

AlphaSum: Size-Constrained Table Summarization using Value Lattices *

K. Selçuk Candan, Huiping Cao, Yan Qi
 Arizona State Univ.
 Tempe, AZ 85283, USA
 {candan,hcao11,yan.qi}@asu.edu

Maria Luisa Sapino
 Univ. di Torino
 Torino, Italy
 mlsapino@di.unito.it

ABSTRACT

Consider a scientist who wants to explore multiple data sets to select the relevant ones for further analysis. Since the visualization real estate may put a stringent constraint on how much detail can be presented to this user in a single page, effective table summarization techniques are needed to create summaries that are both sufficiently small and effective in communicating the available content. In this paper, we first argue that table summarization can benefit from knowledge about acceptable value clustering alternatives for clustering the values in the database. We formulate the problem of table summarization with the help of value lattices. We then provide a framework to express alternative clustering strategies and to account for various utility measures (such as information loss) in assessing different summarization alternatives. Based on this interpretation, we introduce three preference criteria, max-min-util (cautious), max-sum-util (cumulative), and pareto-util, for the problem of table summarization. To tackle with the inherent complexity, we rely on the properties of the fuzzy interpretation to further develop a novel ranked set cover based evaluation mechanism (RSC). These are brought together in an *AlphaSum*, table summarization system. Experimental evaluations showed that RSC improves both execution times and the summary qualities in *AlphaSum*, by pruning the search space more effectively than the existing solutions.

1. INTRODUCTION

Summarization of large data tables is required in many scenarios where it is hard to display complete data sets. [17, 32], for example, leverages table summarization for mobile commerce applications over PDAs, which cannot effectively present a large table of results with their small screens. Table summarization is also useful in various other scenarios, where it is hard or highly-impractical to visualize large data

*Authors are listed in the alphabetical order; Supported by NSF Grant “Archaeological Data Integration for the Study of Long-Term Human and Social Dynamics” (0624341)

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the ACM. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permissions from the publisher, ACM. EDBT 2009, March 24–26, 2009, Saint Petersburg, Russia. Copyright 2009 ACM 978-1-60558-422-5/09/0003 ...\$5.00


Name	Age	Location
John	12	Phoenix
Sharon	19	Los Angeles
Mary	19	San Diego
Peter	22	Baltimore
James	22	Frederick
Alice	27	Baltimore

(a) Data table

Name	Age	Location
-	1*	Southwest
-	2*	Maryland

(b) Data table after 3-summarization on $\langle \text{Age}, \text{Location} \rangle$ using hierarchies in Figure 1

Table 1: (a) A database and (b) A 3-summarized version on the $\langle \text{Age}, \text{Location} \rangle$ attribute pair



(a) Hierarchy for Age

(b) Hierarchy for Location

Figure 1: Hierarchy-based value clustering for attribute Age (a) and Location (b). The directed edges denote the clustering/summarization direction.

sets. Consider, for example, a scientist exploring the *Digital Archaeological Record (tDAR/FICSR)* [26, 27], a digital library which archives and provides access to a large number of (and diverse) data sets, collected by different researchers within the context of different projects and deposited to *tDAR* for sharing. When this scientist poses a search query through *FICSR*, her query might match many potentially relevant databases and data tables. For this scientist to be able to explore the multitude of candidate data resources identified by *tDAR* as quickly and effectively as possible, novel data reduction techniques (such as table summarization) are needed.

In this paper, we first note that visualization process can leverage knowledge about available screen real estate. For example, Table 1 (a) shows a data table consisting of 6 rows. If the user is interested in summarizing the given table based on the attribute pair, $\langle \text{Age}, \text{Location} \rangle$ in such a way that, the summarized table can be visualized in a space that can hold at most 2 tuples, this can be achieved by clustering rows such that each row in the summary corresponds to at least 3 ($= \lceil 6/2 \rceil$) rows in the original table – we refer to this as

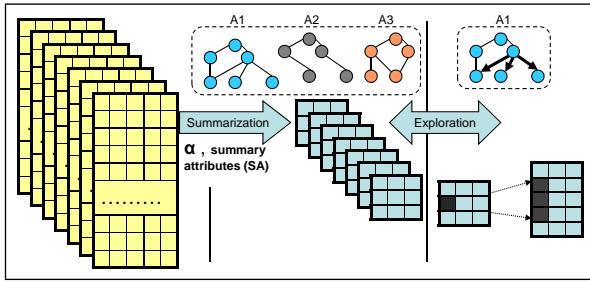
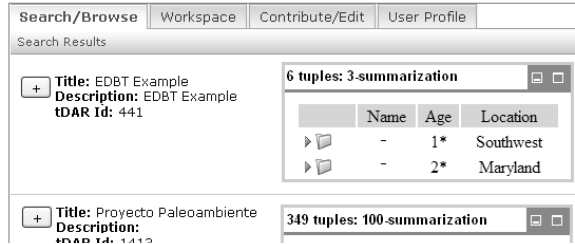



Figure 2: α -summarization helps reduce the table size for quick exploration and enable the user explore the summary as in OLAP



(a) A 3-summarization of Table 1(a) on attribute pair $\langle \text{Age}, \text{Location} \rangle$ using the hierarchies in Figure 1



(b) First level expansion: from (a), expand *Southwest* of the record $(-, 1^*, \text{Southwest})$ for attribute *Location*

(c) Second level expansion: from (b), expand *California* of the record $(-, 1^*, \text{California})$ for attribute *Location*

Figure 3: (a) a 3-summarized version on the $\langle \text{Age}, \text{Location} \rangle$ attribute pair, and (b)-(c) expand different clustering values for the attribute *Location*.

3-summarization. Furthermore, this clustering can be best achieved if additional domain knowledge, which can help choose how to cluster the values, is available. For example, given the value hierarchies in Figure 1(a) for attribute *Age* and Figure 1(b) for attribute *Location*, one can 3-summarize this table as shown in Table 1(b).

1.1 Motivating Example: Exploration of Data Sets

Value hierarchies (such as the one in Figure 1(a)) are commonly used for user-driven data exploration *within* large data sets. For example, OLAP operators (such as *drill-down*, navigating from less specific to more detailed data by stepping down on a given hierarchy, and *roll-up*, which performs aggregation on by climbing up a hierarchy of concepts) commonly assume value hierarchies are given [6].

OLAP-based navigation (using drill-down and roll-up operations), on the other hand, does not take into account the visualization real estate (e.g., the number of tuples to be displayed) available for exploration. For instance, when a user poses a *source location* query to tDAR to identify potential




Figure 4: Lattice-based value clustering graphs for the attribute *Age*.

data sets for further study, the system would return multiple candidate tables to be further explored for relevance. Displaying multiple tables (or even a single large table) on a single result page imposes limits on how many tuples can be presented to the user for each table. As in web search interfaces which provide snippets for each candidate result to help users choose which link to follow, summarization of the candidate tables can help them fit onto the user screen. Thus, if we can create summaries leveraging an underlying value clustering hierarchy, these summaries can be used for exploring data sets by refining the relevant portions of the summary results by moving up-down on the hierarchy as in OLAP’s *drill-down* and *roll-up* (Figure 2).

Consider, for example, the data table in Table 1(a) and the value clustering hierarchies for attributes *Age* and *Location*, shown in Figure 1. Figure 3(a) shows a 3-summarization of this table (on attribute pair $\langle \text{Age}, \text{Location} \rangle$) visualized through our *AlphaSumExplorer* interface. *AlphaSumExplorer* lets users expand on specific attributes or summarized values to explore the table to further examine selected details. Figure 3(b) and 3(c) show a sample sequence of expansions (similar to *drill-down* in OLAP), first on the clustered value “*Southwest*” and then on the clustered value “*California*.” In tDAR, we are using an efficient α -summarization system (*AlphaSum* described in this paper) and the *AlphaSumExplorer* interface for navigating through large data sets obtained through queries over heterogeneous archaeological data sets.

1.2 AlphaSum: Size-Constrained Summarization of Tables

As described above, the main goal of *AlphaSum* is to obtain OLAP-like navigable summaries from large tables. Its goals, however, differ from OLAP in two significant ways:

- First of all, unlike traditional OLAP, *AlphaSum* takes the visualization real estate into account. Given a table with n tuples, the maximum number, $m (= \lceil n/\alpha \rceil)$ of tuples to be visualized to the user as a highest level summary is an input parameters to *AlphaSum*.
- Secondly, *AlphaSum* recognizes that in many applications (even in OLAP [28]), systems may need to take into account the existence of multiple acceptable value clustering *alternatives*. For example, the value lattice shown in Figure 4 provides more clustering alternatives than the simpler hierarchy in Figure 1(a), potentially enabling the applications to be more precise and informative.

When there are alternative clustering strategies for the values in a database, however, not all alternatives might be equally desirable. Let us reconsider Figure 4:

- Clustering the value “12” under “*preteen*” (which corresponds to values $0 \leq \text{age} < 13$) versus “1*” (which




Figure 5: A partial taxonomy of animals

corresponds to values $10 \leq age \leq 19$) may be less desirable because the range of possible values (i.e., the precision lost by the replacement) in the former case is larger than in the latter.

- If the nodes of the lattice do not represent concrete sets, but concepts in an ontology or a *concept* taxonomy, then the semantic difference between terms can be estimated using *structural* and/or *information-based* approaches.

In structure-based methods [23, 25], the distance between two nodes is measured as the sum of the *distance* weights of edges between them: the weight of an edge between a parent concept and a child is measured based on the structural clues available in the lattice, such as depth (the deeper the edge in the lattice, the less information loss associated to the edge; e.g., in Figure 5, *dogs* are more related to *wolves* than *mammals* are to *fishes*) and local density of the concepts (the denser the lattice, the smaller the semantic distance between concepts in that neighborhood of the lattice).

Information-based methods leverage available corpora to measure potential information loss. For example, knowledge about term-occurrence frequencies may help measure information-loss in an information theoretic manner [24] (information content is measured as the negative logarithm of the probability of encountering an instance of the concept in the corpus).

Therefore, when creating an α -summary (where each row corresponds to at least α rows in the original table, as in Table 1(b)) by clustering data values according to the given lattice, we need to account for the loss of precision.

1.3 Contributions

In this paper, we first formulate the problem of α -summarization under weighted value clusterings. Instead of solving a special case of the problem (e.g., assuming a hierarchy instead of a lattice or assuming a specific information loss measure), we introduce a general *utility-weighted value clustering lattice* representation to capture alternatives and choices over these alternatives, we also introduce the concept of utility of a clustering strategy, and formalize the problem of α -summarization, which is the core of the *AlphaSum* system (Section 3). Relying on this, we also introduce three solution preference criteria, *max-min-util*, *max-sum-util*, and *pareto-util* for the problem of α -summarization. We then rephrase the summarization problem into a set cover problem and present a novel *ranked set cover* algorithm that is able to prune the search space significantly to reduce the execution time for α -summarization task (Section 5). In Section 6, we evaluate the proposed approach and show that the

proposed ranked set cover algorithm improves both execution time and result quality.

2. RELATED WORK

In [2], Alfred *et al.* present an approach to data summarization by aggregating data over a relational dataset. A good survey of the database summarization can be found in [29], where an online linguistic table summarization system, *SaintEtiQ*, is also presented. *SaintEtiQ* [29, 30] computes and incrementally maintains a hierarchically arranged set of summaries of the input table. In such a hierarchical structure, any non-leaf summary generalizes the content of its children nodes. TabSum [17] summarizes tables through row and column reduction. To reduce the number of rows, it first partitions the original table into groups based on one or more attribute values of the table, and then collapses each group of rows into a single row relying on the available metadata, such as the concept hierarchy. For column reduction, it simplifies the value representation and/or merges multiple columns into one column. [32] discusses a related approach for refinement of table summaries. Neither of these approaches, however, considers the impact of the imprecision of metadata and the information loss during summarization.

Data compression techniques, like Huffman or Lempel-Ziv, can also be used to reduce the size of the table. A database compression technique is designed in [20] based on vector quantization. Buchsbaum *et al.* [4, 5] address the problem of compressing massive tables through a partition-training paradigm. These methods are not directly applicable, since the compressed tables are not human readable. Histograms can be also exploited to summarize information into a compressed structure. Following this idea, Buccafurri *et al.* [3] introduce a quad-tree based partition schema for summarizing two-dimensional data. Range queries can be *estimated* over the quad-tree since the summarization is a lossy compression. Leveraging the quad-tree structure, [8, 9] proposes approaches to processing OLAP operations over the summary. [9] first generates a 2-dimensional OLAP view from the input multidimensional data and then compresses the 2-dimensional OLAP view by means of an extended quad-tree structure. In fact, any multidimensional clustering algorithm can be used to summarize a table. Such methods, however, cannot take into account specific domain knowledge (e.g. “what are acceptable summarizations, how do they rank?”) that hierarchies and lattices would provide.

OLAP operations, *drill-down* and *roll-up*, which help users navigate between more general and more specific views of the data with the help of given value hierarchies, are also related to table summarization. The concept of imprecision in OLAP dimensions is discussed in [22]. In that framework, a fact (e.g., a tuple in the table) with imprecise data is associated with dimension values of coarser granularities, resulting in the dimensional imprecision. In schema design for traditional relational-OLAP systems, issues around heterogeneous dimensions have been discussed in [14, 13]. A dimension is called heterogeneous if two members in a given category are allowed to have ancestors in different categories. Given a heterogeneous dimension, an aggregate view for a category may not be correctly derived from views for its sub-categories (a summarizability¹ problem). In

¹Summarizability is a notation to denote the conditions under which a value or object can be summarized correctly from a more detailed value or object, based on the given

[28], we supported OLAP operations over imperfectly integrated taxonomies. We proposed a re-classification strategy which eliminates conflicts by introducing minimal imprecision. This approach is complementary to the work presented here: the obtained navigable taxonomy can be taken as the input for table summarization.

The table summarization task is also related to k -anonymization problem, introduced as a technique against linkage attacks on private data [31, 16]. The k -anonymization approach eliminates possibility of such attacks by ensuring that, in the disseminated table, each value combination of attributes is matched to k others. To achieve this, k -anonymization techniques rely on a-priori knowledge about acceptable value generalizations. Cell generalization schemes [1] treat each cell in the data table independently. Thus, different cells for the same attribute (even if they have the same values) may be generalized in a different way. This provides significant flexibility in anonymization, while the problem remains extremely hard (NP-hard [18]) and only approximation algorithms are applicable under realistic usage scenarios [1]. Attribute generalization schemes [31, 16] treat all values for a given attribute collectively; i.e., all values are generalized using the same unique domain generalization strategy. While the problem remains NP-hard [18] (in the number of attributes), this approach saves significant amount of time in processing and may eliminate the need for using approximation solutions, since it does not need to consider the individual values. Most of these schemes, such as Samarati’s original algorithm [31], however, rely on the fact that, for a given attribute, applicable generalizations are in total order and that each generalization step in this total order have the same cost. [31] leverages this to develop a binary search scheme to achieve savings in time. [16] relies on the same observation to develop an algorithm which achieves attribute-based k -anonymization one attribute at a time, while pruning unproductive attribute generalization strategies.

Unlike these solutions to k -anonymization, the summarization algorithm presented in the paper does not rely on any constraints imposed on the shape and weights of the value clustering lattice and, instead, owes its efficiency to a novel ranked set cover based evaluation mechanism (RSC) that is able to prune the search space effectively.

3. PROBLEM FORMULATION

Let us consider a data table, T , and a set, SA , of summarization-attributes. Roughly speaking, our purpose is to find another relation T' which clusters the values in T such that T' satisfies the α -summarization property (where each row in T' clusters at least α rows in T) with respect to the summarization-attributes. In what follows, we first formalize the concept of *weighted clustering lattices*. Based on this, we then specify the problem of α -summarization over weighted clustering lattices.

3.1 Value Clustering Lattices

We model a value clustering lattice as an acyclic graph $M(V, E)$ where V encodes values and clustering-identifiers (e.g., high-level concepts or cluster labels, such as “1*” in Figure 1(a) and Figure 4 and E contains acceptable value clustering relationships.

summarization rules (e.g., value mappings, value lattices, etc.). It has been defined in [33] and used in the context of OLAP.




Figure 6: A weighted lattice for the attribute Age

DEFINITION 3.1 (VALUE CLUSTERING LATTICE).

A value clustering lattice L is a directed acyclic graph $M(V, E)$:

- $v = (id : value) \in V$ where $v.id$ is the node id in the lattice and $v.value$ is either a value in the database or a value clustering encoded by the node.
- $e = v_i \rightarrow v_j \in E$ is a directed edge denoting that the value encoded by the node v_i can be clustered under the value encoded by the node v_j .

Those nodes in V which correspond to the attribute values that are originally in the database do not have any incoming edges in E . There is only one root (i.e., node without any outgoing edge) in the lattice.

The value clustering lattice describes a partial order that reflects acceptable clustering strategies for each value (e.g., Figure 4 and Figure 5).

3.2 Value Clustering

Given an attribute value in the data table T and a weighted value clustering lattice corresponding to that attribute, we can define alternative clusterings as paths on the corresponding lattice:

DEFINITION 3.2 (VALUE CLUSTERING). Given a weighted value clustering lattice L , a lattice node v_j is a clustering of a lattice node v_i , denoted by $v_i \preceq v_j$, if $\exists path p = v_i \rightsquigarrow v_j$ in L . We also say that v_j covers v_i . Note that p may also be empty.

Based on this, in the following, we formalize the concept of *utility* of value clusterings.

3.2.1 Utility of a Value Clustering Strategy

As described earlier, the utility value of a clustering step is domain specific and may represent different reasons one may prefer one clustering alternative over the other. A high utility value indicates a clustering step with low information loss and high precision (utility value 1.0 indicates zero information loss and uncertainty due to clustering; thus is excluded from the model). Any utility below 1.0 denotes a certain degree of knowledge relaxation. For example, under one interpretation, the edge in Figure 4 from “12” to “preteen” would have utility 1/13 (“12” is one out of 13 values the label “preteen” captures). Similarly, the edge from “12” to “1*” has utility 1/10 (Figure 6).

Given the utilities of clustering edges in the lattice, we can also compute the utility of a value clustering strategy, $util(v_i \preceq v_j)$, between v_i and any of its ancestors v_j , by considering all paths between v_i and v_j .

The value clustering utility function used for combining utilities of alternative ways v_i can be re-written into v_j may differ depending on the semantics of the lattice. In the case when precision is defined in terms of set membership, all




Figure 7: An imprecise taxonomy of animals: there is a disagreement among users of the taxonomy as to whether *dogs* should be classified under *small* or *medium mammals*

path utilities are identical, thus the combination function would simply return the utility of any path. For example, in Figure 6, all paths from value, 19, to the clustering value, *, would return the same path utility, 1/121, since 19 is one out of 121 values covered by *.

If different paths between v_i and v_j have different utilities, on the other hand, these may need to be combined using other combination functions. For example, in Figure 7, there are two different paths between “*dogs*” and “*mammals*” representing different views of the scientists using the taxonomy. In this case, replacement of *dogs* with *mammals* needs to consider all interpretations and combine the corresponding utilities. Candidate combination functions include, *minimum*, *maximum*, and *sum*. For example, in Figure 7, *sum* may be suitable: independently of whether “*dogs*” are *small* or *medium mammals*, they are *mammals*. Thus, the imprecision at the lower levels of the lattice should not matter when considering the utility of replacing “*dogs*” with “*mammals*”.

3.2.2 Computing the Value Clustering Utilities

Given a value clustering lattice, \mathcal{L} , which provides a partial order of value clusterings, there may be exponentially many paths between a given pair of values, v_i and v_j . Thus, computing the value clustering utilities on the given lattice may be costly, if performed naively. We note, however, that in many practical cases, value clustering utilities between all pairs of values in the value hierarchy can be computed in polynomial time. Especially in the special case where any replacement path between a given pair of values return the same precision loss, the cost of the computation is simply quadratic in the number of nodes.

The challenge is dealing with the exponential growth in the number of summarization-tuple clustering strategies when the number of summarization-attributes increases.

3.3 Summarization-Tuple Clustering

So far we have focused on the clustering of the value of single summarization-attribute; i.e., the case where the size of the set, SA , of summarization-attributes is one and there is only one clustering domain. However, in general, the size of the summarization-attribute set of a data table can be greater than one. Thus, we first need to define clusterings involving non-singleton summarization-attributes:

DEFINITION 3.3 (SUMMARIZATION-TUPLE CLUSTERING). *Let t be a tuple on attributes $SA = \{Q_1, \dots, Q_q\}$. t' is said to be a clustering of the summarization-tuple (SA-tuple), t , (on attributes SA) iff $\forall i \in [1, q]$*

- $t'[Q_i] = t[Q_i]$, or
- $\exists \text{path } p_i = t[Q_i] \rightsquigarrow t'[Q_i]$ in the corresponding clustering lattice L_i .

We use $t \preceq t'$ as shorthand.

3.3.1 Utility of a SA-Tuple Clustering Strategy

Given the above definition, we can define the utility of a summarization-tuple clustering strategy:

DEFINITION 3.4 (UTILITY OF A SA-TUPLE CLUSTERING.). *Let t and t' be two tuples on attributes $SA = \{Q_1, \dots, Q_q\}$, such that $t \preceq t'$. Then the utility of the corresponding clustering strategy is defined through a monotonic combination function, \otimes , of the utilities of the clustering along each individual summarization-attribute:*

$$\text{util}(t \preceq t') = \bigotimes_{1 \leq i \leq q} \text{util}_i,$$

where

- $\text{util}_i = 1$, if $t'[Q_i] = t[Q_i]$
- $\text{util}_i = \text{util}(t[Q_i] \preceq t'[Q_i])$, otherwise.

When $q=1$ (i.e., there is only one summarization-attribute), $\text{util}(t \preceq t')$ is simply the utility of the corresponding single attribute value clustering strategy. In general, \otimes can be any monotonically increasing function; for example, minimum, multiplication, or summation.

3.4 α -Summarization Problem

Having defined *utilities* of tuple clusterings, we can now define the problem of α -summarization with lattices:

DEFINITION 3.5 (α -SUMMARIZED CLUSTERING). *Given two data tables T and T' (with the same schema), and the summarization-attribute set SA , T' is said to be a α -summarized clustering of T ($T \preceq_\alpha T'$ for short) iff*

- $\forall t' \in T'[SA]$, $|\{t'' | t'' \in T'[SA] \wedge t''[SA] = t'\}| \geq \alpha$ (i.e., t' appears at least α times in T'),
- there is a one-to-one and onto² mapping, μ , from $T'[SA]$ to $T[SA]$, such that
 - $\forall t' \in T'[SA]$, $\mu(t') \preceq t'$.

Here, $T[SA]$ and $T'[SA]$ are projections of the data tables T and T' on summarization-attributes. $T[SA]$ and $T'[SA]$ are multi-sets; i.e., they may contain multiple instances of the same summarization-tuple.

Given a data table T , there can be multiple such α -summarized clustering alternatives of T . We thus define the utility of a α -summarized clustering as follows:

DEFINITION 3.6 (TABLE CLUSTERING UTILITY). *Let T be a table and T' be its α -summarized clustering (i.e., $T \preceq_\alpha T'$). The utility of T' is defined using a monotonic merge function, \uplus , as follows:*

$$\text{util}(T') = \biguplus_{t' \in T'[SA]} \text{util}(\mu(t') \preceq t').$$

Since there can be more than one alternative clustering, our goal is to find one that maximizes the utilities of the underlying clustering.

PROBLEM 1 (UTILITY MAXIMIZING α -SUMMARIZATION). *Given T , SA , \mathcal{L} , α , \otimes , and \uplus , where*

²It is possible to relax the *onto* requirement and allow for some limited number of tuples which are hard to cluster with others, due to outlier values, to remain unsimplified. In this paper, we do not explicitly consider this situation, though it is possible to account for these tuples by representing them through a special form of clustering alternative, with 0 utility value.

- T is a data table with n tuples to be published,
- $SA = \{Q_1, \dots, Q_q\}$ is the set of summarization-attributes,
- $\mathcal{L} = \{L_1, \dots, L_q\}$ is the set of weighted clustering lattices for the attributes in SA ,
- α is the summarization parameter, and
- \otimes and \uplus are the appropriate utility merge functions,

find another data table T' (and the corresponding mapping μ from $T'[SA]$ to $T[SA]$) which is a α -summarized clustering of T , such that there exists no other α -summarized data table T'' (and the corresponding mapping μ_2), such that T'' is preferable over T' .

3.5 α -Summarization Preference Criteria

In this paper, we consider the following preference criteria.

DEFINITION 3.7 (MAX-MIN-UTIL, CAUTIOUS, CRITERION).

Let T be a data table and let T' and T'' be two α -summarized clusterings. T'' is **max-min-util** preferable (cautiously preferable) over T' iff

$$\left(\min_{t' \in T'} \text{util}(\mu_1(t') \preceq t') \right) < \left(\min_{t'' \in T''} \text{util}(\mu_2(t'') \preceq t'') \right).$$

DEFINITION 3.8 (MAX-SUM-UTIL, CUMULATIVE, CRITERION).

Let T be a data table and let T' and T'' be two α -summarized clusterings. T'' is **max-sum-util** preferable (cumulatively preferable) over T' iff

$$\left(\sum_{t' \in T'} \text{util}(\mu(t') \preceq t') \right) < \left(\sum_{t'' \in T''} \text{util}(\mu_2(t'') \preceq t'') \right).$$

EXAMPLE 3.1. The following example illustrates how the Max-sum-util and Max-min-util preference criteria work for the same summarization requirement. Suppose we have a table with six tuples (t_1 to t_6) shown in Figure 8 and we want to obtain a 2-summarization for it.

Let Figure 8(a) show the possible clustering options for these tuples. Figure 8(b)-(d) show three different summarizations where the solid directed line from a tuple to a tuple clustering denotes the summarization selection, whereas the dashed lines represent the discarded clustering options. For example, in Figure 8(b), tc_7 is selected to cluster t_5 (with utility 0.7) and t_6 (with utility 0.1). Although t_4 can be clustered to tc_7 , this summarization strategy clusters it to tc_4 with utility 0.4. So, the lines from t_5, t_6 to tc_7 are solid, but the line from t_4 to tc_7 is dashed.

Figure 8(b) and 8(c) show the Max-sum-util summarization and Max-min-util summarization, respectively. The summarization in Figure 8(d), on the other hand, satisfies neither the Max-sum-util criterion nor the Max-min-util criterion.

Max-min-util criterion is a *cautious* standard which guarantees that the worst of the tuple summarizations has high value by maximizing the lowest of the tuple utilities. In other words, \uplus is the *minimum* function.

Max-sum-util criterion is a *cumulative* standard which considers the table summarization as a whole and requires that the total tuple utility of the resulting table is maximized. In other words, \otimes is *summation*.

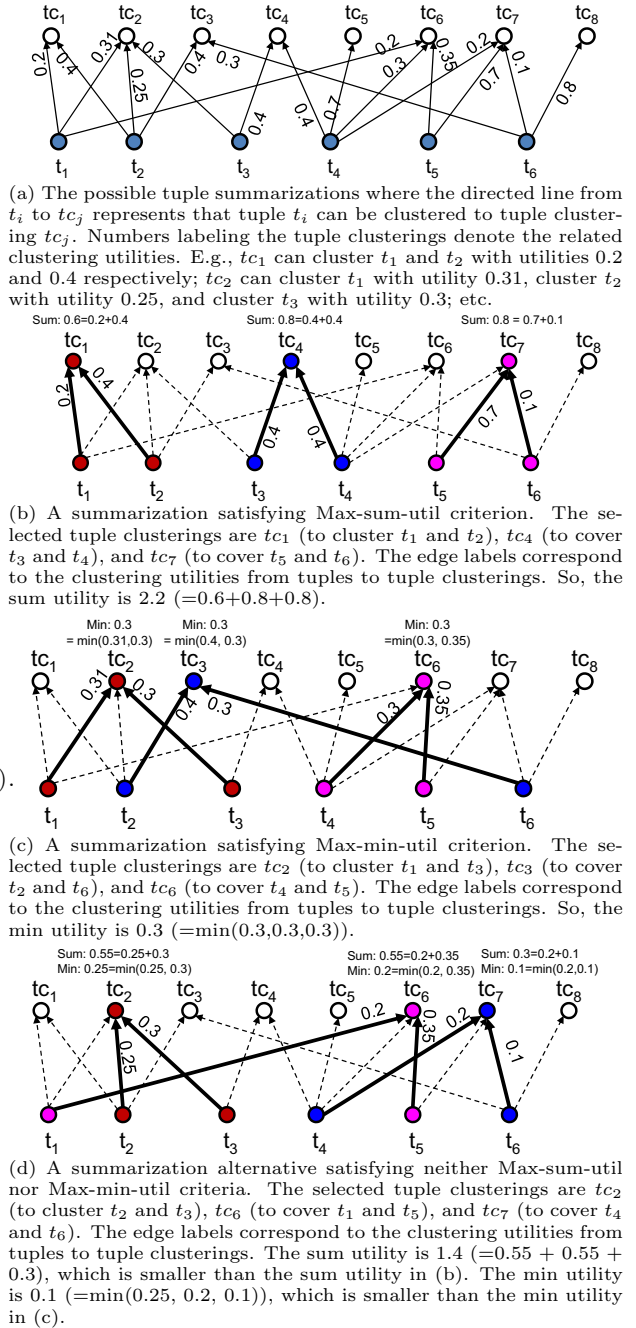


Figure 8: Summarization example with different summarization criteria

4. EXHAUSTIVE FORMULATIONS

We first present exhaustive formulations to the α -summarization problem with weighted lattices. Despite being inefficient, this outlines the overall framework, enables us to study sources for the complexity of the problem, and highlights opportunities for ranked executions. In Section 5, then, we present a novel ranked algorithm that is able to prune the search space significantly to reduce the time for identifying solutions.

4.1 Outline of the Exhaustive Approach

In this subsection, we outline the four steps of an exhaustive approach to the problem. Let us be given a data

table, T , with n tuples, summarization-attributes, $SA = \{Q_1, \dots, Q_q\}$, the set, $\mathcal{L} = \{L_1, \dots, L_q\}$, of clustering lattices for the attributes in SA , and the summarization parameter, α .

Step 1. Computation of the value clusterings: The exhaustive algorithm first considers each value val_i for every attribute $Q_j \in SA$. Let v_i be the lattice node for val_i in the clustering lattice L_j . The clusterings are sorted according to their utility values. Since this step requires at most $|SA|$ all pairs computations, for most practical situations, it is polynomial in time complexity.

Step 2. Computation of the tuple clusterings: For each tuple t in $T[SA]$, the basic algorithm computes the corresponding tuple clusterings and clustering utilities by combining all value clusterings for values $t[Q_1], \dots, t[Q_q]$. Since this step computes the tuple clusterings by combining all possible value clusterings, its worst case space and time complexity is $O(\sum_{i=1}^n \prod_{j=1}^{|SA|} gen_num(t_i, Q_j))$, where $gen_num(t_i, Q_j)$ represents the number of clusterings of value t_i for summarization-attribute Q_j . The complexity term can be approximated as $O(n \cdot vg^{|SA|})$, where vg is the average number of clusterings of the attribute values: the complexity is linear in the number of tuples n , but is exponential in $|SA|$.

Step 3. Preliminary Filtering: In its third step, for each tuple clustering obtained in Step 2, the exhaustive algorithm computes the number of tuples that can be generalized to it. Those tuple clusterings covering less than α tuples are simply removed, as they cannot be involved in any α -summarized result. To identify tuple clusterings covering at least α tuples, the basic algorithm maintains a hash table to keep track of the distinct tuple clusterings. Each item in the hash table is a $\langle tc, cnt \rangle$ pair, where cnt is the number of tuples that tc can cover. While scanning the list of possible clusterings for each tuple, if the clustering has already been recorded (for another tuple), the corresponding count in the hash table is incremented. The space complexity of this step is $O(dist_gen_count)$ where $dist_gen_count$ is the total number of distinct tuple clusterings (which is in the worst case exponential in $|SA|$). The time complexity is $O(n \cdot vg^{|SA|})$; i.e., the time used to scan the tuple clusterings.

Step 4. Once the above three preparatory steps are completed, in its final step the algorithm computes the best α -summarized clustering.

4.2 Set-Covering Formulation

The set-covering problem can be stated as follows: given a finite set S , and a collection, C , of subsets of S , one aims to find a subset, C' , of C , where every element of S belongs to at least one subset in C' and where some objective criterion (such as the cardinality of the set cover or the sum of cardinalities of the subsets in the set cover) is optimized.

We note that, the requirements of the Step 4 of the exhaustive algorithm can be represented as a set covering problem: by treating (a) each tuple clustering tc_i as a set, (b) each tuple t_i as an element, and (c) every clustering relation $t_i \preceq tc_j$ as the set membership relationship $t_i \in tc_j$, we can formulate the α -summarization problem with weighted clustering hierarchies as a k -set-cover problem (for $k = \alpha$).

While the k -set-cover problem is known to be NP-complete, there are various approximation algorithms [7, 21]. Thus the problem under **max-sum-util** preference criterion can be solved through a k -set-cover approximation

Input : Sorted tuple clustering streams ($TCStr_i$ s) for each tuple in T , and summarization parameter α .

Any tuple clustering not covering at least α tuples are dropped from the stream.

Output : TC' such that clustering rules are satisfied

1. $TC_{combined} = \emptyset$;
 2. $PQ = \emptyset$; /*priority queue of tuple clusterings at the boundary*/
 3. Get the first tuple clusterings from each $TCStr_i$;
 4. Insert into $TC_{combined}$;
 5. $initPQ(PQ)$;
 6. repeat
 - (a) if $TC_{combined}$ covers all tuples in T
 - i. $TC' = construct_and_solve(TC_{combined}, \alpha)$;
 - ii. if $TC' \neq \emptyset$,
 - return** (TC' , $extract_mapping(TC')$) and **stop**;
 - (b) $tc_{nextbest} = getFromPQ()$ /*get from the priority queue the tuple clustering with the next best utility value*/
 - (c) $TC_{combined} = TC_{combined} \cup \{tc_{nextbest}\}$
- until $tc_{nextbest} = null$

initPQ(PQ)

1. Get the next tuple clusterings from each $TCStr_i$;
2. Insert them into PQ ;

getFromPQ(PQ)

1. $tc_h = dequeue(PQ)$;
 2. Let $TCStr_i$ be the stream from which tc_h was originally pulled from
 3. Get the next tuple clustering, tc_j from $TCStr_i$
 4. $enqueue(PQ, tc_j)$
 5. **return** (tc_h)
-

Figure 9: Pseudo-code of the ranked algorithm

algorithm. In fact, in the next section, we will show that it is possible to leverage set-covering based formulation of the problem within an incremental framework which produces results more efficiently *and* effectively than a *naive* set-covering formulation.

5. RANKED PROCESSING

In the previous section, we have seen that exhaustive solutions to the proposed problem has three sources of complexity: (a) the number of summarization-attributes increases the number, m , of tuple clusterings exponentially, (b) the set-covering formulation of the problem has an exponentially increasing cost as a function of m . Also, (c) the **max-min-util** preference criterion leads to a non-linear (i.e., more expensive) problem formulation. In this section, we first present a generic ranked algorithm which significantly reduces the number, m , of tuple clusterings which need to be considered. We then develop a novel ranked set-covering algorithm that performs efficiently and more effectively than non-ranked set-covering solutions.

5.1 Ranked Algorithm

The algorithm, outlined in Figures 9 and 10, takes as input (a) a stream of clusterings for each tuple in $T[SA]$ and (b) the summarization parameter α . Tuple clustering streams are taken as input streams in decreasing order according to their *util* values. The algorithm merges the n sorted tuple clustering streams by considering first those tuple clusterings with the highest *util* values (Step 3). Any tuple clustering that is not covering at least α input tuples are eliminated from consideration. To identify tuple clusterings covering at least α tuples, we maintain a hash table of clusterings, which keeps the count of the input streams in which a given clustering is seen. Once we obtain a proper set of clusterings,




Figure 10: Outline of the ranked set cover algorithm

$TC_{combined}$, which covers all the tuples in the data table, we convert this set into set-covering problem (Step 6(a)i). If this step returns a solution, then the ranked algorithm stops (Step 6(a)ii). Otherwise, it continues incrementally considering other tuple clusterings in their utility order. The algorithm uses a priority queue to keep track of the next best clusterings of the tuples in the database (Steps 5 and 6b). This process continues until either all tuple clusterings have been checked or a solution is found.

The correctness of the above ranked algorithm depends on the objective function used. Thus, we need to consider each preference criterion independently. Before we do so, however, we state a fundamental lemma:

LEMMA 5.1 (MONOTONICITY). *If the set of tuple clusterings, $TC_{combined}$, contains a (not necessarily optimal) α -summarized table clustering (i.e., satisfying the α -constraint), then any superset of $TC_{combined}$ also contains the same α -summarized table clustering.*

Max-min-util Criterion.

THEOREM 5.1. *The ranked algorithm is correct w.r.t. the preference criterion, max-min-util.*

Proof sketch A direct corollary of Lemma 5.1 and of the fact that tuple clusterings are considered in decreasing order of utility is that the first solution satisfying the α constraint is also the one which has the largest minimum utility. Any further solution will need to include at least one tuple clustering with a lower utility³, thus cannot improve on this objective function. \square

In fact, since the input streams are sorted in decreasing order of utilities of tuple clusterings, the first α -summarized solution enumerated using ranked processing will also be the one with the maximum minimum utility; i.e., it will implicitly satisfy the max-min-util criterion.

Max-sum-util Criterion. Secondly, let us consider the max-sum-util preference criterion. Figure 11 provides a counter-example illustrating that the above ranked algorithm **does not** provide an optimal solution for this objective function. In this example, the weights on the edges are the corresponding tuple clustering utilities, thus the order in which clustering will be considered is tc_1 , tc_2 , tc_3 and tc_4 . The ranked algorithm will provide a 2-summarized solution (using tc_2 and tc_3) with total utility 1.6, after tc_1 , tc_2 , and tc_3 are seen. However, there is a better solution (using tc_1 and tc_4), with overall utility 1.8.

So, we define another preference criterion.

³For simplicity, we assume utilities are distinct. It is trivial to extend the argument to the cases where there can be multiple clusterings with identical utilities.




Figure 11: An example tuple clustering lattice for which the ranked approach is sub-optimal under max-sum-util criterion

DEFINITION 5.1 (PARETO-UTIL). *Let T be a data table and let T' and T'' be two α -summarized clusterings. T'' is pareto-util preferable over T' iff*

$$\forall_{\mu(t') \equiv \mu_2(t'')} (\text{util}(\mu(t') \preceq t') \leq \text{util}(\mu_2(t'') \preceq t'')) \wedge$$

$$\exists_{\mu(t') \equiv \mu_2(t'')} (\text{util}(\mu(t') \preceq t') < \text{util}(\mu_2(t'') \preceq t'')).$$

Pareto-util criterion prefers solutions where (for any monotonic merge function, \uplus , e.g., minimum or summation) it is not possible to obtain any further improvements without having to make any single tuple worse in utility. The ranked set cover algorithm finds summarizations with Pareto-util.

Pareto-optimal Criterion. The ranked algorithm returns a *pareto-optimal* solution (i.e., a solution where it is not possible to obtain any improvements without rendering any single contributor worse off).

THEOREM 5.2. *The ranked algorithm is correct w.r.t. the preference criterion pareto-util.*

Proof sketch The theorem follows from the fact that, once a solution is found, any future results that will be discovered will use at least one tuple clustering with a lower utility. \square

EXAMPLE 5.1. *Let $\alpha = 2$, the objective function be summation and the preference criterion be Pareto-util. We illustrate the ranked algorithm using the tuples and tuple clusterings in Example 3.1. By removing the tuple clusterings which cover less than 2 tuples and sorting the remaining ones, we get the sorted tuple clustering streams in Figure 12(a). Figure 12(b)-(d) show how the ranked approach run to get the $TC_{combined}$ which can cover all tuples in T . Then the tuples and tuple clusterings are fed to the construct_and_solve function.*

Given the input $TC_{combined} = \{tc_2, tc_3, tc_4, tc_6\}$ (see Figure 12(d)), the resulting summarization is the same to that in Figure 8(c). We note that the sum utility (Pareto-util) of this summarization is 1.85, which is not the Max-Sum-Util solution. In Section 6.2, we show that the Pareto-Util is a very close approximation of the Max-Sum-Util through experiments.

5.2 Ranked Set Cover

In [21], Park and Shim proposed a set cover based approximation algorithm, APPROX-FQ, for cell-based k -anonymity through value suppressions⁴. In particular, [21] shows that it is possible to obtain a $2(1 + \ln 2k)$ -approximation on the

⁴Note that a value suppression can be seen as an extreme degree of generalization, where the value is replaced with the root of the corresponding hierarchy.

TCStr ₁	TCStr ₂	TCStr ₃	TCStr ₄	TCStr ₅	TCStr ₆
tc ₂ (0.31)	tc ₁ (0.4)	tc ₄ (0.4)	tc ₇ (0.6)	tc ₆ (0.35)	tc ₃ (0.3)
tc ₆ (0.2)	tc ₃ (0.4)	tc ₂ (0.3)	tc ₄ (0.4)	tc ₇ (0.2)	tc ₇ (0.1)
tc ₁ (0.2)	tc ₂ (0.25)		tc ₆ (0.3)		

(a) Initially, each $TCStr_i$ is sorted in the ascending order of their clustering utilities.

TCStr ₁	TCStr ₂	TCStr ₃	TCStr ₄	TCStr ₅	TCStr ₆
tc ₂ (0.31)	tc ₁ (0.4)	tc ₄ (0.4)	tc ₇ (0.6)	tc ₆ (0.35)	tc ₃ (0.3)
tc ₆ (0.2)	tc ₃ (0.4)	tc ₂ (0.3)	tc ₄ (0.4)	tc ₇ (0.2)	tc ₇ (0.1)
tc ₁ (0.2)	tc ₂ (0.25)		tc ₆ (0.3)		

(b) After Step 3 and 4, $TC_{combined} = \{tc_1, tc_2, tc_3, tc_4, tc_6, tc_7\}$. No tuple clustering can cover more than at least 2 tuples.

TCStr ₁	TCStr ₂	TCStr ₃	TCStr ₄	TCStr ₅	TCStr ₆
tc ₂ (0.31)	tc ₁ (0.4)	tc ₄ (0.4)	tc ₇ (0.6)	tc ₆ (0.35)	tc ₃ (0.3)
tc ₆ (0.2)	tc ₃ (0.4)	tc ₂ (0.3)	tc ₄ (0.4)	tc ₇ (0.2)	tc ₇ (0.1)
tc ₁ (0.2)	tc ₂ (0.25)		tc ₆ (0.3)		

(c) The algorithm gets the tuple clustering with the next best utility value (0.4). The tuple clusterings that can summarize more than 2 tuples are tc_3 (for t_2 and t_6) and tc_4 (for t_3 and t_4). But not all tuples are covered by such tuple clusterings.

TCStr ₁	TCStr ₂	TCStr ₃	TCStr ₄	TCStr ₅	TCStr ₆
tc ₂ (0.31)	tc ₁ (0.4)	tc ₄ (0.4)	tc ₇ (0.6)	tc ₆ (0.35)	tc ₃ (0.3)
tc ₆ (0.2)	tc ₃ (0.4)	tc ₂ (0.3)	tc ₄ (0.4)	tc ₇ (0.2)	tc ₇ (0.1)
tc ₁ (0.2)	tc ₂ (0.25)		tc ₆ (0.3)		

(d) The algorithm gets the tuple clusterings with the next best utility value (0.3). So far, the tuple clusterings that can summarize more than 2 tuples are tc_2 (for t_1 and t_3), tc_3 (for t_2 and t_6), tc_4 (for t_3 and t_4), and tc_6 (for t_4 and t_5). These tuple clusterings satisfy the condition in Step 6a.

Figure 12: Running example for the ranked algorithm

amount of cell-suppressions by first generating a $(k, 2k - 1)$ -cover (where the sizes of the sets are between k and $2k - 1$) whose suppression length sum is at most $(1 + \ln 2k)$ times that of the optimal solution and then converting the resulting $(k, 2k - 1)$ -cover into a $(k, 2k - 1)$ -partition where all the sets are disjoint. While generating the $(k, 2k - 1)$ -cover for a given table, [21] selects the sets to be processed in the increasing order of *suppression-length per uncovered records*. To further speed up the cover generation, *APPROX-FQ* restricts processing to only those itemsets with at least k support.

The approach underlying the k -anonymization algorithm in [21] is not directly applicable to our problem: in particular, [21] does not consider a lattice (values are either taken as given or simply suppressed) or take into account utilities. In order to be applicable for α -summarization with lattices,

- instead of simply suppressing (values in the cells), the $(k, 2k - 1)$ -cover should generalize (tuples) based on the available lattice information, and
- instead of picking sets based on their suppression length sum, we need to pick sets (tuple clusterings) based on their utilities per uncovered tuple.

The first item above cannot be directly achieved in the suppression based *APPROX-FQ* framework: to merge two sets, *APPROX-FQ* simply unions the corresponding suppression lists. Thus, while *APPROX-FQ* can always return a k -anonymous solution simply increasing the amount of suppression, when generalizations are fed in an incremental manner, the set cover approximation may simply fail. Thus, the set cover may need to be called multiple times, once for each execution of Step 6a). If done in a naive way, this may cause significant overhead in terms of subproblems that are repeatedly re-evaluated as the set grows incrementally.

Input:

- A set $TC = \{tc_1, \dots, tc_m\}$ of tuple clusterings for all tuples $T = t_1, \dots, t_n$,
- Set of tuples $tuples(tc_i)$ for $\forall tc_i \in TC$,
- Summarization parameter α ,
- $TCStr_i$ -s, tuple clustering streams ($TCStr_i$ s) sorted in the descending order of utilities for each tuple in T (see Figure 9).

Output: A cover set $CSet$ such that

- $\forall t \in T, t \in C_i.tuples$ where $C_i \in CSet$. (i.e., every tuple is covered).
- $\forall C_i \in CSet, \alpha \leq |C_i.tuples| \leq 2\alpha - 1$. (i.e., satisfy a loose version of α -summarization).
- $\forall (C_i, C_j) | C_i \in CSet, C_j \in CSet, C_i.tuples \cap C_j.tuples = \emptyset$. (i.e., no tuple is covered by two sets).

ranked-set-cover($TC, tuples(tc_i)$ -s, $\alpha, TCStr_i$ -s)

1. $CSet' = \text{split-cover}(TC, tuples(tc_i)$ -s, $\alpha, T)$;
/*Get tuple covers such that for $\forall C \in CSet', \alpha \leq |C.tuples| \leq 2\alpha - 1$, but each pair (C_i, C_j) in $CSet', C_i.tuples \cap C_j.tuples$ may not be empty.*
2. $CSet = \text{cover-to-Partition}(CSet', \alpha, T, TCStr_i)$;
/*Get tuple covers such that for $\forall C \in CSet, \alpha \leq |C.tuples| \leq 2\alpha - 1$, and each pair (C_i, C_j) in $CSet, C_i.tuples \cap C_j.tuples = \emptyset$ */
3. /*From the set $CSet$, compute the final combined utilities*/
For (each $C \in CSet$)

- (a) Generalize every $t \in C.tuples$ to $C.tc$ with related utility $C.util_t$;

/* C represents a set cover. $C.tc$ is the corresponding tuple clustering, and $C.tuples$ is a set of tuples in T covered by $C.tc$. $C.util$ is the combined utility for generalizing tuple $t \in C.tuples$ to $C.tc$.*/

Figure 13: The main body of ranked-set-cover, called by Step 6a in Figure 9

split-cover($TC, tuples(tc_i)$ -s, α, T);

1. $CSet' = \emptyset; T_{covered} = \emptyset$;
2. while($T_{covered} \neq T$)
 - (a) Let $tuples(tc_i)_{uncovered} = tuples(tc_i) \cap (T - T_{covered})$; (i.e., $tuples(tc_i)_{uncovered}$ contains tuples that are not covered by any cover in $CSet'$)
 - (b) Choose tc_i such that $\frac{util(tc_i)}{\min\{|tuples(tc_i)_{uncovered}|, 2\alpha - 1\}}$ is minimized (or maximized);
/*minimized/maximized means choosing the minimum/maximum ratio; $util(tc_i)$ is the combined utility that tuples in $tuples(tc_i)$ are generalized to tc_i .*/
 - (c) if $|tuples(tc_i)| \leq 2\alpha - 1$ (i.e., tc_i covers less than $2\alpha - 1$ tuples)
 - i. Let $\mathcal{T} = tuples(tc_i)$;
 - (d) else if $|tuples(tc_i)_{uncovered}| > 2\alpha - 1$
 - i. Choose a subset \mathcal{T} of $tuples(tc_i)_{uncovered}$ s.t. $|\mathcal{T}| = 2\alpha - 1$
 - (e) else
 - i. Choose a subset \mathcal{T} of $tuples(tc_i)$ s.t. $|\mathcal{T}| = \max\{\alpha, |tuples(tc_i)_{uncovered}|\}$ and $tuples(tc_i)_{uncovered} \subseteq \mathcal{T}$;
 - (f) Let \mathcal{U} be the related utilities that tuples in \mathcal{T} are generalized to tc_i ;
 - (g) $T_{covered} = T_{covered} \cup \mathcal{T}$;
 - (h) $C_{new} = \langle tc_i, \mathcal{T}, \mathcal{U} \rangle$;
 - (i) $CSet' = CSet' \cup C_{new}$;

Figure 14: The split-cover procedure

On the other hand, we see that it is possible to leverage the algorithmic structure of the set cover solutions to focus the ranked processing to only those parts of the problem which require further evaluation and avoid repeated calls to set cover algorithm. Figures 13 through 16 provide the pseudocodes for the ranked set cover (RSC) algorithm to replace the optimization steps (Step 6a) of the ranked algorithm in Figure 9. Figure 13 provides the main body of the set cover algorithm. As in [21], the algorithm first creates a $(\alpha, 2\alpha - 1)$ -cover (but using tuple clusterings instead of cell-suppressions), then converts this cover into a partition by eliminating overlaps to create a α -summarized table.

```

cover-to-partition( $CSet, \alpha, T, TCstr_i$ )
1. for( $t \in T$ )
  (a) Let  $P$  contain all the  $C_i$ -s s.t.  $t \in C_i.tuples$  and  $C_i \in CSet$ ;
  (b) while  $|P| \geq 2$ 
    i. Randomly pick two covers  $C_i$  and  $C_j$  from  $P$ ;
    ii. if  $|C_i.tuples| > \alpha$  or  $|C_j.tuples| > \alpha$ 
      A. Remove  $t$  from the larger  $C.tuples$ 
      B. update  $C.util$ ;
      C. Remove  $C$  from  $P$ ;
    iii. else
      A.  $C_{new} = \text{get-best-common-tc}(C_i, C_j, TCstr_i)$ ;
      B. Remove  $C_i$  and  $C_j$  from  $P$ ;
      C. Insert  $C_{new}$  to  $P$ ;
2. return  $CSet$ ;

```

Figure 15: The *cover-to-partition* procedure

```

get-best-common-tc( $C_i, C_j, TCstr_i$ s)
1.  $C_{new}.tuples = C_i.tuples \cup C_j.tuples$ ;
2. for(each  $t \in C_{new}.tuples$ )
  (a) Let  $TCStr_t$  be the tuple clustering stream for  $t$ ;
  (b)  $TCstrPointer(t) = \min\{\text{position of } C_i.tc \text{ in } TCstr_i(t), \text{position of } C_j.tc \text{ in } TCstr_j(t)\}$ ;
  /*the position of a  $tc$  in  $TCStr(t)$  is the end of  $TCStr(t)$  if  $tc$  is not in  $TCStr(t)$ */
  Let  $TCStr'_t = \text{tuple clusterings appear after } TCstrPointer(t) \text{ in } TCStr_i$ ;
3. Identify tuple clustering  $tc_i$  whose combined utility computed from  $utils(tc_i)$  is maximal;
4.  $C_{new}.tc = tc_i$ ;
5.  $C_{new}.utils = utils(tc_i)$ ;
6. return  $C_{new}$ ;

```

Figure 16: The *get-best-common-tc* procedure

Figure 14 describes how a $(\alpha, 2\alpha - 1)$ -cover of the given set of tuples is created by iteratively picking those tuple clusterings with minimum or maximum utilities per uncovered tuples for merging. While this step reflects the outline of the $(k, 2k - 1)$ -cover algorithm in [21], there are two fundamental differences: (a) First of all instead of amount of suppression being used as a measure to select the sets to be merged, the algorithm in Figure 14 selects the tuple clusterings based on their degrees of utilities (Step 2b). (b) Secondly, as we will show in Section 6, instead of always picking sets in a way to *minimize* the suppression ratio per uncovered tuples (as in [21]), under different conditions, in Step 2b it may be preferable to *maximize* or *minimize* the utility ratio.

Figure 15 shows the pseudo-code for converting a given $(\alpha, 2\alpha - 1)$ -cover into a $(\alpha, 2\alpha - 1)$ -partition. The algorithm eliminates overlaps between sets included in the $(\alpha, 2\alpha - 1)$ -cover by either removing the tuples in the intersection (if there are sufficiently many other tuples in the sets to ensure α -summarization or (if sets are small) merging them under a new tuple clustering. Note that unlike [21] which simply unions the suppression lists to merge the given two tuple sets (each defined by a list of suppressions), in Step 1(b)iiiA, our cover-to-partition algorithm leverages the knowledge about available tuple clustering options to pick a high utility clustering that covers both clusterings.

The pseudo-code for this process is presented in Figure 16. Step 3 identifies a tuple clustering that covers the two sets of tuples using the tuple clustering streams that were provided as input to the ranked algorithm. Note that relying on the monotonicity of the utility merge functions, this step can be implemented using a ranked top-1 join algorithm [10, 11] on the corresponding (utility-sorted) tuple clustering streams.

6. PERFORMANCE EVALUATION

In this section, we evaluate the proposed ranked scheme for α -summarization. In the results presented in this section, we used real data set (“*Adults*” [15, 19] and [12]). The data




Figure 17: Time for tuple clustering utilities

set contains around $\sim 30K$ tuples. The two *Adults* hierarchies from [15] and [12] differ from each other. Thus, using these two sources as ground truth, we have created a combined *lattice*, where for any node in the lattice with two parents, the utilities of the clustering edge to the two parents are set to 0.25 as opposed to 0.5 for all other clustering edges. The utility of a path is computed by multiplying the utilities of the edges on the path. Experiments with other weighting strategies produced similar utility optimality and execution time behaviors. We compare the behavior of two different algorithms: *naive set cover (modified [21])*, and *ranked set cover* proposed in this paper. Experiments ran on a 2.9GHz Pentium (Windows XP) with 2GB main memory.

6.1 Off-line cost

Figure 17 shows the cost of tuple utility evaluation process. As it can be seen here, the time to compute utility evaluations is largely linear in the number of tuples for this data set. As expected, the computation time increases quickly with $|SA|$.

On the other hand, the tuple utility computations can be performed either off-line (and later re-used for different α -summarizations of the same table) or can be computed in a non-blocking, streaming manner and can be pipelined into the ranked algorithm in run time, without having to wait until all tuple utilities are available.

6.2 Summary utilities

Optimal and pareto-optimal summation utilities

As shown in Theorem 5.1, the ranked algorithm gets optimal results for the set cover formulation with respect to the min utility. However, for the sum utility, the ranked approach only gets a Pareto-optimal solution (Theorem 5.2). This experiment compares the differences between the optimal solution and the Pareto-optimal solution for adult data set⁵ with $|SA| = 3$. In Figure 18, *Sum-Util* represents the optimal solutions which are computed using the exhaustive process and an exact set cover algorithm; *Pareto-Util* denotes the Pareto-optimal solution computed using the ranked algorithm and an exact set cover algorithm (i.e., the *construct_and_solve* in Figure 9). We use Lingo optimizer (LINGO 9.0 solver) as the exact set cover algorithm by formulating the problem as Lingo script. In this experiment, the LINGO non-linear global solver is set to return results within 5% of the optimal. Since the Lingo problem formulation is not in the scope of this paper, it is omitted here.

Figure 18(a) shows the utilities of the Pareto-Util solutions and the Sum-Util solutions. It is apparent that the Pareto-optimal utility (Pareto-Util) is very close to the op-

⁵The results are similar for adult data set with $|SA| = 2$ as well as for synthetic data with different $|SA|$.

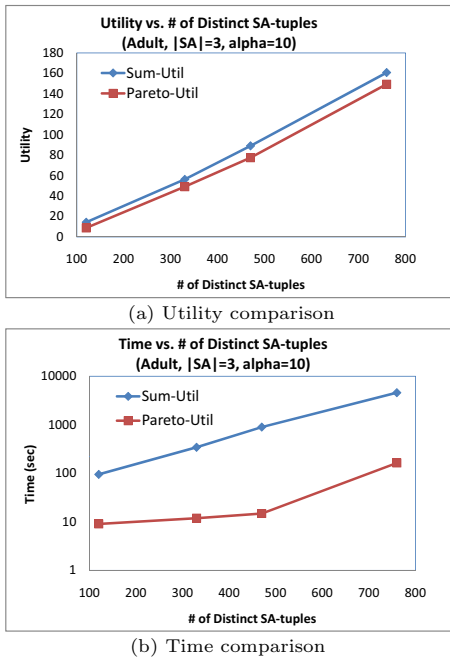


Figure 18: Results for comparing optimal sum-utils and pareto-optimal sum-utils. (a) Pareto-Util computed using the ranked algorithm is close to the optimal Sum-Util. (b) Pareto-Util method is more efficient than the Optimal Sum-Util method.

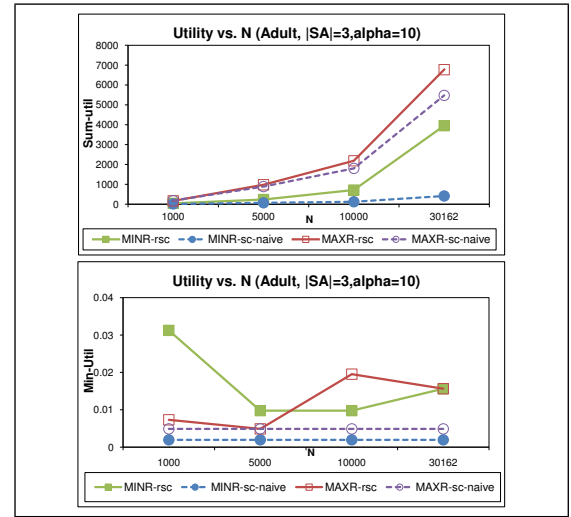
timal solution (Sum-util). Thus the Pareto-optimal solution is a good estimation of the optimal solution. On the other hand, Figure 18(b) exhibits that the Pareto-util solutions can be computed much more efficiently than the Sum-Util solutions. The execution time differs more when the number of tuples increases. Based on the above observations (a) Pareto-Util is similar to Sum-Util and (b) it is much more efficient to get Pareto-Util solution than to compute Sum-Util solution; thus, in the rest of the experiments, we only consider the Pareto-optimal solutions (for sum utility). I.e., the Max-sum always refers to the Pareto-Util.

Native set cover and ranked set cover:

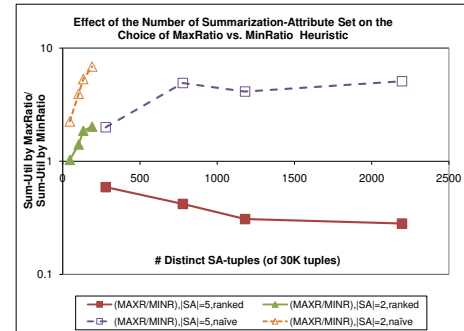
Figure 19(a) compares the ranked set cover approach against the naive-set cover, where all tuple clusterings covering at least α tuples are passed to a version of *APPROX-FQ* [21] modified to take into account clustering lattices and utilities.

The ranked set-cover algorithm performs consistently better than the naive set-cover algorithm, both for max-sum and max-min utility criteria. The *maximum ratio* (*MAXR*) heuristic (which prefers merging high-utility sets (i.e., tuple clusterings) earlier) performs better than the *minimum ratio* (*MINR*) heuristic which would merge low-utility sets earlier. While this particular observation is in line with the premise of [21], which picks clusterings with less suppressions early on, other experiments showed that the parallels with [21] do not always hold in the case of ranked set-cover evaluation. Consider Figure 19(b) which plots the relative overall sum-utility (*maximum ratio* / *minimum ratio*) as a function of $|SA|$ and the number of distinct *SA*-tuples⁶ in the *Adult* database. The naive set-cover algorithm indeed performs better with *maximum ratio* heuristic. On the other hand, especially for large $|SA|$ the ranked algorithms performs bet-

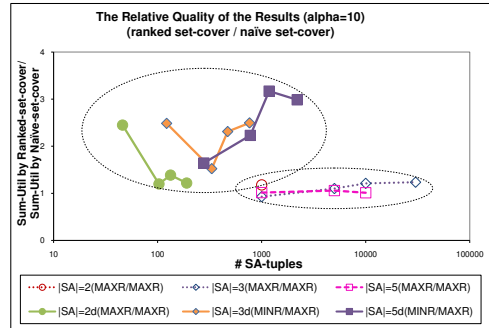
⁶To eliminate any chance of ambiguity due to redundant tuples, we focus only on the distinct *SA*-tuples



(a) Top: Max-sum-util criterion, Bottom: Max-min-util criterion



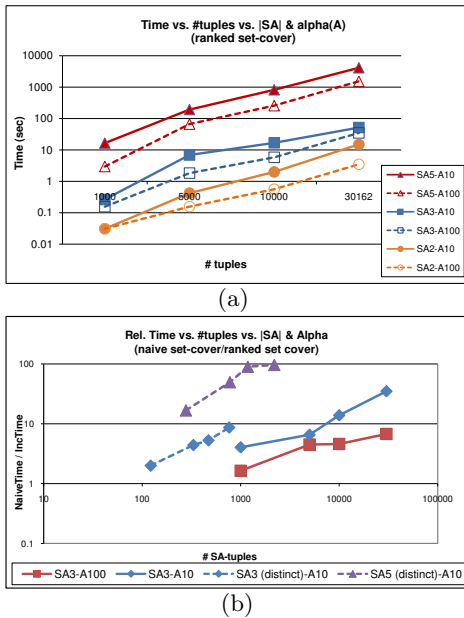
(b) Results with solid lines correspond to ranked set cover. Results with dashed lines correspond to naive set cover



(c) Results with solid lines correspond to cases with distinct *SA*-tuples. To be fair, for each case, the combination that provides the best results (according to (b)) is chosen

Figure 19: Utilities: Naive vs. ranked set cover

ter when sets to be merged are selected based on *minimum ratio* heuristic. This is because the ranked algorithm has not one, but two sources of clusterings: the first one is during *split-cover* (governed by utility-based set selection heuristic) and the second during *cover-to-partition*, where small-sized overlapping covers are merged into partitions through clusterings. Since MINR selects the tuple clusterings with small utilities, clusterings passed to *cover-to-partition* likely have low utilities, but are more general. Thus, at the second step, there is less need for significant clusterings. MAXR optimistically selects high utility covers, but these covers require significant adjustments during the *cover-to-partition* transition. This is more obvious when $|SA|$ is large (i.e., clusterings in second step may involve more loss).



(in naive set cover, SA5-A10, SA5-A100 failed to complete in 3hours)

Figure 20: Time: Naive vs. ranked set cover

Finally, Figure 19(c) compares the result qualities of the ranked set cover and naive set cover algorithms. The ranked set-cover algorithm provides significantly better results, especially when the input summarization-tuples are distinct. When the tuples have large degrees of duplicates, the relative qualities returned by the two algorithms become comparable. However, as we will see next, even in this case, the ranked set cover will pay off in terms of gains in execution time and result utilities.

6.3 Run-time cost

We have seen that the result qualities of the proposed ranked set-cover method can surpass those of naive set-cover algorithms. Figure 20 shows that this benefit does not come with any execution time overhead. The execution time of the ranked set cover algorithm is largely linear in the number of $|SA|$ -tuples (Figure 20(a)), while (especially when α is low and the input contains distinct summarization-tuples) the naive set cover algorithm quickly becomes multiple orders slower (upto 100X in these experiments).

7. CONCLUSIONS

In this paper, we formulated the problem of value clustering-based α -summarization of data in the presence of weighted value clustering hierarchies. We provided a mechanism to express alternative strategies and to account for various utility measures in summarization alternatives. We presented a novel ranked set cover based algorithm which reduces the complexity of the problem. Experiment results showed significant gains in execution time when the proposed ranked approach is used. We note that the utility measure considered here reflects the information loss due to the generalizations on the given value lattices. Our future work will also consider other measures of summarization effectiveness, including loss in the information theoretical sense as well as the diversity of the tuples in the summary.

8. REFERENCES

[1] G. Aggarwal, K. K. Tomas Feder, R. Motwani, R. Panigrahy, D. Thomas, and A. Zhu. Approximation algorithms for k-anonymity. *Journal of Privacy Technology*, 2005.

[2] R. Alfred and D. Kazakov. Data summarization approach to relational domain learning based on frequent pattern to support the development of decision making. *ADMA*, 2006.

[3] F. Buccafurri, F. Furfaro, D. Sacca, and C. Sirangelo. A quad-tree based multiresolution approach for two-dimensional summary data. In *SSDBM'2003*.

[4] A. L. Buchsbaum, D. F. Caldwell, K. W. Church, G. S. Fowler, and S. Muthukrishnan. Engineering the compression of massive tables: an experimental approach. In *SODA*, 2000.

[5] A. L. Buchsbaum, G. S. Fowler, and R. Giancarlo. Improving table compression with combinatorial optimization. *J. ACM*, 50(6):825–851, 2003.

[6] S. Chaudhuri and U. Dayal. An overview of data warehousing and olap technology. *SIGMOD Record*, 26(1):65–74, 1997.

[7] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. The MIT Press, 2001.

[8] A. Cuzzocrea, F. Furfaro, and D. Saccà. Hand-olap: A system for delivering olap services on handheld devices. In *ISADS '03*

[9] A. Cuzzocrea, D. Saccà, and P. Serafino. A hierarchy-driven compression technique for advanced olap visualization of multidimensional data cubes. In *DaWaK*, pages 106–119, 2006.

[10] R. Fagin. Combining fuzzy information from multiple systems. In *PODS*, 1996.

[11] R. Fagin, A. Lotem, and M. Naor. Optimal aggregation algorithms for middleware. *J. Comput. Syst. Sci.*, 2003.

[12] B. C. M. Fung, K. Wang, and P. S. Yu. Top-down specialization for information and privacy preservation. In *ICDE*, 2005.

[13] C. A. Hurtado, C. Gutierrez, and A.O.Mendelzon. Capturing summarizability with integrity constraints in olap. *ACM Trans. Database Syst.*, 30(3):854–886, 2005.

[14] C.Hurtado and A.Mendelzon. Reasoning about summarizability in heterogeneous multidimensional schemas. *ICDT* 2001.

[15] V. S. Iyengar. Transforming data to satisfy privacy constraints. In *Proc. of ACM SIG-KDD*, pages 279–288, 2002.

[16] K. LeFevre, D. J. DeWitt, and R. Ramakrishnan. Incognito: Efficient full-domain k-anonymity. *SIGMOD* 2005.

[17] M.-L. Lo, K.-L. Wu, and P. S. Yu. Tabsum: A flexible and dynamic table summarization approach. *ICDCS*, 00:628, 2000.

[18] A. Meyerson and R. Williams. On the complexity of optimal k-anonymity. *PODS* 2004.

[19] M. E. Nergiz and C. Clifton. Thoughts on k-anonymization. *Data Knowl. Eng.*, 63(3):622–645, 2007.

[20] W. K. Ng and C. V. Ravishankar. Relational database compression using augmented vector quantization. *ICDE* 1995.

[21] H. Park and K. Shim. Approximate algorithms for k-anonymity. In *SIGMOD '07*, pages 67–78, 2007.

[22] T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. Supporting imprecision in multidimensional databases using granularities. In *SSDBM*, pages 90–101, 1999.

[23] R. Rada, H. Mili, E. Bicknell, and M. Blettner. Development and application of a metric on semantic nets. *IEEE Transactions on Systems, Man and Cybernetics*, 19(1), 1989.

[24] P. Resnik. Semantic Similarity in a Taxonomy: An Information-Based Measure and its Application to Problems of Ambiguity in Natural Language. *JAIR*, Vol.11, 1999.

[25] R. Richardson and A.F. Smeaton. Using WordNet in an Knowledge-Based Approach to Information Retrieval. Working paper CA-1294, Dublin City Univ., Dublin, 1994.

[26] Y. Qi, K.S. Candan, and M.L. Sapino. Feedback-based InConsistency Resolution and Query Processing on Misaligned Data Sources. *SIGMOD* 2007

[27] Y. Qi, K.S. Candan, M.L. Sapino, and K.W. Kintigh. Integrating and Querying Taxonomies with QUEST in the Presence of Conflicts. *SIGMOD* 2007.

[28] Y. Qi, K. S. Candan, J. Tatemura, S. Chen, and F. Liao. Supporting olap operations over imperfectly integrated taxonomies. In *Proc. of ACM SIGMOD*, 2008.

[29] R. Saint-Paul, G. Raschia, and N. Mouaddib. General purpose database summarization. *VLDB* 2005.

[30] R. Saint-Paul, G. Raschia, and N. Mouaddib. Database summarization: The saintetiq system. *ICDE* 2007.

[31] P. Samarati. Protecting respondents' identities in microdata release. *TKDE*, 13(6):1010–1027, 2001.

[32] K.-L. Wu, S.-K. Chen, and P. S. Yu. Dynamic refinement of table summarization for m-commerce. *WECWIS* 2002.

[33] M. Rafanelli, A. Shoshani. STORM: A Statistical object representation model. *SSDBM90*.