common attribute values prevalent in many tuples and may select rules containing both high- and low-valued tuples.

Diversified top- k [31]: Qin et al. [31] formulated the top- k result diversification problem: given a relation S where each element has a score and any two elements have a similarity value between them, output at most k elements such that any two selected elements are dissimilar (similarity $>$ a threshold τ), and maximize the sum of the scores of the selected elements. [31] considers diversification, but it does not consider result summarization using $*$ -values (it chooses individual representative elements instead). In addition to lacking high level properties, this adapted process would possibly lose the holistic picture since some low-valued elements may be assigned to the chosen representatives from the top elements.

DisC diversity [8]: Drosou and Pitoura [8] proposed *DisC* diversity: given a set of elements S , the goal is to output a subset S' of smallest size such that all the elements in S are *similar* to at least one element in S' (*i.e.*, have distance at most a given threshold τ), whereas no two elements in S' are similar to each other (distance is $>$ τ). Here diversification can be achieved similar to [31]. However, it ignores the values or relevance of the elements (unlike us or [31]), and has no bound on the number of elements returned (unlike us, [31, 24]). Therefore this approach may not be useful when the user wants to investigate a small set of answers, and it does not provide a summary of common properties of high valued tuples.

Classification and clustering: Classification and clustering have been extensively studied in the literature. Various classification algorithms like Naive Bayes Classifier [27] and decision trees [32] are widely used and are easy to implement. A simpler variation of our problem—separating top- L elements from others—can be cast as a classification problem. However, this formulation would completely ignore values of elements outside top L , whereas our

problem considers all element values and uses the top- L elements only as a coverage constraint. One could also formulate the problem of clustering the top- L elements and apply the standard k -means algorithm [20] and its variants (e.g., [21, 42]). However, such algorithms do not produce clusters with simple and concise descriptions, and their clustering criteria do not consider values of elements outside top L . Therefore, it is necessary to find a new approach other than traditional clustering and classification.

Other work on result diversification, summarization, and exploration: Diversification of query results has been extensively studied in the literature for both query answering in databases and other applications [4, 2, 17, 49, 45, 15, 44, 10, 48, 33, 3, 1, 36, 7, 41, 31, 8, 12, 47, 40, 46, 35, 23, 22]. These include the *MMR (Maximal Marginal Relevance)*-based approaches, the *dispersion* problem studied in the algorithm community, *diverse skyline*, summarization in *text and social networks*, relational data summarization and OLAP data cube exploration among others. The MMR-based and dispersion approaches consider diversification of results, outputting a small, diverse subset of relevant results, but do not summarize all relevant results. Others focus on various application domains and all have problem definitions different from this work.

3. PRELIMINARIES

Let R be a relation with attributes $\mathcal{A}$, which can either be an input table (base relation) or a derived relation possibly coming from a complex sub-query involving multiple tables. Let $\mathcal{A}_{groupby} \subseteq \mathcal{A}$ be a set of *grouping attributes* used in the group-by clause where $|\mathcal{A}| = m$. Let $aggr$ be any aggregate function allowed by SQL that outputs a real number. Therefore, we are considering a query Q of the form:

```
SELECT  $\mathcal{A}_{groupby}$ , aggr as val
FROM R <base relation or output of a sub-query>
GROUP BY  $\mathcal{A}_{groupby}$ 
ORDER BY val DESC
```

We denote the output of this query as S , where each tuple in S is called an *original element*. Here val denotes the *score* or *value* of each output tuple in S signifying relevance or importance of the tuple in response to the input query³. Usually, the query will output n tuples (i.e., $|S| = n$). Even if the number of attributes m in the group-by clause is small, n might be large due to large domains of the participating attributes. Therefore, a user frequently runs a top- L query to retrieve the top- L tuples (denoted by S_L^*) with highest scores (adding a LIMIT L clause to the above query).

³Our work is also applicable to the settings where the scores of tuples do not come from an SQL query (e.g., are given by a domain expert).

Clusters. To display a solution with relevance, diversity, and coverage, our output is provided in *two layers*: the top layer displays a set of *clusters* that hide the values of some attributes by replacing them with *don't-care* (*) values, and the second layer contains the original elements covered by them.

For every original element t in the output S of Q , let $\text{val}(t)$ denote the value or score of t . Other than the value, each $t \in S$ has m attributes $A_1, \dots, A_m$ with active domains $D_1, \dots, D_m$ respectively, A *cluster* C on S has the form: $C \in \prod_{i=1}^m D_i \cup \{*\}$. Let $\mathcal{C}$ denote the set of all clusters for relation S . We assume that the m attributes $A_1, \dots, A_m$ have a predefined order, and therefore we omit their names to specify a cluster. For instance, for $m=4$ attributes A_1, A_2, A_3, A_4 , $(a_1, b_1, *, *)$ implies that $(A_1 = a_1) \wedge (A_2 = b_1)$, and the values of A_3 and A_4 are don't-care (*). We denote the value of an attribute A_i of C by $C[A_i]$; where $C[A_i] \in D_i \cup \{*\}$, $i \in [1, m]$. In particular, each element t in S also qualifies as a cluster, which is called a *singleton cluster*.

A cluster C *covers* another cluster C' if $\forall i \in [1, m]$, $C[A_i] = *$ or $C[A_i] = C'[A_i]$. Since each element t in S is also a cluster, each cluster C covers some elements from S . Further, the notion of coverage naturally extends to a subset of clusters $\mathcal{O}$. For $C \in \mathcal{C}$, $\text{cov}(C) \subseteq S$ denotes the elements covered by C , and for $\mathcal{O} \subseteq \mathcal{C}$, $\text{cov}(\mathcal{O}) \subseteq S$ denotes the elements covered by at least one cluster in $\mathcal{O}$, i.e., $\text{cov}(\mathcal{O}) = \bigcup_{C \in \mathcal{O}} \text{cov}(C)$. Figure 3a shows two clusters $C_1 = (*, *, c_1, d_1)$, $C_2 = (a_2, b_1, *, d_1)$, and the elements they cover. Note that two clusters may have overlaps in elements they cover. Here C_1, C_2 have overlap on the tuple (a_2, b_1, c_1, d_1) .

Distance function. While the distance between two elements is straightforward (the number of attributes where their values differ⁴), the distance between two clusters has several alternatives due to the presence of the don't care (*) values. We define the distance between two clusters as the number of attributes where they do not have the same value from the domain. The distance function can be shown to be a metric and it exhibits monotonicity property (discussed in Section 4) that we use in our algorithms.

Definition 3.1. The distance $d(t, t')$ between two elements t, t' is the number of attributes where their values differ, i.e., $d(t, t') = |\{i \in [1, m]: t[A_i] \neq t'[A_i]\}|$. The distance between two clusters C, C' is the number of attributes where either (i) at least one of the values is *, or (ii) the values are different in C, C' : $d(C, C') = |\{i \in [1, m]: C[A_i] = *, \text{ or } C'[A_i] = *, \text{ or } C[A_i] \neq C'[A_i]\}|$.

In Figure 3a, the distance between $C_1 = (*, *, c_1, d_1)$ and $C_2 = (a_2, b_1, *, d_1)$ is 3 due to the presence of *-s in A_1, A_2, A_3 . Intuitively, the distance between two clusters is the maximum possible distance between any two elements that these two clusters may contain, and therefore is measured by counting the number of attributes where they do not agree on a value from the domain. The distance function can also be explained in terms of similarity measures between two tuples or clusters: if the distance between two clusters is $> D$, then

⁴In this paper we focus on categorical attributes; other distance functions suitable for numeric attributes is a direction for future work (Section 9).

the number of common attribute values between them is $m - D$ where m is the total number of attributes.

4. FRAMEWORK

In this section, we discuss the technical details for our framework: Section 4.1 formally defines the optimization problem. Section 4.2 discusses the semilattice structure and properties of the clusters, and Section 4.3 discusses the complexity of the optimization problem.

4.1 Optimization Problem Definition

For a cluster C , let $\text{avg}(C)$ denote the average value of all the elements contained in C , i.e., $\text{avg}(C) = \frac{\sum_{t \in \text{cov}(C)} \text{val}(t)}{|\text{cov}(C)|}$. Similarly, for a set of clusters $\mathcal{O}$, $\text{avg}(\mathcal{O})$ denotes the average value of the elements covered by $\mathcal{O}$.

DEFINITION 4.1. Given relation S with original tuples and their values, size constraint k , coverage constraint L , distance constraint D , and set $\mathcal{C}$ of possible clusters for S , a subset $\mathcal{O} \subseteq \mathcal{C}$ is called a feasible solution if all the following conditions hold: (1) (**Size** k) The number of clusters in $\mathcal{O}$ is at most k , i.e., $|\mathcal{O}| \leq k$. (2) (**Coverage** L) $\mathcal{O}$ covers all top- L elements in S , i.e., $S_L^* \subseteq \text{cov}(\mathcal{O})$. (3) (**Distance** D) The distance between any two clusters C_1, C_2 in $\mathcal{O}$ is at least D , i.e., $d(C_1, C_2) \geq D$. (4) (**Incomparability**) No clusters in $\mathcal{O}$ cover any other cluster in $\mathcal{O}$ {equivalently, the clusters should form an antichain in the semilattice discussed in Section 4.2}. The objective (called Max-Avg) is to find a feasible solution $\mathcal{O}$ with maximum average value $\text{avg}(\mathcal{O})$.

The first three conditions in the above definition correspond to the input parameters, whereas the last condition eliminates unnecessary information from the returned solution. All these three parameters, k , D , and L , are optional and can have a default value; e.g., the default value of k can be n , if there is no constraint on the maximum number of clusters that can be shown. If maintaining diversity in the answer set is not of interest, then D can be set to 0. Similarly, if coverage is not of interest, L can be set to 0 (to display a set of clusters with high overall value), or 1 (to cover the element with the highest value in S), or to k (to cover the original top- k elements from S). To maintain all the constraints, the chosen clusters may pick up some *redundant elements* $t \in S_L^*$ that do not belong to the top- L elements.

The optimization objective, called Max-Avg, intuitively highlights the important attribute-value pairs across all tuples with high values in S , even if they are outside the top- L elements.⁵ In any solution, the value of each covered element contributes only once to the objective function, hence the selected clusters in $\mathcal{O}$ do not get any benefit by covering the elements with high value multiple times. In fact, the optimal solution when $D = 0$ and $k > L$ is obtained by selecting top- k original elements. The optimal solution considers the average

⁵We also investigated an alternative objective called Min-Size that minimizes the number of redundant elements. However it may miss some interesting global properties covering many high-valued elements in S , and is less useful for summarization.

value instead of their sum since otherwise, always the *trivial solution* $(*, *, \dots, *)$ covering all elements and satisfying all constraints will be chosen.

4.2 Semilattice on Clusters and Properties

A *partially ordered set* (or, *poset*) is a binary relation over a set of elements that is reflexive ($a \leq a$), antisymmetric ($a \leq b, b \leq a \Rightarrow a = b$), and transitive ($a \leq b, b \leq c \Rightarrow a \leq c$). A poset is a *semilattice* if it has a *join or least upper bound* for any non-empty finite subset. The coverage of elements described in Section 3 naturally induces a semilattice structure on our clusters $\mathcal{C}$, where for any two clusters $C, C' \in \mathcal{C}$, $C \leq C'$, $C \sqcup C' = C'$ if and only if C' covers C , i.e., $\text{cov}[C] \subseteq \text{cov}[C']$. If $C \leq C'$, then C' is called an *ancestor* of C in the semilattice, and C is a *descendant* of C' . Equivalently, if a cluster C_{up} covers another cluster C_{down} by replacing exactly one attribute value of C_{down} by the don't care value (*), then we draw an edge between them, and put C_{up} at one level higher than C_{down} in the semilattice (this gives a *transitive reduction* of the poset). Level ℓ of the semilattice is the set of clusters with exactly ℓ^* values. Figure 3b shows the semilattice structure of $\mathcal{C}$ that has two attributes A_1 and A_2 , where the domains are $D_1 = \{a_1, a_2\}$ and $D_2 = \{b_1, b_2\}$. The distance function described in Section 3 has a nice monotonicity property that we use in devising our algorithms in Section 5 (proof is in Appendix A.1):

Proposition 4.2. (Monotonicity) Let $\mathcal{O}$ be a set of clusters. Let λ be the minimum distance between any two clusters in $\mathcal{O}$ as defined in Definition 3.1, i.e. $\lambda = \min_{C, C' \in \mathcal{O}} d(C, C')$. Let $\mathcal{SC}' = (\mathcal{SC} \setminus \{C_1\}) \cup \{C_2\}$, where a cluster C_1 is replaced by another cluster C_2 such that C_2 covers C_1 (i.e., C_2 is an ancestor of C_1 in the semilattice). Let λ' be the minimum distance in $\mathcal{SC}'$, i.e., $\lambda' = \min_{C, C' \in \mathcal{SC}'} d(C, C')$. Then $\lambda' \geq \lambda$.

Assuming the semilattice structure in Figure 3b, note that $\{(a_1, b_2), (*, b_1)\}$ satisfies the distance constraint for $D = 2$. If we replace (a_1, b_1) by one of its ancestors $(a_1, *)$, the new two clusters $\{(a_1, *), (*, b_1)\}$ also satisfies the constraint for $D = 2$.

4.3 Complexity Analysis

The optimization problem can be solved in polynomial time in *data complexity* [39] if the size limit k is a constant. This is because we can iterate over all possible subsets of the clusters of size at most k , check if they form a feasible solution, and then return the one with the maximum average value. However, this does not give us an efficient algorithm to meet our goal of interactive performance. For example, if $k = 10$, the domain size of each attribute is 9, and the number of attributes is 4, the number of clusters (say N) can be 10^4 , and the number of subsets will be of the order of $N^k = 10^{40}$.

When k is variable, the complexity of the problem may arise due to any of the four factors in Definition 4.1: the size constraint k , the coverage parameter L , the distance parameter D , and the incomparability requirement that the output clusters should form an antichain. Due to multiple constraints, it is not surprising that in general, even checking if there is a non-trivial feasible solution is NP-hard. In particular, when $k \leq L$, simply the requirement of covering L original elements by k clusters in a feasible solution lead to NP-hardness without any other

constraints. However, in the case when $k = L$ (the user is willing to see L clusters), the decision and optimization problems become relatively easier.

In Appendix A.2 we show the following: (1) $k = L, D = 0$: Top- k elements give the optimal solution, since adding any redundant element worsens the Max-Avg objective. (2) $k = L$, arbitrary D : A non-trivial feasible solution always exists, since we can pick arbitrary ancestors of each top- L element from level $D - 1$ satisfying all the constraints. However, the optimization problems are NP-hard. (3) $k < L, D = 0$: Even checking whether a nontrivial feasible solution exists is NP-hard. (4) $k < L$, arbitrary D : The same hardness as above holds.

Although the optimization problem shows similarity with set cover, for a formal reduction, we need to construct an instance of our problem by creating a set of tuples and ensure that the ‘sets’ in this reduction conform to a semi-lattice structure. To achieve this, we give reductions from the **tripartite vertex cover** problem that is known to be NP-hard [25], and construct instances S with only $m = 3$ attributes. The NP-hardness proof the optimization problem for $k = L$ is more involved than the NP-hardness proof for the decision problem for $k < L$, since in the former case the coverage constraint with $k = L$ does not lead to the hardness.

5. ALGORITHMS

Given that the optimization problem for the case $k = L$, and even the decision problem for the case $k < L$, are NP-hard, we design efficient heuristics that are implemented in our prototype and are evaluated by experiments later. Not only finding provably optimal solutions for our objectives is computationally hard, but designing efficient heuristics for these optimization problems is also non-trivial. The optimization problem in Definition 4.1 has four orthogonal objectives for feasibility: incomparability, size constraint k , distance constraint D , coverage constraint L . In addition, the chosen clusters should have high quality in terms of their overall average value. In Section 5.1, we discuss the Bottom-Up algorithm that starts with L singleton clusters satisfying the coverage constraint, and merges clusters greedily when they violate the distance, incomparability, or the size constraints. Then in Section 5.2, we discuss an alternative to Bottom-Up that we call the Fixed-Order algorithm that builds a feasible solution incrementally considering each of the top- L elements one by one. In general, Bottom-Up gives better quality solution and as discussed in Section 6, is amenable to processing of multiple parameter settings as precomputation, whereas Fixed-Order is more efficient, hence in Section 5.3 we describe a Hybrid algorithm combining these two.

5.1 The Bottom-Up Greedy Algorithm

Here we start with L singleton clusters with the top- L elements as our current solution $\mathcal{O}$, which satisfies the coverage and incomparability constraints, but may violate size and distance constraints. Then we iteratively merge clusters in two phases: the *first phase* ensures that no two clusters in $\mathcal{O}$ are within distance D of each other, the *second phase* ensures that the number of clusters is k or less. The following invariants are maintained by the algorithm at all time steps: (1) (*Coverage*) Clusters in $\mathcal{O}$ cover the top- L answers. (2) (*Incomparability*) No cluster in $\mathcal{O}$ covers another. (3) (*Distance*) The minimum distance among the pairs of

clusters in $\mathcal{O}$ never decreases. During the execution of the algorithm, the only operation is *merging of clusters*, therefore, the coverage invariant above is always maintained. Further, the Merge procedure described below maintains the incomparability invariant.

The Merge($\mathcal{O}, C_1, C_2$) procedure. Given two clusters $C_1, C_2 \in \mathcal{O}$, the Merge($\mathcal{O}, C_1, C_2$) procedure replaces C_1, C_2 by a new cluster $C_{new} = \text{LCA}(C_1, C_2)$, their *least common ancestor*, and also removes any other cluster in $\mathcal{O}$ that is also covered by C_{new} . LCA(C_1, C_2) is computed simply by replacing by * any attribute whose values in C_1, C_2 differ. For instance, the LCA of $(a_1, *, c_1, *)$ and $(a_1, b_2, c_2, *)$ is $(a_1, *, *, *)$. Further, if another cluster $(a_1, b_3, *, *)$ belongs to $\mathcal{O}$, Merge would also remove this cluster, since it is covered by $(a_1, *, *, *)$.

Algorithm 1 The Bottom-Up algorithm

Input: Size, coverage, and distance constraints k, L, D
1: $\mathcal{O} =$ set of L singleton clusters with the top- L elements.
2: /* First phase to enforce distance */
3: **while** $\mathcal{O}$ has two clusters with distance $< D$ **do**
4: Let P_D be the pairs of clusters in $\mathcal{O}$ at distance $< D$.
5: Perform UpdateSolution($\mathcal{O}, P_D$).
6: **end while**
7: /* Second phase to enforce size limit k , almost the same as above except
all pairs of clusters are considered. */
8: **while** $|\text{soln}| > k$ **do**
9: Let P_{all} be the all pairs of clusters in $\mathcal{O}$.
10: Perform UpdateSolution($\mathcal{O}, P_{all}$).
11: **end while**
12: **return** $\mathcal{O}$
13: **Procedure** UpdateSolution($\mathcal{O}, P$)
14: **Input:** current solution $\mathcal{O}$, a set of pairs P of clusters
15: $(C_1, C_2) = \text{argmax}_{(C_1, C_2) \in P} \text{avg}(\mathcal{O} \cup \text{LCA}(C_1, C_2))$
16: Perform Merge($\mathcal{O}, C_1, C_2$).

In addition to maintaining the coverage condition, the merging process does not add any new violations to the distance condition in $\mathcal{O}$. This follows from the monotonicity of the distance condition given in Proposition 4.2. However, due to the merging process, the value of the solution may decrease, since $\text{LCA}(C_1, C_2)$ covers all the elements covered by C_1, C_2 and all other clusters that are removed from $\mathcal{O}$, and can potentially cover some more.

The bottom-up algorithm is given in Algorithm 1. The UpdateSolution($\mathcal{O}, P$) procedure used in this algorithm takes the current solution $\mathcal{O}$ and a set of pairs P to be considered for merging, and greedily merges a pair. The first and second phases of Algorithm 1 are very similar, the only difference being the pairs of clusters P they consider for merging. In the first phase, only the pairs with distance $< D$ are considered, whereas in the second phase, all pairs of clusters in $\mathcal{O}$ are considered for merging.

We also implemented and evaluated other variants of bottom-up algorithms: (i) when we start at the clusters at level $D - 1$ (instead of individual top- L tuples that satisfy the distance constraint), and (ii) when we greedily merge pairs C_1, C_2 with maximum value of $\text{avg}(\text{LCA}(C_1, C_2))$ (instead of maximum average value of the overall solution after merging). Both these variants had efficiency and quality comparable or worse than the basic Bottom-Up algorithm as observed in our experiments.

5.2 The Fixed-Order Greedy Algorithm

The Fixed-Order algorithm maintains a set of clusters $\mathcal{O}$, and considers top- L elements in descending order by value. It decides whether the next element is already covered by an existing cluster in $\mathcal{O}$ or can be added as is (satisfying D and k constraints); otherwise Fixed-Order merges the element with one of the existing clusters in greedy fashion. All the constraints (k , D , and incomparability of clusters) are maintained after each of the top- L is processed, so at the end the coverage on top- L is satisfied too. Fixed-Order considers a smaller solution space than Bottom-Up, since it processes each top- L element in an online fashion, and therefore may return a solution with worse value. However, instead of all pairs of initial clusters (quadratic in number of clusters) it considers each cluster only once (linear), so has better running time than Bottom-Up. Details and pseudocode for Fixed-Order are shown in Appendix A.4.

We also consider two variants of Fixed-Order and evaluate them later in experiments: i) *k-means*-Fixed-Order, where we first run the *k*-means clustering algorithm [20] (with random seeding) on the top L elements, find the minimum pattern covering all elements in each of the resulting clusters, and make Fixed-Order process these k patterns first before moving on to the top L elements (in descending-value order); ii) *random*-Fixed-Order, where we first pick k element at random from the top L elements to process first, before moving on to the remaining top L elements (still in descending-value order). Both variants introduce some randomness in the results, and *k-means*-Fixed-Order has considerable higher initial processing overhead. However, as we shall see in Section 7, they do not produce higher-quality results.

5.3 The Hybrid Greedy Algorithm

Bottom-Up tends to produce results with higher quality than Fixed-Order, and can process multiple k , D values at the same time as discussed in Section 6, but usually requires more iterations than Fixed-Order. In order to get a good trade-off between these factors, we introduce the hybrid algorithm. It has two phases - the Fixed-Order phase and Bottom-Up phase. For a given k , L , and D , the first phase for Hybrid is the same as Fixed-Order, but with a larger number of $c \times k$, $c > 1$ is a constant, initial singleton clusters. After covering all top- L elements in $c \times k$ clusters, Hybrid goes into the Bottom-Up phase to reduce the number of clusters from $c \times k$ to k using the Merge procedure that can collect redundant elements. Like Bottom-Up, Hybrid also helps in incremental computation for different choices of parameters as discussed in the next section.

6. INTERACTIVE PARAMETER SELECTION

One of the main challenges in a system with multiple input parameters is choosing the input parameter combination carefully to help the user explore new interesting scenarios in the answer space. To help the user choose interesting values of k , L , D , we provide an overall view of the values of the solutions (average value of all element covered by the clusters chosen by our algorithm) that at the same time precomputes the results for certain parameter combinations and helps in interactive exploration. In Section 6.1 we describe the visualization facilitating parameter selection, in Section 6.2 we discuss how the

precomputation is achieved to plot these graphs, and in Section 6.3 we discuss a number of optimizations for interactive performance of our approach.

6.1 Visual Guide for Parameter Selection

Figure 2 gives an example of visualization showing the overview of the solutions that is generated for each chosen value of L , and illustrates the values of the solutions for a range of choices on D and k . The y -axis shows the average value of the tuples covered by the chosen clusters by our algorithm (Definition 4.1), the value of k (in a chosen range) varies along the x -axis, and different lines correspond to different values of D (also in a chosen range).

With the help of this visualization, the user can avoid selecting certain *uninteresting* or *redundant* parameter combinations. For example, with the visualization in Figure 2, a user can quickly see that the bottom-left region (where $k = 2, 3$) is uninteresting, with low average values. The user also sees that certain ranges of parameter settings are not worth exploring as they do not affect the solution quality or in very predictable ways: *e.g.*, for $D = 1$, the range of $k > 12$ yields almost the same solution quality, while for $k \in [2, 9]$, the quality changes predictably with k . On the other hand, the “knee points” (*e.g.*, $k = 9, 11$ for $D = 1$) suggest good choices of parameters. The visualization also reveals the trade-off between different choices of D ; *e.g.*, at $k = 9$, the user can decide between a solution set with a higher value ($D = 1$) or more diversity ($D = 2$). Note that in Figure 2, curves for different D values may overlap, which suggests ranges of D values with little impact on solution quality, allowing the user to work on “bundles” of D values instead of individually. (If the user cares about curves for individual D values, the user can click on the legend on the right to hide particular curves to reveal others that overlap.)

6.2 Incremental Computation and Storage

To be able to generate plots in Figure 2, one obvious approach is running an algorithm from Section 5 for all combinations of k and D given an L value. However, for interactive exploration, this approach is sub-optimal. The *Hybrid* algorithm (and Bottom-Up) exhibits *two levels* of incremental properties that help in computing the solutions for a range of k, D values in a batch.

In *Hybrid*, for a given value of L , the Fixed-Order phase outputs a set of initial clusters that can be used for all combinations of k, D , and therefore, this step can run only once. Remembering this intermediate solution, the Bottom-Up phase can run for all D values from the stored status. For each D , it computes results for all k values (ranging from the maximum to the minimum value) since in every round of iteration, two clusters are merged to reduce the number of clusters by one. The procedure for this incremental computation is shown in Figure 4a. In the following, we discuss how we materialize and index solutions for efficient retrieval.

Retrieval Data Structure. The computed solutions for different k, D values serve as pre-computed solutions when the user wants to inspect the solution in detail for a certain choice of k, L, D . The obvious solution for storage is to record the set of output clusters for every choice of (k, D) . However, we implemented a combined retrieval data structure for storage

that is both space and time efficient based on the following observation in the execution of Hybrid (and Bottom-Up) algorithm:

Proposition 6.1. (Continuity) Given solution cluster lists $\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_r$ where r rounds are executed, for any cluster $c \in \mathcal{O}_a$ where $1 \leq a < i$, once c is removed from $\mathcal{O}_i$ at the end of round i (because of merging), for all $j > i$, $c \notin \mathcal{O}_j$.

In other words, once a cluster is merged and therefore vanishes from the set of clusters in the solution, it never comes back. Hence, if $\mathcal{O}_{L,D,k}$ denotes the solution for a given combination of L, D, k , the set of values of k for which a cluster $c \in \mathcal{O}_{L,D,k}$ forms a continuous interval.

Therefore, instead of storing the set of clusters for all values of D, k given an L value (where the solutions may have substantial overlap), we use an *interval tree*[6] S_D for each value of D to store the range of k for which a cluster appears in $\mathcal{O}_{L,D,k}$ storing only the maximum (or starting) and minimum (or ending) k value for this cluster (see Figure 4b). It reduces the number of solutions (sets of clusters) to be stored from $O(N_k \times N_D)$ (where N_k and N_D denote the total number of k and D values under consideration respectively) to $O(N_D)$.

Further, the interval tree data structure supports efficient retrieval in time $O(\log N_k)$ [6].

6.3 Optimizations

A number of additional optimizations are implemented to make the system efficient and interactive as described below.

Delta judgment. In every iteration (called *round*) of greedy cluster merging in the Hybrid (and Bottom-Up) algorithm, clusters are merged such that the average value of the clusters in the resulting solution is maximized using the UpdateSolution function (Algorithm 1). Let $\mathcal{O}_i$ be the set of clusters at the end of a round i , $T_i = \text{cov}(\mathcal{O}_i)$ be the tuples covered by $\mathcal{O}_i$, $v_i = \text{avg}(\mathcal{O}_i)$ be the average value of $\mathcal{O}_i$, and $T_c = \text{cov}(c)$ be the tuples covered by a given cluster c . The naive way of executing UpdateSolution in round $i + 1$ involves comparing the tuple list T_c of a given cluster $c (= \text{LCA}(C_1, C_2))$ as mentioned in Algorithm 1) and the current set of covered tuples T_i , finding out new tuples in $T_c \setminus T_i$ to obtain $T_i \cup T_c$ as potential T_{i+1} , and recalculating the objective $\text{avg}(T_i \cup T_c)$ based on the new tuples. However, it takes a huge amount of time doing all the tuple-wise comparison for all possible clusters that are eligible to be merged in this round. Instead, we incrementally keep track of the marginal benefit (as sum and count to compute the average) that a cluster c brings to the new solution $\mathcal{O}_{i+1}$ compared to $\mathcal{O}_i$ as follows (pseudocode in Algorithm 2).

The basic idea is that the improvement in the total average value that a cluster c brings to solution $\mathcal{O}_i$ is due to the tuples in $T_c \setminus \mathcal{O}_i$, and that it brings to $\mathcal{O}_{i-1}$ is due to the tuples in $T_c \setminus \mathcal{O}_{i-1}$. The difference can be computed by keeping track of the new tuples that appear in $T_i \setminus T_{i-1}$, and comparing them with the tuples in T_c . In addition, we incrementally store i, c, sum and i, c, count (the sum of values and the count of tuples in $T_c \setminus T_i$ incrementally computed from $i-1, c, \text{sum}$ and $i-1, c, \text{count}$). Hence the tentative new average value of the solution

$\mathcal{O}_{i+1}$ if we add c to $\mathcal{O}_i$ can be computed as $v_{i+1} = \frac{v_i \times |T_i| + \Delta_{i,c,sum}}{|T_i| + \Delta_{i,c,cnt}}$. This optimization evaluates the UpdateSolution procedure efficiently since the above computations need comparisons between (i) the list containing $T_j \setminus T_{j-1}$ and (ii) T_c , and $T_j \setminus T_{j-1}$ is likely to be much smaller than T_j . This gives 30x speedup in our experiments.

Algorithm 2 The Delta-Judgment Procedure

Input: Marginal score benefit Δ_{sum} , marginal amount benefit Δ_{cnt} , round indicator i indicating when were Δ values last updated, current round $j + 1$, difference list $T_{(j,j-1)} = T_j \setminus T_{j-1}$

```

1: /*  $\Delta_{i-1,c,sum} = \Delta_{sum}$  and same for  $\Delta_{i-1,c,cnt}$  in this case */
2:  $v_{j+1}$  = new score to be calculated.
3:  $v_j$  = current score.
4: /* Marginal benefits are far outdated, -1 is default value;  $i \leq j - 1$  means  $\Delta_{i-1,c,sum}$  and  $\Delta_{i-1,c,cnt}$  cannot be updated directly using  $T_{(j,j-1)}$  */
5: if  $i = -1$  or  $(j - i \geq 1)$  then
6:    $\Delta_{cnt} = \sum(val(T_j \setminus T_c))$ 
7:    $\Delta_{sum} = |T_j \setminus T_c|$ 
8:   /* Marginal benefits were updated last round (round  $j$ ) and can use  $T_{(j,j-1)}$  for comparison */
9: else if  $i = j$  then
10:   $\Delta_{cnt} = \sum(val(T_{(j,j-1)} \setminus T_c))$ 
11:   $\Delta_{sum} = |T_{(j,j-1)} \setminus T_c|$ 
12:  /* Marginal benefits were updated in the same round */
13: else if  $i = j + 1$  then
14:  /* Do nothing. It is up-to-date. */
15: end if
16:  $v_{j+1} = (v_j \times |\mathcal{O}_j| + \Delta_{sum}) / (|\mathcal{O}_j| + \Delta_{cnt})$ 
17: /* Update the round indicator */
18:  $i = j + 1$ 
19: return  $v_{j+1}$ 

```

Cluster generation and mapping to tuples. The semilattice structure on the clusters given an L value is required to run our algorithms that may contain a number of clusters in a naive implementation. To reduce this space to contain only the relevant clusters, clusters are first generated by each tuple in top- L , which ensures that each generated cluster is a possible cluster covering at least one tuple in top- L . Besides, we need to maintain mappings between clusters and the tuples they contain, for which tuples generate *matching expressions* for their target clusters and search through the cluster list (instead of starting with a cluster and searching for matching tuples). Experiments in Section 7.3 shows the benefit - 100x—1000x speedup in running time.

Hash values for fields. The value of an attribute is often found to be text (or other non-numeric value). While storing information on the clusters, we maintain hashmaps for each field between actual values and integer hash values, and store the hash values inside each cluster (mapped back to the original values in the output). This optimization reduces the running time of the order of 50x.

7. EXPERIMENTS

We develop an end-to-end prototype with a graphical user interface (GUI) to help users interact with the solutions returned by our two-layered framework. The prototype is built using Java, Scala, and HTML/CSS/JavaScript as a web application based on Play Framework 2.4, and it uses PostgreSQL at the backend (screenshots of the graphical user interface can be found in the demonstration paper for our system [43]). In this section we

experimentally evaluate our algorithms using our prototype by varying different parameters (Section 7.1), and then test the precomputation and guidance performance (Section 7.2). The effects of optimizations are given in Section 7.3, scalability of our algorithms for a larger dataset is discussed in Section 7.4.

Datasets. In most of the experiments, we use the MovieLens 100K dataset [38, 37, 19]. We join all the tables in the database (for movie-ratings, users, their occupation, etc) and materialize the universal table as *RatingTable*. Each tuple in this rating table has 33 attributes of three types: (a) *binary* (e.g., whether or not the movie is a comedy or action movie), (b) *numeric* (e.g., age of the user), and (c) *categorical* (e.g., occupation of the user). We join the tables as a precomputation step to avoid any interference while measuring the running time of our algorithms.

The other dataset we use is TPC-DS benchmark [28] primarily for evaluating scalability of our algorithms. The table we materialized via generator is *Store_Sales*, which contains 23 attributes and 2,880,404 tuples in total. The aggregate queries used for these two datasets (average rating for MovieLens and average net profit for TPC-DS) can be found in Appendix A.8. All experiments were run on a 64-bit Ubuntu 14.04.4 LTS machine, with Intel Core i7-2600 CPU (4096 MB RAM, 8-core, 3.40GHz).

7.1 Varying Parameters

Unless mentioned otherwise, the three algorithms from Section 5 are compared in this section: (i) Bottom-Up, (ii) Fixed-Order, (iii) Hybrid. In the plots showing the values, we also include (iv) Lower Bound: value of the trivial (feasible) solution (a single cluster with don't-care * values for all attributes) as a baseline.

Comparison with baselines. We compare our algorithms with two baselines: the brute-force algorithm considers all possible cluster combinations, and outputs the global optimal; the lower-bound algorithm simply returns the trivial answer containing one single cluster with all attributes as “*”, which is always feasible for any value of k, L, D . We also consider two variants of Fixed-Order: random and k -means, discussed in Section 5.2. Figure 5a shows the running time for $L = 5, D = 3$ and $k = 2, 3, 4$ (lower-bound is omitted because it returns trivial answers). Even with such small parameter values, the brute-force algorithm is not practical: e.g., at $k = 4$, it takes more than 2.5 hours. Figure 5b compares the average values produced by different algorithms. Since the random and k -means variants of Fixed-Order are randomized, we report their average values over 100 runs each. From Figure 5b, we see that the results of Fixed-Order and its variants are comparable with brute-force's, and are much better than the trivial solution. Another observation is that neither random nor k -means variant improves the quality of plain Fixed-Order. Further, they introduce more variance in the result quality (0.033 for random and 0.045 for k -means in terms of combined standard deviation), and slightly increase the running time. Therefore, in the rest of the section, we focus on the plain Fixed-Order algorithm.

Effect of size parameter k . Figure 6a shows the running time varying k . The running time of Fixed-Order is the best as it never considers more than k candidate merges per step; in contrast, Bottom-Up may consider a quadratic number of candidate merges per step and it is

slower than Fixed-Order as a consequence. Hybrid is in the middle for runtime as expected. Furthermore, $D = 3$ helps bound the size of C_ℓ and hence the cost of computing the set cover. The running time tends to decrease with bigger k for both Fixed-Order and Bottom-Up; the reason is that fewer merges are needed to reach the desired k . However, for Hybrid, since larger k makes the candidate pool larger and might bring in more calculation in the second phase (Bottom-Up phase), the run time for Hybrid tends to get closer to Bottom-Up.

The average value of Fixed-Order is lower than the value of Bottom-Up or Hybrid as explained in Section 5, although gets better with larger k in Figure 6b.

Effect of coverage parameter L Figure 6c shows that running time of all algorithms increase as the number of elements to be covered L increases. Since Fixed-Order depends linearly on L , it is less affected by L , whereas Bottom-Up treats individual elements as clusters and may incur quadratic time w.r.t. L . For Hybrid, with the restriction of the size of the candidate pool determined by k , the run time increase is slower than Bottom-Up and is comparable with Fixed-Order. Note that in Figure 6d, the upper bound decreases since with L increasing, the average value of the top- L elements decreases. All three algorithms seem to be close in terms of average values, but Bottom-Up has the highest value most of the times and Hybrid usually gets results close or equal to Bottom-Up.

Effect of distance parameter D In Figure 6e, Fixed-Order is mostly unaffected by D since the distance value is checked only once when an element is considered. Hybrid is relatively constant as well given that when the distance check starts, the number of unchecked tuples is limited by the candidate pool. For Bottom-Up as D increases, the run time drops first and then climbs. It may be caused by the existence of a balance point on number of calculations between distance insurance (phase 1 in Bottom-Up) and greedy merge (phase 2 in Bottom-Up).

The average value of the output (the value of objective function) is highest when $D = 1$ (since singleton clusters are collected for $L = k = 20$), then drops with D going up as shown in Figure 6f.

Effect of number of attributes m Varying the number of grouping attributes m also illustrates the effect of varying input data size. Since our algorithms run on the output of an aggregate query, as m increases, our input data size $|S| = n$ is likely to increase (for the m values in Figure 6g and 6h, the size of the input ranges from 140 to 280). When a new query comes, the system performs an initialization step of constructing clusters and the semi-lattice structure. This initialization time is shown in Figure 6g. This step is performed only once per query, varying k and D does not need another initialization. Our implementation takes from 10ms when $m = 4$ to about 1s when $m = 10$. Note that this is the number of group-by attributes in the top- k aggregate query, not in the original dataset. So it is likely to have a small value 10. Figure 6h has the running time of the algorithms for $k = L = 20$, $D = 3$ and shows that all the algorithms return results in real time (in a few ms) after the initialization step.

7.2 Cost and Benefit of Precomputation

The performance evaluation for precomputation is shown by varying k , L and D separately, and comparing the running time of Hybrid between precomputation implementation and non-precomputation (single) implementation.

Effect of size parameter k . In this experiment, $L = 1000$, $D = 2$ and $N = 2087$ are fixed. Five k values are chosen: 5, 10, 20, 50, 100. The running time result is shown in Figure 7a: the initialization time hardly changes with k growing since k does not affect the initialization process. Given that a larger k requires less operations in Bottom-Up phase to reach the target k , the running time for the algorithm (Hybrid) has a descending trend.

Effect of coverage parameter L . The fixed parameters are $k = 20$, $D = 2$, and $N = 2087$. Three L values are selected for the experiment: $L = 200, 500, 1000$. The running time results for single version and precomputation version are presented in Figures 7c and 7d. Both implementations have rising trend with respect to L and share similar initialization times as expected. Although under the same parameter combinations, algorithm runtime for single implementation is much lower than precomputation time in the other implementation (about 1/3 to 1/4), but the retrieval time for precomputation implementation is extremely short (tens of milliseconds), which can make up for the time in multiple runs.

Effect of total elements N . Here three parameters k , L , D are fixed as $k = 20$, $L = 500$ and $D = 2$. We varied total input elements to test the system's performance with relatively higher capacity: $N = 927, 2087$ and 6955 . The running time result is shown in Figures 7e and 7f. The changing trends are similar with those in Figures 7c and 7d, but a significant increase for the initialization time can be observed with N growing. This is caused by materializing more possible clusters brought by variety of tuples.

Single run vs. multiple runs. From Figure 7c, 7d, 7e and 7f, the information is enough for comparing precomputation and non-precomputation versions on both single run and multiple runs scenario - For a single run, precomputation process is unused, making precomputation version the slower and more expensive choice; For multiple runs with similar setup, the precomputation version has increasingly more benefits brought by the rapid retrieval process taking tens of milliseconds. In order to provide a quantitative comparison, we provide Figure 7b with $N = 6955$: if only a single run is required, the single version of Hybrid is clearly faster and cheaper than the precomputation version. However, when the third run finishes, the precomputation version is already faster; When all six runs finish, the single version takes about two times in terms of running time compared with the precomputation version.

Timing for Guidance Visualization. We evaluated the running time for the generation of guidance visualization under different queries. The generation times are similar among different number of attributes - 20–40 milliseconds when the number of attributes is from 4 to 10 with $N = 2087$ in MovieLens dataset, meeting the requirement for interactive performance.

7.3 Benefit of Optimizations

Cluster generation and mapping to tuples. Since L is the only factor that affects the initialization time when the input size N is fixed, in this experiment, L varies among 200, 500 and 1000 while others are fixed: $k = 20$, $D = 2$, $N = 2087$. The result is presented in Figure 8a. Only the running time of initialization is drawn because the optimizations in this section only affect the initialization time. The optimizations - cluster generation and cluster-tuple mapping - provide significant performance improvement by cutting down the running time from $> 100s$ for $L = 1000$ to 0.5s.

Delta Judgment. The effect by introducing *Delta Judgment* is shown in Figure 8b. Given that L is also the most effective variable to affect the running time, the experimental settings are the same as the experiment for Figure 8a. However, only the running time of the algorithm is plotted since *Delta Judgment* has no effect on the running time of initialization. The result in Figure 8b shows that the *Delta Judgment* successfully improves the algorithm's efficiency from 4.6s to 0.15s when $L = 1000$, which is the slowest case in the experiment in this section.

7.4 Scalability with a Larger Dataset

In order to evaluate the scalability of our algorithms we perform an experiment with TPC-DS dataset on *Store_Sales* table. The parameters are set to $k = 20$, $D = 2$ and $N = 47361$. Coverage parameter L varies among 500, 1000 and 2000. Both single and precomputation version are evaluated using this set of parameters. From the results shown in Figure 9a and Figure 9b, the initialization time is interactive - about 1s for the the largest parameters: $L = 2000$ and $N = 47361$. However, even for the single version, the running time of the algorithm increases to more than 1s compared with 200ms from results in Figures 7e and 7f, and for the precomputation version it increases to $\sim 2.5s$. Although the running time increases, the total running time ($\sim 3.5s$) for precomputation is still interactive. Note that the size of the answers (N) output by a query is likely to be much smaller than the size of the dataset, even for a big dataset.

8. USER STUDY AND SURVEY

We conducted a user study with the following high-level goals: (1) to compare our approach with an alternative that adapts decision trees [32] and (2) to evaluate the utility of user-specified parameters in our approach. Specifically, we want to know: (1) whether our new problem formulation provides any advantage over adapting existing methods to the same usage scenarios; and (2) whether allowing user-specified parameters in our problem formulation is warranted in order to capture the range of different usage scenarios and/or user preferences. In addition, we informally solicited feedback during the demonstration of our system at *SIGMOD* 2018 [43] to assess the effectiveness of our interactive features in Sections 6.

8.1 User Study Setup

Dataset and queries. All data are drawn from the *MovieLens RatingTable* as described in Section 7. Queries are based on the same aggregate query template introduced as in Example

1.1, with an additional WHERE condition and variations in query constants and group-by attributes across user tasks.

Adapted decision tree. As discussed in Section 2, no existing method suits our problem setting. After exploring various possibilities, we decided to adapt the method of decision trees [32] as it offers the closest match with our application scenarios. The structure of a decision tree naturally induces summaries of top- L tuples in the form of predicates, which are easier for users to interpret than other classifiers. It is also discriminative, as opposed to simply running clustering algorithms over the top- L tuples while ignoring low-value tuples. We use the standard implementation provided by Python’s scikit-learn package [30]; we tune the height of the decision tree such that the number of “positive” leaf nodes (wherein top- L tuples are the majority) as close as possible to, but no greater than, k . Note that the cluster patterns under this approach can be more complex than ours, as they may involve non-equality comparisons and negations. This additional complexity increases discrimination, but makes the patterns more difficult to interpret and internalize—a hypothesis we shall test with our study.

Tasks. Each study subject is asked to carry out three groups of tasks (*task groups*): (i) *varying-method*, (ii) *varying-k*, and (iii) *varying-D*. The first group is designed to compare our approach and decision trees. The last two are designed to evaluate the utility of making parameters k and D in our approach specifiable by users.⁶ To account for the possible learning effect, we sequence the task groups differently among study subjects—half go through the sequence varying-(method, k , D), while the remaining go through varying- (k , D , method).

Before each task group, we familiarize the subject with the aggregate query result as well as the tasks; Then, we give the subject a series of questions, organized into three sections in order. Each question asks the subject to classify a given tuple, whose value is hidden, into one of three categories: “top” (value among the top L of all tuples), “high” (value above or equal to the average, but outside the top L), and “low” (value below average). The three sections are based on the same “working set” of clusters, but differ in the information the subject can access:

- *Patterns-only*, 6 questions: The subject can see the clusters and their associated patterns, but not the membership within clusters or the table of all query result tuples. This section is designed to test how well the cluster patterns help users understand the data.
- *Memory-only*, 6 questions: The subject cannot access any information; all questions must be answered from memory. This section is designed to test the extent to which users can internalize the insights learned from the cluster patterns for later use. We ensure that these six tuples are distinct from those chosen before.

⁶We do not evaluate the utility of making L user-specifiable, as it should be evident that what “top” tuples mean depends on the situation—e.g., a small L means interest in characterizing really high-valued tuples, while a larger L means interest in tuples whose values are “good enough.”

- *Patterns+members*, 8 questions: The subject can see the clusters patterns as well as the covered result tuples. This section is designed to test how our full-fledged cluster UI can help user explore data. The 8 tuples are chosen and reordered randomly from the 12 tuples used in the previous two sections.

After these three sections, we present two sets of clusters: one is the working set, the other is obtained under a different setting (but for the same aggregate query and L) for comparison. We then ask the subject to choose a preferred set for the tasks just performed. For a *varying-method* task group, the cluster to compare is produced by decision trees, under the same k setting (D does not apply to decision trees); For a *varying-k* task group, the cluster to compare is produced by our approach under another k , while other parameters remain the same; For a *varying-D* task group, the cluster to compare is produced by our approach under another D .

Participants and assignment of tasks. There are 16 participants - 14 of them are graduate students at Duke University (12 in computer science and 2 others), while the remaining 2 are Duke undergraduates. They all have some prior experience working with tabular data and are capable of handling tasks in our user study.

Recall that each of the three task groups compares two sets of clusters. There are $2^3 = 8$ possible assignments in total. We assign two subjects to each of these 8 possibilities, each goes through one of the two task group sequences. Finally, we ensure that tuples in our questions are equally distributed among all subjects.

Metrics. We record the time for each subject to complete each of the three sections in each of the three task groups. We evaluate the accuracy of answers using the standard accuracy measure of $\frac{TP + TN}{TP + FP + FN + TN}$ based on confusion matrices [11], and we define two variants:

T-accuracy focuses on discerning the top tuples from the rest, where “positive” means being in top L ; *TH-accuracy* focuses on discerning the top and high tuples from the low ones, where “positive” means being in either top or high category.

8.2 User Study Results

Table 1 summarizes both the quantitative results (subjects’ performance in terms of time and accuracy for classifying tuples into categories) and qualitative results (subjects’ preferences between the clustering outputs compared) of our user study.

Varying-method task group. For this task group, we set $L = 50$, $k = 10$, $D = 1$ for our approach, and $L = 50$, $k = 10$ for the method based on decision trees. For this scenario, tree depth of 7 gives exactly 10 positive leaf nodes.

First, note that among the three sections, memory-only is the fastest, patterns-only is considerably slower, and patterns+members is the slowest. This observation holds both for our approach and for decision trees (as well as under each setting of other task groups). This universal trend can be intuitively explained by the fact that users tend to spend more time on a question if more information is presented to them.

As for accuracy, patterns+members has the highest accuracy, and patterns-only is usually no worse than memory-only. This trend also makes intuitive sense as users are generally able to achieve higher accuracy if aided with more information. Across settings, patterns+members is always nearly perfect, as expected.

Comparing our approach and decision trees in terms of time spent by study subjects, our approach is consistently faster over the three sections. The biggest advantage is seen in the patterns-only section, suggesting that our patterns are much easier to apply. The advantage is less pronounced in the other two sections. For patterns+members, a possible explanation is that users spend bulk of the time examining detailed memberships. For memory-only, our conjecture is that decision tree patterns are so difficult to recall that our subjects realized quickly that spending more time did not help.

In terms of accuracy, our approach is better than decision trees for the patterns-only and memory-only sections (recall that patterns+members is always nearly perfect across settings). It is understandable for decision trees to have lower TH-accuracy, because they are trained to separate only the top tuples from the rest, while our approach considers the values of all tuples covered by the patterns. On the other hand, while the T-accuracy for decision trees is good for patterns-only, it drops significantly for memory-only, because decision tree patterns are difficult for users to memorize. In comparison, the accuracy of our approach degrades very little from patterns-only to memory-only, which is evidence that users can internalize insights from our simple patterns very well.

Finally, when asked which method they prefer, the overwhelming majority of the subjects (14/16) chose our approach over decision trees. The key reason cited was the simplicity of our patterns.

Varying- k task group. In this task group, we fix $L = 30$ and $D = 1$, and compare $k = 5$ vs. $k = 10$. Note that with the bigger k , we expect to have more clusters with more specific patterns, leading to higher discrimination but more complex summaries.

From Table 1, we see that the bigger k leads to more time spent as long as patterns are accessible to the subjects, i.e., for patterns-only and patterns+members. However, for memory-only, the bigger k actually results in less time spent; one conjecture is that complex summaries are so difficult to recall from memory that some subjects simply stopped trying and resorted to guessing. This observation is consistent with the low accuracies seen under the bigger k for memory-only, further discussed below.

In terms of accuracy, favor turns from the smaller k to the bigger k for patterns-only and patterns+members, pointing to a clear tradeoff between time and accuracy. On the other hand, for memory-only, the trend is reversed: accuracies under the bigger k drop dramatically and become lower than under the smaller k , because the subjects had trouble recalling the summaries from their memory. In comparison, under the smaller k , accuracies for memory-only are at least as good as those for patterns-only.

Finally, when asked whether they prefer the smaller or bigger k , a slight majority of the subjects prefer the bigger, but still a significant fraction (7/16) prefer the smaller. There is no clear winner here, unlike the case for the varying-method task group.

Varying- D tasks. We fix $L = 10$ and $k = 7$, and compare $D = 1$ vs. $D = 3$. $D = 1$ represents a looser constraint, and in this case leads to detailed summaries and higher discriminative power; the trade-off, of course, is that patterns appear less diverse.

As we can see from Table 1, the bigger D leads to faster answer speed and higher accuracy in most cases, with just two exceptions: the smaller D is more accurate in terms of T-accuracy for patterns-only, and it is faster for patterns+members. Both can be explained by the fact that, here some clusters produced by the bigger D happen to have more general patterns and cover more tuples. Without access to cluster membership, T-accuracy would suffer because these clusters may cover some high-valued (but necessarily top-valued) tuples. With access to cluster membership, T-accuracy would not be a problem, but more tuples take longer to examine.

Although the performance results appear to favor the bigger D (looser constraint), preferences are divided. A majority of the subjects do prefer the bigger D , but still a sizable number of them (6/16) prefer the smaller D , which produces more diverse patterns.

Learning effect. We assess the possible learning effect by comparing the quantitative result within one experimental sequence (varying-method first, then varying- k and varying- D), and compare with the results in Table 1. The differences are minor, and the relative ordering of approaches by performance largely stays the same, so the conclusions drawn above still stand. Details are shown in Appendix A.9.

8.3 Informal User Survey Results

To measure the effectiveness of the interactive feature described in Sections 6, we asked attendees who visited our demo booth at *SIGMOD* 2018 to fill out an informal survey. We received 18 responses, and the results are summarized below:

Did you find the visualizations helpful?	Yes, very much	Yes	Not that much	Not at all
For parameter selection	4	13	1	0

The vast majority of the responses are positive. Some constructive criticisms were offered too. One pointed out that the visualization for guiding interactive parameter selection still required extensive explanation before users can understand and benefit from it. Another pointed out that instead of showing all choices of k and D in this visualization, it might be possible to use the data behind this visualization to narrow down the choices further.

8.4 Summary and Discussion

The high-level findings are: (1) our approach is more suitable to the designed tasks than the decision trees, thanks to the simplicity of our patterns by design; (2) while more specific and

detailed clusters can offer better accuracy, this advantage dissipates when users no longer see the cluster patterns directly, because they are much less memorable; (3) parameters k and D affect the complexity of our clustering results and present various trade-offs (e.g., accuracy vs. efficiency), so users have different preferences.

It is also worth noting that while we did not explicitly compare with the approach of simply showing the top L tuples with no summarization at all, which can be seen as an extreme case where $k = L$ and $D = 1$. Hence, the general observation we made when comparing parameter settings applies here too: showing the top L tuples alone would provide the most detailed information, but that would be very difficult to use and memorize.

9. CONCLUSIONS

In this paper, we presented a framework for summarization and exploration of high-valued aggregate query answers maintaining properties like diversity, small size, and coverage of original top elements. We studied optimization problems for these tasks, developed efficient algorithms and optimizations, evaluated the approach using real and benchmark datasets, and showed that our implementation is capable of interactive exploration in real time. We also conducted a user study showing that our constraints are useful and our results are preferred. There are several directions for future work. While we mainly focused on categorical attributes, for numeric attributes one can consider other distance functions (e.g., L_p norms) and description of clusters (e.g., ranges). Due to the presence of multiple competing constraints and complex objective function, obtaining formal approximation guarantees will be a challenging theoretical problem. One can also consider objective functions other than average and see how our algorithms (that can be adapted to other objectives) perform. Other sophisticated visualizations to better help the users can also be explored.

APPENDIX

A. APPENDIX

A.1 Proof of Proposition 4.2

PROOF OF PROPOSITION 4.2. Consider any cluster $C \in SC$ other than C_1 . Suppose $d(C_1, C) = \ell$ i.e., ℓ of the attributes contribute to the distance function. Fix such an attribute A . There are three possibilities. (1) $C_1[A] = C[A] = *$. If $C_2[A] = *$, it contributes 1 to $d(C_2, C)$. If $C_2[A] \neq *$, i.e., an attribute value, then also it contributes 1 to $d(C_2, C)$. (2) $C_1[A] \neq *$, $C[A] = *$, and $C_1[A] \neq C[A]$. If $C_2[A] = *$, it contributes 1 to $d(C_2, C)$. If $C_2[A] \neq *$, i.e., an attribute value, then it must be the same as $C_1[A]$, and therefore contributes to $d(C_2, C)$. (3) $C_1[A] = *$ and $C_1[A] \neq *$. Then $C_2[A] = C_1[A] = *$ and it contributes 1 to $d(C_2, C)$. (3) $C_1[A] \neq *$ and $C[A] = *$. Then either $C_2[A] = C_1[A]$ or, $C_2[A] = *$. In either case, it contributes 1 to $d(C_2, C)$. Summing over all A and considering all $C, \lambda' \neq \lambda$. $\square$

A.2 NP-hardness Proofs

THEOREM A.1. *The optimization problem for the objective of Max-Avg for the case when k and $D > 0$ is NP-hard.*

Proof. The reduction is from the problem of finding a minimum vertex cover in a tri-partite graph G with partitions (X, Y, Z) , which has shown to be NP-hard in [25]. The goal in this problem is to decide if the input graph has a vertex cover of size M , i.e., a subset of vertices $T \subseteq X \cup Y \cup Z$ and $|S| = M$ such that any edge in G has at least one endpoint in T ⁷. First we give the proof for Max-Avg. Suppose G has N_e edges and N_v vertices.

Given such a graph G and a bound M , we construct a database instance S as follows. There are three attributes A_X, A_Y, A_Z . (i) *Top- L tuples from the edges of G* : Any edge of the form (x, y) , $x \in X, y \in Y$ forms two tuples (x, y, Z_{xy}^1) and (x, y, Z_{xy}^2) , where Z_{xy}^1, Z_{xy}^2 are two unique values of the A_Z attribute for the edge (x, y) in G . Similarly, an edge (y, z) , $y \in Y, z \in Z$ forms two tuples (X_{yz}^1, y, z) , (X_{yz}^2, y, z) , and an edge (x, z) , $x \in X, z \in Z$ forms a tuple (x, Y_{xz}^1, z) , (x, Y_{xz}^2, z) . The weights of these tuples are 1. (ii) *Redundant tuples from the vertices of G* : For any vertex $x \in A_X$ in G , create a redundant tuple $(x, \gamma_x^2, \gamma_x^3)$. These redundant tuples have weight 0. Similarly, form redundant tuples for vertices $y \in A_Y$: $(\gamma_y^1, y, \gamma_y^3)$, and for vertices $z \in A_Z$: $(\gamma_z^1, \gamma_z^2, z)$ with weight 0. (iii) *More redundant tuples*: For each Z_{xy}^1 , form $N_r = 2 * N_e * N_v$ redundant tuples of the form $(-, -, Z_{xy}^1)$ with weight 0, where the positions with $-$ are filled with unique attribute values. Similarly, form redundant N tuples for each of $Z_{xy}^2, Y_{xz}^1, Y_{xz}^2, X_{yz}^1$, and X_{yz}^2 , placing these attribute values in their corresponding positions.

We set $k = M, L = 2 \times N_e, D = 3$. Note that only the tuples from the edges of the form $(x, y, Z_{xy}^1), (x, y, Z_{xy}^2), (x, Y_{xz}^1, z), (x, Y_{xz}^2, z), (X_{yz}^1, y, z), (X_{yz}^2, y, z)$ form the top- L original elements and have to be covered. We claim that G has a vertex cover of size M if and only if S has a solution, a set of clusters SC , of value $\frac{2N_e}{2N_e + M}$, where N_e is the number of edges in G .

(only if) Suppose G has a vertex cover T of size $M' \leq M$. For any $x \in X \cap T$, choose the cluster $(x, *, *)$ in SC ; similarly for $y \in Y \cap T$ and $z \in Z \cap T$, choose the clusters $(*, y, *)$ or $(*, *, z)$ respectively in SC . These clusters have mutual distance = 3, are incomparable, have size M , and cover all top- L elements. Each such cluster also covers a redundant element (with γ attribute value) of value 0. Therefore, the value of the solution is

$$\frac{2N_e \times 1 + M' \times 0}{2N_e + M'} \geq \frac{2N_e}{2N_e + M}.$$

(if) Suppose S has a solution SC of value $\frac{2N_e}{2N_e + M}$. Without loss of generality, any cluster in SC covers at least one of the top- L elements with value 1, otherwise it can be discarded without increasing the value of the average or size/distance/coverage of the solution (the

⁷[16] gives a reduction from the tri-partite vertex cover problem for *size-constrained weighted set cover* (given weights on the subsets, a size constraint k , a coverage fraction s , the goal is to return up to k sets that together contain at least sn elements and whose sum of weights is *minimal*). In contrast, in our setting, the weights are assigned on elements (not on subsets); the goal is to select at most k subsets with *maximum* value with the distance and other restrictions, that cover top- L original elements; and we show NP-hardness of the decision problem even if the elements are unweighted.

redundant elements have value 0 and can only reduce the average). Also note that the trivial solution $(*, *, *)$ cannot be chose, since it has value

$$\frac{2N_e + 0}{2N_e + N_v + 2N_e * N_r} \leq \frac{2N_e}{2N_e + 2N_e * N_r} = \frac{1}{1 + N_r} = \frac{1}{1 + 2N_e N_v}, \text{ which is strictly less than the}$$

$$\text{assumed value of } SC \geq \frac{2N_e}{2N_e + M}.$$

(A) *None of the chosen clusters in SC can be of the form $(*, *, Z_{xy}^1)$ (similarly for $Z_{xy}^2, Y_{xz}^1, Y_{xz}^2, X_{xz}^1$).* Suppose one cluster in SC is $(*, *, Z_{xy}^1)$. Then it covers N_r redundant tuples that are not covered by any other cluster in SC. Suppose SC has N' redundant tuples all together from all other clusters. Then the average value of SC is

$$\frac{2N_e + 0}{2N_e + N' + N_r} \leq \frac{2N_e}{2N_e + N_r} = \frac{2N_e}{2N_e + 2N_e N_v} = \frac{1}{1 + N_v}, \text{ which is strictly less than } \frac{2N_e}{2N_e + M}, \text{ the}$$

assumed value of SC since $M \geq N_v$.

Next we argue that each cluster in SC can have exactly two *-s, combining with (A) above, must be of the form $(x, *, *)$, $(*, y, *)$, $(*, *, z)$.

(B) *The clusters in SC cannot have zero *-s.* Suppose without loss of generality that for a top-L tuple (x, y, Z_{xy}^1) , the singleton cluster (x, y, Z_{xy}^1) has been chosen in SC. Due to the incomparability condition, none of $(x, y, *)$, $(x, *, *)$, $(*, y, *)$ can belong to SC. Hence, to cover the other top-L tuple (x, y, Z_{xy}^2) , one of (x, y, Z_{xy}^2) , $(x, *, Z_{xy}^2)$, $(*, y, Z_{xy}^2)$ has to be chosen $(*, *, Z_{xy}^2)$ cannot be chosen due to (A) above. However, these three clusters have distance 1, 2, 2 respectively from (x, y, Z_{xy}^1) violating the distance constraint $D = 3$.

(C) *The clusters in SC cannot have one *-s.* (i) Suppose for a top-L tuple (x, y, Z_{xy}^1) , the cluster $(x, y, *)$ has been chosen in SC. Due to the incomparability condition, none of $(x, *, *)$, $(*, y, *)$ can belong to SC, and any cluster with 1 or zero * to cover top-L tuples from edges of the form (x, y') or (x', y) in G will have distance ≥ 2 with $(x, y, *)$, violating $D = 3$. If $(x, *, Z_{xy}^1)$ is chosen (same for $(*, y, Z_{xy}^1)$), to cover the other top-L tuple (x, y, Z_{xy}^2) , one of (x, y, Z_{xy}^2) , $(x, *, Z_{xy}^2)$, $(*, y, Z_{xy}^2)$ has to be chosen. Since the first two have distance 1 and 2 respectively from $(x, *, Z_{xy}^1)$ violating the distance constraint $D = 3$, $(*, y, Z_{xy}^2)$ must also belong to SC. However, $(x, *, Z_{xy}^1)$ and $(*, y, Z_{xy}^2)$ together rule out covering clusters for other top-L tuples from edges of the form (x, y') or (x', y) in G , which will have distance ≥ 2 from either of these two clusters, violating $D = 3$.

Hence the clusters in SC must be of the form $(x, *, *)$, $(*, y, *)$, or $(*, *, z)$, which corresponds to a vertex cover in G . Suppose there are K clusters in SC. Then the value of SC is $\frac{2N_e + 0}{2N_e + K} \geq \frac{2N_e}{2N_e + M}$ by our assumption, hence $K \leq M$, and G has a vertex cover of size at most M . $\square$

THEOREM A.2. The decision version of whether a feasible non-trivial solution exists for the problem defined in Definition 4.1 is NP-hard, even with three attributes, $D = 0$, $L = n$, and uniform weights of the elements⁸.

Proof Proof of Theorem A.2. The reduction is again from the problem of finding a minimum vertex cover in a tri-partite graph G with partitions (X, Y, Z) (see the proof of Theorem A.1).

Given such a graph G and a bound M , we construct a database instance S as follows. There are three attributes A_X, A_Y, A_Z . Any edge of the form (x, y) , $x \in X, y \in Y$ forms a tuple (x, y, Z_{xy}) , where Z_{xy} is a unique value of the A_Z attribute for the edge (x, y) in G . Similarly, an edge (y, z) , $y \in Y, z \in Z$ forms a tuple (X_{yz}, y, z) , and an edge (x, z) , $x \in X, z \in Z$ forms a tuple (x, Y_{xz}, z) . The total number of tuples in the relation S is n and each tuple has the same weight 1. We set $k = M$, $L =$ the number of edges in G , and claim that G has a vertex cover of size M if and only if S has a non-trivial solution, a set of clusters SC , of size k .

(only if) Suppose G has a vertex cover T of size M . For any $x \in X \cap T$, choose the cluster $(x, *, *)$ in SC ; similarly for $y \in Y \cap T$ and $z \in Z \cap T$, choose the clusters $(*, y, *)$ or $(*, *, z)$ respectively in SC . Since T is a vertex cover, SC will cover all tuples in S and has size $M = k$.

(if) Suppose S has a non-trivial solution SC of size $k = M$. The clusters in SC can have $*$ in zero, one, or two positions (all three positions cannot be $*$ since SC is a non-trivial solution). We will argue that any cluster in SC can be replaced by a cluster with two $*$ and a vertex from G forming another feasible non-trivial solution without increasing the size of SC (the size may decrease). Consider any tuple of the form $t = (x, y, Z_{xy})$ in S (the other two cases (x, Y_{xz}, z) and (X_{yz}, y, z) follow similarly). If t is covered by a cluster of the form (x, y, Z_{xy}) , $(x, y, *)$, $(x, *, Z_{xy})$, $(*, y, Z_{xy})$, or $(*, *, Z_{xy})$, such clusters cannot cover any other tuple in S since Z_{xy} is unique, so replace such clusters by $(x, *, *)$ or $(*, y, *)$. After this is repeated for all tuples in S , the only types of clusters remaining in SC be of the form $(x, *, *)$, $(*, y, *)$, or $(*, *, z)$, which corresponds to a vertex cover in G , and the size of the solution has not increased. $\square$

A.3 Architecture and GUI

The architecture of the system is shown in Figure 10. The interface on the browser sends HTTP requests to the server based on the actions of the user. For different requests, the server uses the router module to determine which action in the controller module will be performed. After an action processes the request, it sends the result to the view module. Then the view module partially renders the page and sends it back to the browser.

The interface consists of two parts: a *tool panel* on the left, and a *working space* on the right. When the user opens the GUI, it connects to the DBMS through the router and the controller modules. The user can see all the databases in the DBMS, and get a preview of any table in any database using the tool panel to have an idea on the schema and data. Then the user enters an aggregate SQL query (in the query box), the parameters k, L, D (either in the query

⁸The top- k original elements may not constitute an optimal solution for $D = 0$ when $L > k$.

box or in their respective boxes), and clicks on ‘go’. Then the *SQL result* section below the SQL input area is updated with the output.

The action for computing the output clusters has three stages. *First*, it retrieves the original SQL result from PostgreSQL. Then it initializes or updates the cache, which stores some data structures that are frequently used while computing the clusters. If the underlying aggregate query is new, the system will fully update the cache. If only parameters k , L , or D are changed, the system may partially update the cache. The system uses both memory and PostgreSQL as cache. *Second*, based on the input parameters, the system selects an algorithm to compute and optimize the result. *Third*, the system assembles the result as a JSON string and sends it back to the browser. Finally, the SQL result section on the GUI is updated without refreshing this page.

A.4 Omitted Pseudocodes from Sections 5.2

Algorithm 3 describes details of the Fixed-Order algorithm introduced in Section 5.2

Algorithm 3 The Fixed-Order algorithm
Require: Size, coverage, and distance constraints k, L, D
1: $\mathcal{O} = \emptyset$.
2: for $i = 1$ to L do
3: Let t_i be the i -th top element.
4: if $t_i \in \text{cov}(\mathcal{O})$ then
5: $\text{/* } t_i \text{ is already covered, do nothing. */}$
6: else if $ \mathcal{O} < k$ then
7: if t_i is at distance $\geq D$ from all clusters in $\mathcal{O}$ then
8: $\mathcal{O} = \mathcal{O} \cup \{t_i\}$.
9: else
10: Let P_D be the pairs of clusters $(C, \{t_i\})$ where $C \in \mathcal{O}$ such that the distance between C, t_i is $< D$.
11: Perform $\text{UpdateSolution}(\mathcal{O}, P_D)$.
12: end if
13: else/* merge t_i with one of the clusters in $\mathcal{O}$ */
14: Let P be the pairs of clusters $(C, \{t_i\})$, $\forall C \in \mathcal{O}$.
15: Perform $\text{UpdateSolution}(\mathcal{O}, P)$.
16: end if
17: end for
18: return $\mathcal{O}$

A.5 Qualitative Evaluation with Other Related Approaches

For the query in Example 1.1, we here compare the results adapting approaches in three related papers [24, 31, 8] that are most relevant to this work (*i.e.*, consider either summarization or diversification). Comparison with the MMR-based λ -parameterized diversification framework (no summarization) from [41] is also presented in this section. We use the brute-force or an efficient algorithms from these papers using $k = 4$, $D = 2$, and $L = 10$. As expected, the objectives proposed in these papers do not serve the purpose of summarizing aggregate query answers that we study in this paper.

A.5.1 Comparison with smart drill-down [24]—The goal in [24] is to find an ordered set of k rules (clusters with $*$ in our framework) R with maximum score $\text{score}(R) = \sum_{r \in R} \text{MCount}(r, R) \times W(r)$, where the marginal count $\text{MCount}(r, R)$ denotes the number of tuples covered by r but not covered by preceding rules in R , and $W(r)$ denotes the number of non- $*$ attributes in r . The two parameters focus on diversity (by preferring rules with high

MCount) and goodness (by preferring more specific rules) of rules. To compare it with our framework with *relevance* (summarizes aggregate results where tuples have values), we update the scoring function as $\text{score}(R) = \sum_{r \in R} \text{MCount}(r, R) \times W(r) \times \text{val}(r)$, where val denotes the average value of the elements covered by r that are uncovered so far by previous rules. Our framework also considered *coverage* of top- L elements, so we evaluated a greedy algorithm (shown to perform well in [24]) on two inputs: (i) all elements in the aggregate query results, and (ii) top- L elements and obtained the following results:

hdecade	agegrp	gender	occupation	avg score
Smart drill-down on top-10 elements				
*	20s	M	*	3.47
*	10s	M	Student	3.63
1995	30s	F	Educator	3.70
Smart drill-down on all elements				
1995	*	M	*	3.18
*	20s	M	*	3.47
1995	*	F	*	3.24
*	10s	M	Student	3.63

Comparing the above tables with Figure 1b and 1c, we see that the average score of the rules or clusters is much less than our output. Although the above rules capture (20, M) that is prevalent in the top-10 results in Figure 1a, it is not a characteristics of the top-valued elements; *e.g.*, elements ranked 44, 46, 49 (in total 18 out of 50 tuples) satisfy this rule leading to a low score.

A.5.2 Comparison with diversified top-k [31]—Given the elements with scores, using our terminology, the goal in [31] is to find at most k elements such that the distance between any two chosen elements is at least D and the sum of scores is maximized. This notion considers *diversity* and *relevance*. To also include coverage like ours, we ran it on top- L elements and got the following answers (by a brute-force implementation since the goal was qualitative evaluation). We show both actual value of the representative elements (score) as well as the average value of elements within distance $D - 1$ of these chosen elements (avg score) below since intuitively the chosen elements cover these elements.

hdecade	agegrp	gender	occu.	score	avg score
Diversified top- k on top-10 elements					
1975	20s	M	Student	4.24	3.71
1980	20s	M	Programmer	4.13	3.77
1980	10s	M	Student	3.97	3.69
1995	30s	F	Educator	3.70	3.52

Although [8] optimizes for a different optimization goal, the above results illustrate that it does not perform well for our goal of summarization of top-valued elements with diversity. *First*, the average values of the “clusters” formed by the representative elements shown above are less than the values given by our approach in Figure 1b. These chosen values include the original top-10 elements within distance 1, but also include many elements with low values within such circle. For instance, the first element above covers 6 elements (including itself) and covers the 28th element with value 3.12. *Second*, the representative of a cluster is one of the original elements, and we are not getting summarized common properties of the cluster using *-attribute values.

A.5.3 Comparison with DisC diversity [8]—Given the distance parameter D and a set of elements P , the goal of [8] is to find a *DisC diverse subset* S^* of minimum size such that each element in P is at most distance D from some element in S^* , and no two elements in S^* are within distance D of each other. Like [31], this diversity notion naturally includes summarization, since an element can be assigned to the *cluster* corresponding to a point in S^* at distance D . To also include the notion of *relevance* and *coverage of top- L* , we ran it too on top- L elements (a brute-force implementation) and obtained the following results. Like [31], [8] gives clusters with smaller scores than ours (e.g., the first tuple “covers” eight elements where the last one is of rank 28 and has value 3.31), and do not exhibit the common properties by * values.

hdecade	agegrp	gender	occu.	score	avg score
DisC diversity on top-10 elements					
1980	20s	M	Student	3.91	3.81
1985	10s	M	Student	3.76	3.66
1995	30s	F	Educator	3.70	3.52
1985	20s	M	Engineer	3.65	3.62

A.5.4 Comparison with MMR-based λ -parameterized approach [41]—Running an MMR-based λ -parameterized approach in [41], we obtain the following results for different values of λ . Note that this is a result diversification problem and not a result summarization problem, therefore does not include a coverage or summary of the top answers, or an average score. Below, higher λ denotes higher diversity.

hdecade	agegrp	gender	occu.	score
$\lambda = 0$				
1975	20s	M	Student	4.24
1980	20s	M	Programmer	4.13
1980	10s	M	Student	3.96
1980	20s	M	Student	3.91
$\lambda = 0.2, 0.5, 0.8$				

1975	20s	M	Student	4.24
1980	20s	M	Programmer	4.13
1980	10s	M	Student	3.96
1995	30s	F	Educator	3.70
$\lambda = 1.0$				
1985	20s	M	Programmer	3.86
1980	20s	M	Engineer	3.83
1985	10s	M	Student	3.76
1995	30s	F	Educator	3.70

A.6 Extension for range values

We introduced *don't care* value (*) in Section 3. However, for numerical attributes (such as age) and text attributes (such as date) with given hierarchies, ranges values can be an interesting option to present. Aiming at this goal, we set up a tree structure to represent the given hierarchy. Examples of trees for numerical attributes and text attributes are shown in Figure 11 and Figure 12. Individual values of the attribute serve as leaves, and possible ranges serve as nodes of the tree structure.

To get the proper range required for building and calculating clusters, we just need to get the LCA of candidate leaves and nodes in the tree structure. The LCA node is unique for certain candidates and can be found directly in the tree. For example, for the example given in Figure 11, to get the union of [20,40) and 55, we need to find the LCA node of them, which is [20, 60). A $\log(n)$ algorithm[18] is available to find the LCA node, where n is the total number of nodes.

Our framework can currently handle a given concept hierarchy in the form of a tree. How to automatically build such concept hierarchies will be an orthogonal future research direction.

A.7 Details for Visualizing Successive Changes

Figure 1 shows how a user can inspect the clusters and the elements they contain in the form of tables in our framework. However, for interactive exploration, it is also important that the user can see how the old solution changes to a new solution (if k , D , or L is updated). When the user explores the solution space updating an input parameter, in some scenarios the solution can change marginally, whereas in others it can change drastically. To help the user understand how two consecutive solutions compare with each other, our framework produces a visualization showing the old and new solution, and how the tuples in these clusters are redistributed (also the size of the cluster, the fraction of top- L tuples contained in them, etc.). For example, Figure 13 shows that if $k = 4$ in Example 1.2 is changed to $k = 3$, then two of the clusters will merge to form the new solution.

To support this comparison, we display the successive solutions and their overlaps (Section A.7.1), and formulate it as an optimization problem for a clean display (Section A.7.2).

A.7.1 Visualizing Changes—An example visualization is shown in Figure 14. Each box corresponds to a cluster in the solution. The left hand side boxes (in green) are result clusters for the previous run, and the right hand side boxes (in yellow) are clusters under the current parameters. The width of each box is proportional to the number of tuples contained in the cluster. Clusters connected with bands or ribbons contain shared tuples. The thicker a band is in the middle, the more common tuples it contains⁹. The parts of the boxes in darker color correspond to the fraction of top- L tuples contained in these clusters. Hovering the mouse over different regions shows details.

A.7.2 Optimizing Cluster Placement—A careful ordering of the cluster boxes helps in displaying a cleaner visualization comparing two consecutive executions. For instance, Figure 15 is an example where placement of the clusters (boxes) leads to more *crossing* of the bands, whereas Figure 14 shows a better placement. The placement can have more effect when the number of clusters (the value of k) is larger. Therefore, we formulate the cluster placement problem as an optimization problem that intends to minimize crossing of the bands as follows.

Optimization problem. Let $\mathcal{O}_a$ and $\mathcal{O}_b$ denote the old and new set of clusters respectively in two consecutive runs containing cluster $\mathcal{O}_a = \{c_{a1}, c_{a2}, \dots, c_{am}\}$ and $\mathcal{O}_b = \{c_{b1}, c_{b2}, \dots, c_{bn}\}$.

Let m_{ij} denote the number of shared tuples between clusters c_{ai} and c_{bj} , and M denote the set of all such m_{ij} values. We assume that the first clusters in both sides are placed at the same vertical position, and there are no gaps or overlaps between two adjacent clusters on either side. Consider two orderings of clusters on the left hand side and right hand side respectively in terms of their starting positions $P_a = P_{a1}, P_{a2}, \dots, P_{am}$ (a permutation of $[0, m - 1]$), and $P_b = p_{b1}, p_{b2}, \dots, p_{bn}$ (a permutation of $[0, n - 1]$). We define a weighted *earth mover's distance*[34] d_{ij} between one left cluster c_{ai} and one right cluster c_{bj} to evaluate the amount of crossing due to a single band (from c_{ai} to c_{bj}) as:

$$d_{ij} = m_{ij} \times |p_{ai} - p_{bj}|$$

Since $\mathcal{O}_a$ is the previous cluster set, P_a is fixed and given. The goal of the optimization problem for this visualization is to output a good ordering P_b for $\mathcal{O}_b$, and is formulated as follows:

Definition A.3. Optimization for placement of clusters. Given old and new clusters $\mathcal{O}_a, \mathcal{O}_b$, their overlaps M , and ordering P_a of the clusters on the left hand side, find an ordering P_b of the clusters on the right hand side that minimizes $D = \sum_{i=1}^m \sum_{j=1}^n d_{ij}$.

⁹This visualization shows some similarity with SANKEY diagrams that are widely used in the field of energy and material flow management[26]. However, popular SANKEY diagram libraries (e.g., d3-sankey) focus more on managing placement among columns (horizontal positioning). For vertical positioning, they do multiple iterations to re-position objects to achieve a satisfying visualization. We do not need to consider multiple columns, and our visualization focuses on vertical positioning and ordering of objects

Optimal solution using bipartite matching. The above optimization problem can be reduced to the minimum cost perfect matching problem in a complete bipartite graph as follows. We form a weighted complete bipartite graph $G(U \cup V, E)$, where the n nodes in U correspond to the n clusters in $\mathcal{O}_b$, and the n nodes in V correspond to the positions $1 \cdots n$.

An edge (u, v) denotes the possibility when cluster c_{bu} is placed in position $v \in [1, n]$. The weight of the edge (u, v) is the cost $\sum_{i=1}^m d_{iu} = \sum_{i=1}^m (m_{iu} \times |p_{ai} - (v-1)|)$ (if c_{bu} is placed in position v , there will be $v-1$ clusters before it, and its position will be $v-1$), *i.e.*, the total contribution of cluster c_{bu} in the optimization objective in Definition A.3 if it is placed in position $v \in [1, n]$. Since $\mathcal{O}_a$ and P_a are given and fixed, this weight can be computed in polynomial time as a precomputation step for each cluster in $\mathcal{O}_b$ and each position. A matching gives a positioning P_b on the clusters in $\mathcal{O}_b$, and the minimum cost perfect matching, which has a polynomial time algorithm[14], gives an optimal solution to our optimization problem.

We also studied an alternative formulation of the above optimization problem, where instead of the *width* of the boxes being proportional to the number of tuples in clusters, the *height* is proportional to the number of tuples, *i.e.*, the starting positions P_a and P_b are no longer permutations of $[0, m-1]$, $[0, n-1]$ but also depend on the height of the individual clusters. However, we found that this variant is NP-hard by a reduction from the *earliness-tardiness job scheduling problem*[13]. The proof and details of this alternative formulation are deferred to the extended version of this paper.

A.7.3 Performance of Comparison Visualization—We tested the running time for calculating and generating the visualization for both the applied algorithm [14] and the brute-force algorithm under $k = 10$, $L = 15, 20$ and $D = 2$ in Movielens dataset with $N = 2087$. In this test, both algorithms have similar figure drawing time (20ms) since they have identical data for figure generation (both of them get the optimal answer), but difference between the calculation time is enormous—the bipartite matching algorithm takes less than 10ms while brute-force takes more than 2s.

The quality of visualizations produced by bipartite matching is shown in Figure 16a and Figure 16b, as the “matched visualization,” in comparison to the “default visualization.” For the default visualization, we use the sequences of clusters as returned by successive runs of the clustering algorithm, where clusters for both sides are ordered by value. Parameter sets are $D = 2$, $(k, (L_1, L_2)) = (5, (8, 10)), (10, (15, 20))$ and $(20, (30, 40))$ where L_1 and L_2 are L s for the two answer sets. Figure 16a shows that bipartite matching is very effective in reducing the “clutter” in visualization, as measured by our distance metric in Section A.7.2 (note that the distances are generally not comparable among different k s). In addition to this metric, we also counted the number of crossings among bands (connections between left and right clusters) and plot the result in Figure 16b. It is clear that our approach also succeeds in cutting down the amount of crossings.

A.7.4 Informal User Study—As part of the informal user study we conducted at SIGMOD 2018 mentioned in Section 8, the satisfactory towards this comparison view is shown as below.

Did you find the visualizations helpful?	Yes, very much	Yes	Not that much	Not at all
For parameter selection	4	13	1	0
For comparing old/new clusters	7	11	0	0

It suggests that the comparison view is helpful and useful in real life scenario. One constructive suggestion upon this view is that it would be better if the user can view tuples inside each cluster and band after clicking.

A.8 Aggregate Queries Used in Experiments

The aggregate queries we use for the MovieLens dataset has the following form:

```
SELECT ⟨grouping attributes⟩, avg (rating) as val
FROM RatingTable
GROUP BY ⟨grouping attributes⟩
HAVING count (*) > 50
ORDER BY val DESC
```

The aggregate queries we use for TPC-DS related experiments share the same form with the example given above.

```
SELECT ⟨grouping attributes⟩, cast (avg (net_profit) as int) as val
FROM store_sales
GROUP BY ⟨g rouping attributes⟩
HAVING count (*) > 10
ORDER BY val DESC
```

A.9 Detailed User Study Setup

Dataset and queries. All data are drawn from the *MovieLens RatingTable* as described in Section 7. Queries are based on the same aggregate query template introduced therein, with an additional WHERE condition and variations in query constants and group-by attributes across user tasks;

Adapted decision tree. As discussed in Section 2, no existing method suits our problem setting. After exploring various possibilities, we decided to adapt the method of decision

trees [32] as it offers the closest match with our application scenarios. The structure of a decision tree naturally induces summaries of top- L tuples in the form of predicates, which are easier for users to interpret than other classifiers. It is also discriminative, as opposed to simply running clustering algorithms over the top- L tuples while ignoring low-value tuples. Finally, it is possible to control the complexity of the tree. We used the standard decision tree implementation provided by Python’s scikit-learn package [30]; given k , the maximum number of clusters to produce, we tune the height parameter of the decision tree such that the number of “positive” leaf nodes (wherein top- L tuples are the majority) as close as possible to, but no greater than, k .

Note that the cluster patterns under this approach can be more complex than ours, as they may involve non-equality comparisons and negations. This additional complexity increases the discriminative power, but makes the patterns more difficult for users to interpret and internalize—a hypothesis that we shall test with our study.

Tasks. Each study subject is asked to carry out three groups of tasks: the *varying-method group*, *varying- k group*, and *varying- D group*. The third first group is designed to compare our approach and decision trees. The last two are designed to evaluate the utility of making parameters k and D in our approach specifiable by users.¹⁰ To account for the possible effect of users learning and getting better with our approach, we sequence the task groups differently among study subjects—half of them go through the sequence (varying-method, varying- k , varying- D), while the remaining half go through (varying- k , varying- D , varying-method).

All tasks within one group are based on the same aggregate query. Before beginning the task group, we familiarize the subject with the aggregate query and result as well as the tasks to perform; we also show all query result tuples in a table, with top L tuples highlighted for convenience. Then, we remove the table of all query result tuples, and give the subject a series of questions, organized into three sections, in order. Each question asks the subject to classify a given tuple, whose value is hidden, into one of three categories: “top” (the tuple is among the top L of all query result tuples), “high” (the tuple has value above or at the average among all tuples, but is outside the top L), and “low” (the tuple has below-average value). The three sections are based on the same “working set” of clusters, but differ in what information the subject can access when answering questions:

- *Patterns-only*, 6 questions: When answering these questions, the subject can see the clusters and their associated patterns, but not the membership within clusters or the table of all query result tuples. This section is designed to test how well the cluster patterns help users understand the data. The 6 tuples to be classified are chosen randomly and evenly across the top, high, and low categories, and are ordered randomly. We do not reveal to the subject how these tuples are distributed among the three categories, as with questions in other sections below.

¹⁰We do not evaluate the utility of making L user-specifiable, as it should be evident that what “top” tuples mean depends on the situation—e.g., a small L means the user is interested in characterizing really high-valued tuples, while a larger L may mean the user is interested in tuples whose values are “good enough.”

- *Memory-only*, 6 questions: The subject can see neither the clusters or the table of all query result tuples; all questions must be answered from memory. This section is designed to test the extent to which users can internalize the insights learned from the cluster patterns for later use. The 6 tuples are chosen in the same way as in the patterns-only section, but we ensure that they are distinct from those chosen before.
- *Patterns+members*, 8 questions: The subject can see the clusters, their associated patterns, as well as the result tuples they cover; but the table of all query result tuples remains inaccessible. This section is designed to test how our full-fledged cluster UI can help user explore data. The 8 tuples are chosen and re ordered randomly from the 12 tuples used in the previous two sections, such that 4/2/2 are from the top/high/low categories, respectively.

After these three sections are done, to conclude the task group, we present two sets of clusters side-by-side: one is the working set that the subject has been using, and the other one is obtained under a different setting (but for the same aggregate query and same L) for comparison. We then ask the subject to choose which set of clusters would be preferred for the tasks just performed.

- For a *varying-method* task group, the clusters to compare are produced by our approach (using Hybrid) and by the method of decision trees, under the same k setting (D does not apply to decision trees).
- For a *varying-k* task group, the clusters to compare are produced by our approach under two different k settings, while other parameters remain the same.
- For a *varying-D* task group, the clusters to compare are produced by our approach under two different D settings, while other parameters remain the same.

Participants and assignment of tasks. There are 16 participants in total. 14 of them are graduate students at Duke University (12 in computer science and 2 others), while the remaining 2 are Duke undergraduates. They have varying degrees of knowledge about databases and SQL language, but all have some prior experience working with tabular data and are capable of handling all tasks in our user study.

Recall that each of the three task groups compares two sets of clusters. While every subject sees both sets at the end of the task group, the questions earlier in the task group are based on one working set chosen between the two. There are a total of $2^3 = 8$ possibilities for assigning working sets to the three task groups. We assign two subjects to each of these 8 possibilities. As discussed earlier, to account for the learning effect, we make one of these subjects go through the sequence (varying-method, varying- k , varying- D) and the other (varying- k , varying- D , varying-method). Finally, we ensure that tuples used in our questions appear equal number of times over tasks across all subjects.

Metrics. We record the time it takes for each subject to complete each of the three sections in each of the three task groups. We evaluate the accuracy of answers to the questions using the standard accuracy measure of $\frac{TP + TN}{TP + FP + FN + TN}$ based on confusion matrices [11], but

we define two variants: *T-accuracy* focuses on the ability to discern the top tuples from the rest, where “positive” means being in top L ; *TH-accuracy* focuses on the ability to discern the top and high tuples from the low ones, where “positive” means being in either top or high category.

A.10 Detailed Analysis for Learning Effect in User Study

We give some brief conclusions on learning effect in Section 8 in the main paper. In this subsection, we present the quantitative result within one experimental sequence (varying-method first, then varying- k and varying- D) in Table 2. Our conclusion is that the learning effect does not have a huge impact on the leadership within each task group.

The learning effect takes place with the time for each question - since users in Table 2 work with varying-method first, it takes 20% more time for patterns-only questions; when it comes to varying- D , since users in this group are already familiar with tasks, it takes slightly less time in all three tasks.

A.11 Additional Related Work

In addition to the papers discussed in Section 2, diversification of query results has been extensively studied in the literature for both query answering in databases and other applications [4, 2, 17, 49, 45, 15, 44, 10, 48, 33, 3, 1, 36, 7, 41, 31, 8]. One of the formalisms to capture both *diversity and relevance* of a resultset is to balance these two objectives using a trade-off parameter λ specified by the user. This approach, called *MMR (Maximal Marginal Relevance)* aims to reduce redundancy while maintaining relevance of the chosen outputs for the input query, and was first used for re-ranking retrieved documents and in selecting appropriate passage for text summarization [4]. Gollapudi and Sharma [17] studied three variants of the objective function (max-sum, max-min, mono) based on the MMR criterion. Deng and Fan [7] studied the data complexity and combined complexity for these problems. Vieira et al. [41] conducted an experimental study of existing and new algorithms for the max-sum objective defined in [17] with some small modifications. Fraternali et al. [12] studied this objective for diversification of objects in a low-dimensional vector space. We compare results from Vieira et al. [41] with our work in Appendix A.5.4.

The *diversity* criterion has been studied algorithmically as the *facility dispersion problem* [33]. [3] studied the *max-sum dispersion problem* and the *max-sum diversification problem* (as in [17]) when the value of a subset of elements $w(S')$ is given by a monotone submodular function. Abbassi et al. [1] studied the diversity maximization of a set of points under matroid constraints. *Diverse skyline* [36] is another related direction.

Zhu et al. [48] proposed a ranking algorithm with applications in text summarization and social network analysis. Vee et al. [40] studied the problem of computing diverse query results for non-aggregate queries in online shopping applications. In the area of Information Retrieval (IR), Zheng et al. [47] studied search result diversification. using λ -parameterized MMR objective function [17, 41], but their “diversity score” is defined as the sum (over possible topics) of product of importance of a subtopic to the input query and how much a document covers this topic.

Chen and Li [5] considered the problem of categorizing query answers using clusters on a navigational tree by exploiting the query history of the users when different users have diverse preferences. Other approaches include relational data summarization [46] and Web table search taking into account schema/instance diversity, table popularity, and redundancy [29]. For result summarization and exploration in databases, Gebaly et al. [9] considered summarization of attributes using * values to find factors that affect a binary (non-aggregate) attribute. Sarawagi explored (e.g., [35]) sophisticated OLAP operators for helping the user visit unvisited interesting parts in a data cube.

10. REFERENCES

- [1]. Abbassi Z, Mirrokni VS, and Thakur M. Diversity maximization under matroid constraints. In KDD, pages 32–40, 2013.
- [2]. Agrawal R, Gollapudi S, Halverson A, and Jeong S. Diversifying search results In Proceedings of the second ACM international conference on web search and data mining, pages 5–14. ACM, 2009.
- [3]. Borodin A, Lee HC, and Ye Y. Max-sum diversification, monotone submodular functions and dynamic updates. In PODS, pages 155–166, 2012.
- [4]. Carbonell J and Goldstein J. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In SIGIR, pages 335–336, 1998.
- [5]. Chen Z and Li T. Addressing diverse user preferences in sql-query-result navigation. In SIGMOD, pages 641–652, 2007.
- [6]. Cormen TH. Introduction to algorithms. MIT press, 2009.
- [7]. Deng T and Fan W. On the complexity of query result diversification. ACM TODS, 39(2):15:1–15:46, 5 2014.
- [8]. Drosou M and Pitoura E. Disc diversity: result diversification based on dissimilarity and coverage. PVLDB, 6(1):13–24, 2012.
- [9]. El Gebaly K, Agrawal P, Golab L, Korn F, and Srivastava D. Interpretable and informative explanations of outcomes. PVLDB, 8(1):61–72, 9 2014.
- [10]. Fan W, Wang X, and Wu Y. Diversified top-k graph pattern matching. PVLDB, 6(13):1510–1521, 2013.
- [11]. Fawcett T. An introduction to roc analysis. Pattern recognition letters, 27(8):861–874, 2006.
- [12]. Fraternali P, Martinenghi D, and Tagliasacchi M. Top-k bounded diversification. In SIGMOD, pages 421–432, 2012.
- [13]. Garey MR, Tarjan RE, and Wilfong GT. One-processor scheduling with symmetric earliness and tardiness penalties. Mathematics of Operations Research, 13(2):330–348, 1988.
- [14]. Geomans MX. Lecture notes on bipartite matching. Massachusetts Institute of Technology, 2009.
- [15]. Gkorgkas O, Vlachou A, Doukeridis C, and Nørsvåg K. Finding the most diverse products using preference queries. In EDBT, pages 205–216, 2015.
- [16]. Golab L, Korn F, Li F, Saha B, and Srivastava D. Size-constrained weighted set cover. In ICDE, pages 879–890, 2015.
- [17]. Gollapudi S and Sharma A. An axiomatic approach for result diversification. In WWW, pages 381–390, 2009.
- [18]. Harel D and Tarjan RE. Fast algorithms for finding nearest common ancestors. SIAM Journal on Computing, 13(2):338–355, 1984.
- [19]. Harper FM and Konstan JA. The movielens datasets: History and context. ACM Trans. Interact. Intell. Syst, 5(4):19:1–19:19, Dec. 2015.
- [20]. Hartigan JA. Clustering algorithms. 1975.
- [21]. Huang Z. Extensions to the k-means algorithm for clustering large data sets with categorical values. Data mining and knowledge discovery, 2(3):283–304, 1998.

- [22]. Jagadish H, Madar J, and Ng RT. Semantic compression and pattern extraction with fascicles. In VLDB, volume 99, pages 7–10, 1999.
- [23]. Jagadish H, Ng RT, Ooi BC, and Tung AK. Itcompress: An iterative semantic compression algorithm. In Data Engineering, 2004. Proceedings. 20th International Conference on, pages 646–657. IEEE, 2004.
- [24]. Joglekar M, Garcia-Molina H, and Parameswaran AG. Interactive data exploration with smart drill-down. In ICDE, pages 906–917, 2016. [PubMed: 28210096]
- [25]. Llewellyn DC, Tovey CA, and Trick MA. Erratum: Local optimization on graphs. Discrete Applied Mathematics, 46(1):93–94, 1993.
- [26]. The SM sankey diagram in energy and material flow management. Journal of industrial ecology, 12(1):82–94, 2008.
- [27]. Murphy KP. Naive bayes classifiers. University of British Columbia, 18, 2006.
- [28]. Nambiar RO and Poess M. The making of tpc-ds. In Proceedings of the 32nd international conference on Very large data bases, pages 1049–1058. VLDB Endowment, 2006.
- [29]. Nguyen TT, Nguyen QVH, Weidlich M, and Aberer K. Result selection and summarization for web table search. In ICDE, pages 231–242, 2015.
- [30]. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, et al. Scikit-learn: Machine learning in python. Journal of machine learning research, 12(Oct):2825–2830, 2011.
- [31]. Qin L, Yu JX, and Chang L. Diversifying top-k results. PVLDB, 5(11):1124–1135, 2012.
- [32]. Quinlan JR. Induction of decision trees. Machine learning, 1(1):81–106, 1986.
- [33]. Ravi SS, Rosenkrantz DJ, and Tayi GK. Heuristic and special case algorithms for dispersion problems. Operations Research, 42(2):299–310, 1994.
- [34]. Rubner Y, Tomasi C, and Guibas LJ. The earth mover’s distance as a metric for image retrieval. International journal of computer vision, 40(2):99–121, 2000.
- [35]. Sarawagi S. User-adaptive exploration of multidimensional data. In VLDB, pages 307–316, 2000.
- [36]. Tao Y. Diversity in skylines. IEEE Data Eng. Bull, 32(4):65–72, 2009.
- [37]. <http://grouplens.org/datasets/movielens/>.
- [38]. <http://movielens.org>.
- [39]. Vardi MY. The complexity of relational query languages (extended abstract). In Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing, STOC ‘82, pages 137–146, 1982.
- [40]. Vee E, Srivastava U, Shanmugasundaram J, Bhat P, and Amer-Yahia S. Efficient computation of diverse query results. In ICDE, pages 228–236, 2008.
- [41]. Vieira MR, Razente HL, Barioni MCN, Hadjieleftheriou M, Srivastava D, Traina C, and Tsotras VJ. On query result diversification. In ICDE, pages 1163–1174, 2011.
- [42]. Wagstaff K, Cardie C, Rogers S, Schrödl S, et al. Constrained k-means clustering with background knowledge. In ICML, volume 1, pages 577–584, 2001.
- [43]. Wen Y, Zhu X, Roy S, and Yang J. Qagview: Interactively summarizing high-valued aggregate query answers. In Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10–15, 2018, pages 1709–1712, 2018.
- [44]. Xin D, Cheng H, Yan X, and Han J. Extracting redundancy-aware top-k patterns. In KDD, pages 444–453, 2006.
- [45]. Yu C, Lakshmanan L, and Amer-Yahia S. It takes variety to make a world: Diversification in recommender systems. In EDBT, pages 368–378, 2009.
- [46]. Zaharioudakis M, Cochrane R, Lapis G, Pirahesh H, and Urata M. Answering complex sql queries using automatic summary tables. In SIGMOD, pages 105–116, 2000.
- [47]. Zheng W, Wang X, Fang H, and Cheng H. Coverage-based search result diversification. Information Retrieval, 15(5):433–457, 2012.
- [48]. Zhu X, Goldberg AB, Gael JV, and Andrzejewski D. Improving diversity in ranking using absorbing random walks. In HLT-NAACL, pages 97–104, 2007.

- [49]. Ziegler C-N, McNee SM, Konstan JA, and Lausen G. Improving recommendation lists through topic diversification. In WWW, pages 22–32, 2005.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Rank	hdec	agegrp	gender	occupation	val
1	1975	20s	M	Student	4.24
2	1980	20s	M	Programmer	4.13
3	1980	10s	M	Student	3.96
4	1980	20s	M	Student	3.91
5	1985	20s	M	Programmer	3.86
6	1980	20s	M	Engineer	3.83
7	1985	10s	M	Student	3.77
8	1985	20s	M	Student	3.76
.....					
43	1995	30s	M	Marketing	3.02
44	1995	20s	M	Technician	2.92
45	1995	30s	M	Entertainment	2.91
46	1995	20s	M	Executive	2.91
47	1995	30s	F	Librarian	2.84
48	1995	30s	M	Student	2.81
49	1995	20s	M	Writer	2.51
50	1995	20s	F	Healthcare	1.98

(a) Top-8 and bottom-8 tuples with values as score

hdec	agegrp	gender	occupation	avg_val	
1975	20s	M	Student	4.24	▼
1980	*	M	*	3.96	▼
1985	20s	M	Programmer	3.86	▼
1985	*	M	Student	3.76	▼

(b) Clusters with average score

hdec	agegrp	gender	occupation	avg_val	rank
1975	20s	M	Student	4.24	▼
1975	20s	M	Student	4.24	1
1980	*	M	*	3.96	▼
1980	20s	M	Programmer	4.13	2
1980	10s	M	Student	3.96	3
1980	20s	M	Student	3.91	4
1980	20s	M	Engineer	3.83	6
1985	20s	M	Programmer	3.86	▼
1985	20s	M	Programmer	3.86	5
1985	*	M	Student	3.76	▼
1985	10s	M	Student	3.77	7
1985	20s	M	Student	3.76	8

(c) Original result tuples (with ranks) in the clusters

Figure 1:

Illustrating our framework for $k = 4$, $L = 8$, $D = 2$. In general, original result tuples outside top- L can belong to the output clusters, although here we happen to have a solution that covers just the top- L result tuples.

Figure 2:
Visualization for parameter selection: how results vary for different k and D (some lines overlap).

Figure 3:
(a) Example clusters, and (b) semilattice on clusters.

Figure 4:
Incremental computation and interval structure.

Figure 5:
Comparison with brute-force.

Figure 6:

Experimental results varying parameters. The default values of parameters are $m = 8$, $k = 3$, $L = 40$, $D = 3$.

Figure 7:
Experimental results varying parameters, and with or without precomputation

Figure 8:
Experiment on effects of optimizations

Figure 9:
TPC-DS experimental results varying parameters and with/without precomputations.

Figure 10:
Diagram of System Architecture.

Figure 11:
Range tree example on age

Figure 12:
Range tree example on date

Figure 13:
Visualizing changes between two consecutive clustering in Example 1.2: from $k = 4$ to $k = 3$.

Figure 14:
Solution comparison: clean visualization

Figure 15:
Solution comparison: cluttered visualization

Figure 16:
Experiment on Comparison Visualizations

Table 1:

Summary of results from the user study. Times are in seconds, and accuracies are between 0 and 1; we report average and standard deviation over all subjects. Better performances (shorter times and higher accuracies) and stronger preferences are highlighted with box enclosures, unless the advantage is too small.

Task group		Varying-method		Varying- k		Varying- D	
		Decision tree	Our method	$k = 5$	$k = 10$	$D = 1$	$D = 3$
Patterns-only	Time/question	25.7 ± 6.6	23.5 ± 6.5	19.6 ± 6.6	22.5 ± 5.9	13.5 ± 3.1	9.6 ± 2.9
	T-accuracy	0.792 ± 0.121	0.854 ± 0.132	0.771 ± 0.139	0.833 ± 0.088	0.771 ± 0.087	0.750 ± 0.108
	TH-accuracy	0.646 ± 0.075	0.854 ± 0.075	0.625 ± 0.121	0.708 ± 0.121	0.771 ± 0.087	0.813 ± 0.098
Memory-only	Time/question	9.6 ± 3.8	8.3 ± 2.9	11.1 ± 4.3	9.8 ± 4.2	8.4 ± 2.1	6.6 ± 3.4
	T-accuracy	0.625 ± 0.150	0.792 ± 0.121	0.771 ± 0.139	0.667 ± 0.125	0.646 ± 0.116	0.667 ± 0.108
	TH-accuracy	0.625 ± 0.174	0.792 ± 0.121	0.667 ± 0.108	0.625 ± 0.121	0.771 ± 0.139	0.875 ± 0.083
Patterns+members	Time/question	22.9 ± 4.2	23.7 ± 2.4	21.0 ± 6.2	22.5 ± 4.2	13.5 ± 1.9	14.3 ± 3.3
	T-accuracy	0.922 ± 0.165	0.953 ± 0.061	0.938 ± 0.088	0.953 ± 0.061	0.906 ± 0.104	0.953 ± 0.061
	TH-accuracy	0.750 ± 0.088	0.844 ± 0.104	0.938 ± 0.088	0.97 ± 0.054	0.844 ± 0.054	0.922 ± 0.087
Overall preferred		12.5%	87.5%	43.8%	56.2%	37.5%	62.5%

Table 2:

Summary of results from the user study when varying-method group goes first. Times are in seconds, and accuracies are between 0 and 1; we report average and standard deviation over all subjects. Better performances (shorter times and higher accuracies) and stronger preferences are highlighted with box enclosures, unless the advantage is too small.

Task group		Varying-method		Varying- k		Varying- D	
		Decision tree	Our method	$k = 5$	$k = 10$	$D = 1$	$D = 3$
Patterns-only	Time/question	28.8 ± 7.0	28.2 ± 4.9	19.9 ± 8.1	24.8 ± 7.0	11.8 ± 2.5	9.7 ± 3.4
	T-accuracy	0.833 ± 0.204	0.833 ± 0.118	0.750 ± 0.144	0.833 ± 0.118	0.750 ± 0.083	0.792 ± 0.138
	TH-accuracy	0.708 ± 0.072	0.917 ± 0.083	0.583 ± 0.144	0.750 ± 0.186	0.792 ± 0.072	0.833 ± 0.167
Memory-only	Time/question	9.8 ± 3.0	10.3 ± 2.1	13.8 ± 4.2	11.5 ± 5.3	7.5 ± 0.3	5.5 ± 3.9
	T-accuracy	0.542 ± 0.072	0.708 ± 0.138	0.667 ± 0.204	0.583 ± 0.186	0.667 ± 0.204	0.625 ± 0.138
	TH-accuracy	0.625 ± 0.217	0.917 ± 0.083	0.708 ± 0.138	0.667 ± 0.138	0.792 ± 0.217	0.875 ± 0.138
Patterns+members	Time/question	23.6 ± 5.4	25.3 ± 1.6	17.1 ± 3.1	31.7 ± 4.4	13.5 ± 2.4	13.2 ± 1.9
	T-accuracy	0.875 ± 0.217	0.938 ± 0.063	0.969 ± 0.054	0.969 ± 0.054	0.906 ± 0.104	0.938 ± 0.063
	TH-accuracy	0.750 ± 0.088	0.844 ± 0.136	0.938 ± 0.063	0.97 ± 0.054	0.844 ± 0.054	0.969 ± 0.054
Overall preferred		12.5%	87.5%	43.8%	56.2%	37.5%	62.5%