
Algorithm for Interpretation of Multi-valued Taxonomic Attributes in Similarity-Based Fuzzy Databases

M. Shahriar Hossain and Rafal A. Angryk

Department of Computer Science, Montana State University,
Bozeman, MT 59717-3880, USA
{mshossain, angryk}@cs.montana.edu

Abstract. In this work we are analyzing our ability to discover knowledge from multi-valued attributes (often referred in literature on fuzzy databases as collections [1-3]), that have been utilized in fuzzy relational database models [4-7] as a convenient way to represent uncertainty about the data recorded in the data tables. We present here implementation details and extended tests of a heuristic algorithm, which we used in the past [8-11] to interpret non-atomic values stored in fuzzy relational databases. In our evaluation we consider different data imprecision levels, as well as diverse shapes of fuzzy similarity hierarchies.

Keywords: fuzzy relational databases, fuzzy collections, non-atomic symbolic values, data mining.

1 Introduction

The dynamic growth of data mining research and applications [12] carries significant importance for researchers working in areas related to fuzzy databases. Successful future of unconventional database models becomes more and more dependent on the development of consistent and time-efficient mechanisms for mining the data such databases are capable to store and process.

Fuzzy databases let the user reflect uncertainty about inserted information via the insertion of multiple descriptors in every column of the data table. At the same time, majority of currently available data mining algorithms allow the user to deal only with atomic attribute values. This leads to the challenge of developing a mechanism allowing consistent interpretation of the non-atomic values stored in the fuzzy databases. This process can be interpreted as mapping (i.e. defuzzification) of fuzzy tuples (containing collections of values) from tables in fuzzy relational databases into sets of atomic records, which are compatible with 1NF definition and can be analyzed using regular data mining techniques. We need to explain, that the term “defuzzification” used in this paper should be interpreted in a broader context than typically used in the fuzzy sets’ community (i.e. transformation of a fuzzy set to a single, crisp value). In this work we will focus on problem of interpretation of multi-valued, nonnumeric entries in the similarity-based fuzzy relational databases. We have decided to choose this topic for our research as these types of entries appear in large number of applications, starting from multiple online surveys about customers’ preferences (e.g. *mark*

types of food you like), and finishing on reports from medical examinations (e.g. *identify areas of pain, describe its severity, etc.*). In this paper we will present a heuristics allowing to transfer non-atomic symbolic values to the singleton forms, that can be then interpreted by many regular data mining algorithms. At the very beginning of this article, however, we want to emphasize that the point of this paper is not to argue about accuracy or efficiency of the approach that has been used, but rather to raise awareness of the problem, and to show that a defuzzification of uncertain data in similarity-based fuzzy databases is achievable. In the next section, we provide a brief review of the necessary background that covers the following areas: (1) fuzzy database model incorporating usage of non-atomic values to reflect uncertainty, (2) taxonomic symbolic variables and some of their properties, (3) our heuristics for interpretation of non-atomic entries in the fuzzy tuples. In the section 3, we present implementation details of our approach, discuss the artificial data sets we created for testing, and finally discuss results we obtained. Finally, in section 4, we briefly summarize conclusions coming from our investigation and point out new directions of future research.

2 Background

2.1 Fuzzy Database Model

There are two fundamental properties of fuzzy relational databases, proposed originally by Buckles and Petry [4-5] and extended further by Sheno and Melton [6-7]: (1) utilization of non-atomic attribute values to characterize features of recorded entities we are not sure of, and (2) ability of processing the data based on the domain-specific and expert-specified fuzzy relations applied in place of traditional equivalence relations.

In fuzzy database model it is assumed that each of the attributes has its own fuzzy similarity table, which contains degrees of similarity between all values occurring for the particular attribute. Such tables can be also represented in the form of the similarity hierarchies, named by Zadeh [13] *partition trees* that show how the particular values merge together as we decrease similarity level, denoted usually by α . Examples of partition trees for the domains *Country* and *Food-Type* are presented in the figures 1 and 2, respectively.

2.2 Taxonomic Symbolic Variables and Their Analysis

In [14], Bock and Diday included a comprehensive collection of recent works related to the extraction of statistical information from symbolic data. Based on the classification presented in the book, the data type discussed in this paper should be characterized as *generalized taxonomic symbolic variables*. These types of data entries are expected to carry no quantitative meaning, but yet the values can be nonlinearly ordered in a form of rooted, hierarchical tree, called by Diday [14] a *taxonomy*. Such variables are called *taxonomic* or *tree-structured* variables. Zadeh's partition trees [13], as presented in the figures 1 and 2, represent such taxonomies, allowing us to incorporate fuzzy similarity relations not only to data querying, as it has been used in the past, but also to the data analysis.

In [14], chapter six [15] has been devoted almost entirely to the problem of derivation of basic description statistic (in particular: medians, modes and histograms) from multi-valued symbolic data. The authors propose transformation of non-atomic values to collection of pairs in the format of (ξ, φ) , where ξ stands for a singleton (i.e. atomic) symbolic value, and φ represents ξ 's observed frequency. This is achieved by extending the classical definition of frequency distribution. For a multi-valued variable, the new definition states that the number reflecting the *observed frequency* (i.e. count of ξ 's appearances in the data set) can be a positive real, instead of a positive integer, as is the case for each single-valued variable.

The observed frequency distribution of a multi-valued variable Z , for a single data point, can be now defined as the list of all values in the finite domain of Z , together with the percentage of instances of each of the values in this data point. In [15], the authors mention that other definition of frequency distributions can be proposed, suggesting taking into account the natural dependencies between the symbolic values, as reflected by the provided attribute's taxonomy. This observation provided us with motivation for the work presented below.

2.3 Similarity-Driven Vote Distribution Method for Interpretation of Non-atomic Values

Following the rationale presented in the previous section, we attempted to develop a simple method to transfer non-atomic values in the fuzzy records to the collection of pairs including atomic descriptors and their fractional *observed frequencies*. In our work, however, we wanted to utilize background knowledge about attributes' values, which is stored in fuzzy databases in form of fuzzy similarity relations [13]. In this section we present a simple example to introduce our approach.

We want a reader to assume for a moment that he/she needs to find a drugs' dealer who, as a not-confirmed report says (i.e. our fuzzy tuple), was recently seen in $\{Canada, Colombia, Venezuela\}$. The most trivial solution would be to split the count of observed frequency equally among all inserted descriptors, that is to interpret the entry as the following collection $\{Canada|0.333, Colombia|0.333, Venezuela|0.333\}$. This approach however does not take into consideration real life dependencies, which are reflected not only in the number of inserted descriptors, but also in their similarity (represented by a taxonomy of attribute values, which reflects our pre-defined fuzzy similarity relation/table).

In our work we used a simple heuristics [8-11] letting us to replace the even distribution of a vote with a nonlinear spread, dependent both on the similarity of inserted values and on their quantity. Using the partition tree built from the fuzzy similarity table (grey structure in Fig. 1), we can extract from the set of the originally inserted values those concepts which are more similar to each other than to the remaining values. We call them *subsets of resemblances* (e.g. $\{Colombia, Venezuela\}$ from the above example). Then we use them as a basis for calculating a distribution of a database record's fractions. An important aspect of this approach is extraction of the *subsets of resemblances* at the lowest possible level of their common occurrence, since the nested character of fuzzy similarity relation guarantees that above this α -level they are going to co-occur regularly.

Our heuristic algorithm is pretty straightforward. Given (1) a set of values inserted as a description of particular entity’s attribute, and (2) a hierarchical structure reflecting Zadeh’s partition tree [13] for the attribute; we want to extract a table, which includes (a) the list of all subsets of resemblances from the given set of descriptors, and (b) the highest level of α -proximity of their common occurrence. We then use the list to fairly distribute fractions of the original fuzzy database record.

Our algorithm uses preorder recursive traversal for searching the partition tree. If any subset of the given set of descriptors occurs at the particular node of the concept hierarchy we store the values that were recognized as similar, and the corresponding value of α . An example of such a search for subsets of resemblances in a tuple with the values $\{Canada, Colombia, Venezuela\}$ is depicted in Fig. 1. Numbers on the links in the tree represent the order in which the particular subsets of similarities were extracted.

After extracting the *subsets of resemblances*, we apply a summarization of α values as a measure reflecting both the frequency of occurrence of the particular attribute values in the *subsets of resemblances*, as well as the abstraction level of these occurrences. Since during the search the country *Canada* was reported only twice, we assigned it a grade 1.4 (i.e. $1.0+0.4$). For *Colombia* we get: $Colombia(1.0 + 0.8 + 0.4) = Colombia|2.2$, and for the last value: $Venezuela(1.0 + 0.8 + 0.4) = Venezuela|2.2$.

At the very end we normalize grades assigned to each of the entered values: $Canada \ (1.4/5.8) = Canada \ |0.24$, $Colombia \ |(2.2/5.8) = Colombia \ |0.38$, $Venezuela \ |(2.2/5.8) = Venezuela \ |0.38$. This leads to the new distribution of the record’s fractions, which, in our opinion, more accurately reflects real life dependencies than a linear-split approach.

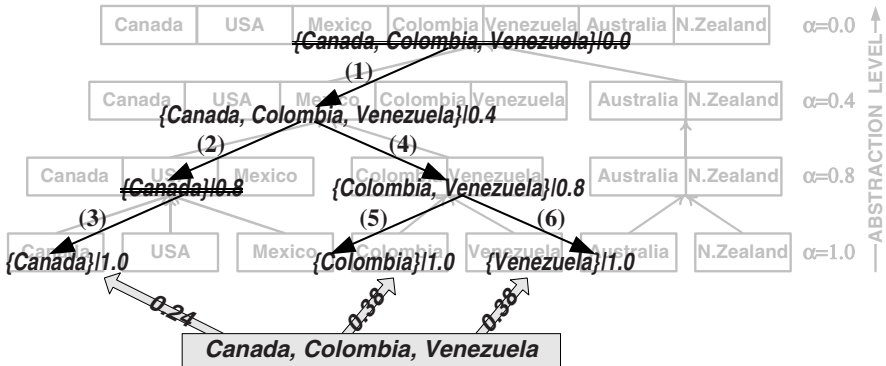


Fig. 1. Subsets of Resemblances extracted from the Partition tree

3 Implementation of the Defuzzification Algorithm

The implementation of our defuzzification system is based on a class named *TreeNode*, which contains only a single node of the fuzzy similarity hierarchy, pointers to its immediate descendants, and some public methods to access them. Each




Fig. 2. Object representation of a partition tree

descendant is an instance of the same class `TreeNode`. Hence, a `TreeNode` contains pointers to descendent `TreeNode`. Descendents that do not point to any other child node basically points to null `TreeNode` ($\{\Phi\}$). This is depicted in Fig. 2.

The `TreeNode` class is given below for better illustration.

```

public class TreeNode {
    int childCount = 0;
    Vector nodeData = new Vector();
    Vector childrenNodes = new Vector();
    /* THIS IS NOT A BINARY TREE. SO IT CAN CONTAIN ANY NUMBER
    OF CHILDREN. childrenNodes CONTAINS ELEMENTS OF TYPE
    TreeNode */
    public TreeNode (Vector colon) {
        /* CREATES A NEW INSTANCE OF TreeNode*/
        nodeData = (Vector) colon.clone();
        childrenNodes = new Vector(); }
    public void addChild(TreeNode childNode){
        childrenNodes.add(childNode);
        childCount++; }
    public int getTotalChildren() {
        /* e.g., IN THE CASE OF A NODE WITH THREE DESCENDENTS,
        THIS FUNCTION WOULD RETURN 3.*/
        return childCount; }
    public Vector getNodeVector () {
        /* RETURNS THE PARENT NODE */
        return nodeData; }
    public void setNodeAsLeaf () {
        childrenNodes.clear(); }
    public Vector getChildrenVector () {
        /* RETURNS THE CHILDREN NODES AS A VECTOR*/
        return childrenNodes; }
    public TreeNode getChild (int i){
        /* RETURNS i-th CHILD WHERE EACH CHILD ITSELF IS A
        TreeNode*/
        return ( (TreeNode) childrenNodes.get(i) ); }
} //END OF public class TreeNode

```

The basic algorithm uses preorder recursive tree traversal (depth-first search, DFS) for searching matching subsets in the partition tree. Our goal is to find at each node of the partition tree the largest matching subset. The high level outline of the algorithm is portrayed in the following code. `List` and `TreeNode` parameters of the `partitionTreeTraversal` function are passed by reference where other parameters are passed by value. The method `clone()` of an object returns entirely a new instance of the cloned object to avoid changes to a referenced parameter while the referenced parameter is necessary to be kept intact for the upper levels of recursion. A `Vector` can store a collection of entered attribute values, e.g. *{French, Greek, Italian}* and it can be used as an instance of `searchVector` of the algorithm. Besides, a `Vector` can also be a collection of other `Vectors`. On the other hand, a `List` is a collection of `Vectors` e.g., during the first call of `partitionTreeTraversal` the parameter `sList` should contain a sorted (descending order, based on the size of subsets) list of all the possible subsets of `searchVector` except the empty subset ($\{\Phi\}$). The operation denoted by “-” in the algorithm below is considered a regular *SetDifference* operation.

```

public Vector partitionTreeTraversal(List sList, TreeNode
masterNode,
searchVector, int level){
    List subsetList = sList.clone();
    int totalChildren = masterNode.getTotalChildren();
    Vector masterVector = masterNode.getNodeVector();
    Vector subsetVector = GET THE LONGEST SUBSET FROM subset-
List THAT                               IS FOUND IN THE ROOT OF
masterNode;

    if (subsetVector.size()==0){
        return { $\Phi$ };
    }
    else if (subsetVector.size()==searchVector.size()){
        UPDATE THE CORRESPONDING ENTRY FOR EACH ELEMENT OF
searchVector or      subsetVector WITH CORRESPONDING  $\alpha$ -VALUE
OF level-TH LEVEL ;
    }
    else{ //IF PARTIALLY AVAILABLE
        ADD CORRESPONDING  $\alpha$ -VALUE OF level-TH LEVEL TO THE
CORRESPONDING      ENTRY FOR EACH ENTITY OF subsetVector ;
    }

    Vector resultVector = searchVector.clone();
    for (int i=0; i<totalChildren; i++){
        Vector temp =
partitionTreeTraversal( subsetList, master-
Node.getChild(i),      subsetVector,
level+1 );
        subsetList = subsetList - (ALL SUBSETS THAT HAVE AT LEAST
ONE                      ENTITY OF temp);
        resultVector = resultVector - temp;
        if (resultVector == { $\Phi$ }) //NOTE:1

```

¹ Do not traverse other branches because all entities are already found in the previously traversed branches.

```

        break ;
    }
    if (totalChildren==0) /*IF masterNode IS A LEAF*/
        return subsetVector;
} //END OF public void partitionTreeTraversal

```

3.1 Analysis of the Algorithm

Let us assume that the average branching factor of the partition tree is b and the number of abstraction levels in the tree is d . If each node has b descendents, then the root (*level 1*) has 1 node, *level 2* has b nodes, *level 3* has b^{3-1} nodes, ..., *level d* has b^{d-1} nodes. Hence the total number of nodes, N in the tree is: $N = \sum_{k=1}^d b^{k-1}$, leading to

$Nb = \sum_{k=1}^d b^k$, where k reflects the current *level* of the *partitionTreeTraversal*

algorithm. The last equation produces the following result:

$Nb - N = N(b-1) = b^d - 1$, thus $N = \frac{b^d - 1}{b - 1}$. Therefore, in the worst case when the

searchVector is dispersed in all the leaf nodes of the partition tree at d^{th} level, *partitionTreeTraversal* has to visit a total of $\frac{b^d - 1}{b - 1}$ nodes which would re-

sult in a time complexity of $O(b^d)$ which should not be alarming despite of exponential growth, as typically $b \ll N$ and $d \ll N$ in real-life data mining. The worst case search scenario occurs when the imprecise attribute of the fuzzy record, stored in *searchVector*, contains every element of the root node of the partition tree; in our example – if the *searchVector* is $\{French, Greek, Italian, Chinese, Sushi, Barbeque, Burger\}$. In such case the traversal would occur to every subtree of the hierarchy of Fig. 2. Such situation, however, should appear rather rarely as this type of entries is used in fuzzy relational database only if total lack of knowledge concerning an attribute value occurs.

Obviously, the best case is when the entire *searchVector* is always found in the first descendent of every node along the traversal trail starting from the root. This case is possible only when the *searchVector* is atomic (as for example, $\{French\}$, considering the partition tree of Fig. 2) causing a traversal of a total of d nodes and would result in a time complexity of $O(d)$, where d denotes the depth of the tree. A heuristics of having the longest node (containing maximum number of elements in a certain level) always first may increase the probability that a concept is always discovered in the left subtree.

The function *partitionTreeTraversal* is written in such a manner that it avoids unnecessary branches while traversing the tree. For example, if the top level *searchVector* is $\{French, Sushi, Barbeque\}$ then the traversal would follow the path depicted in Fig. 3 where the DFS traversal is marked by numbers 1 to 10. The partition tree is shaded in Fig. 3 and the search vector is drawn in black. From this, it becomes evident that the average case time complexity of the algorithm is highly

dependent on the dispersed characteristic of the searchVector. At *level 1* the whole searchVector is found; the algorithm then searches the entire searchVector in the left-most branch of *level 1*, but it gets only a subset of searchVector, this traversal is propagated upto *level 4* to update the summarization of corresponding α -value for only relevant part of the original searchVector, i.e. for {*French*}.

Level 3 contains pointers to children that are { Φ }, hence a traversal of *level 4* indicates search success of an atomic concept in *level 3*. The algorithm then searches for {*Sushi, Barbeque*} in the middle branch at *level 2*, where again a subset {*Sushi*} is detected. The search fails in *level 3*, in the left branch as it contains only {*Chinese*}, the algorithm immediately returns to *level 2* for the next branch and succeeds for {*Sushi*} that goes upto *level 4*. The algorithm then continues the searching taking {*Barbeque*} alone, at the right most branch of *level 1* in the same way. It is evident from Fig. 3 that the unnecessary branches are always omitted by the partitionTreeTraversal algorithm. So the performance of the algorithm depends on the distribution of the searchVector in the partition tree. Average case is dominated by the probability of the existence of a searchVector in a certain node at certain level.

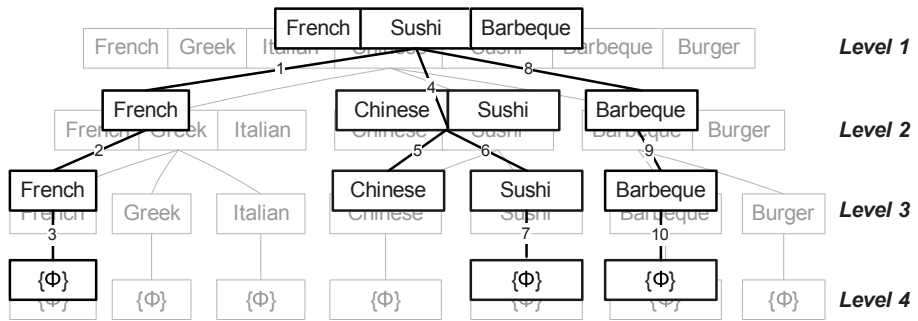


Fig. 3. Path traversed when the top level searchVector is {*French, Sushi, Barbeque*}. The partition tree is drawn in light gray. The DFS traversal is marked with numbers from 1 to 10.

3.2 Impact of the Character of the Partition Tree on the Algorithm Performance

The performance of the partitionTreeTraversal algorithm is influenced by two factors: (1) degree of imprecision reflected by the number of entered attribute values (i.e. by their number, and their similarity), and (2) the character of the partition tree specified for this attribute. This section includes some tests on the performance of the proposed algorithm when applied with different types of similarity trees.

Data Set and Partition Trees. The experiments have been conducted using some artificial datasets and different types of partition trees. Data were picked randomly from the domain (domain size = 32 symbolic values), where the domain is defined as the set of all values specified at the lowest (i.e. $\alpha=1.0$) level of the partition tree. Number

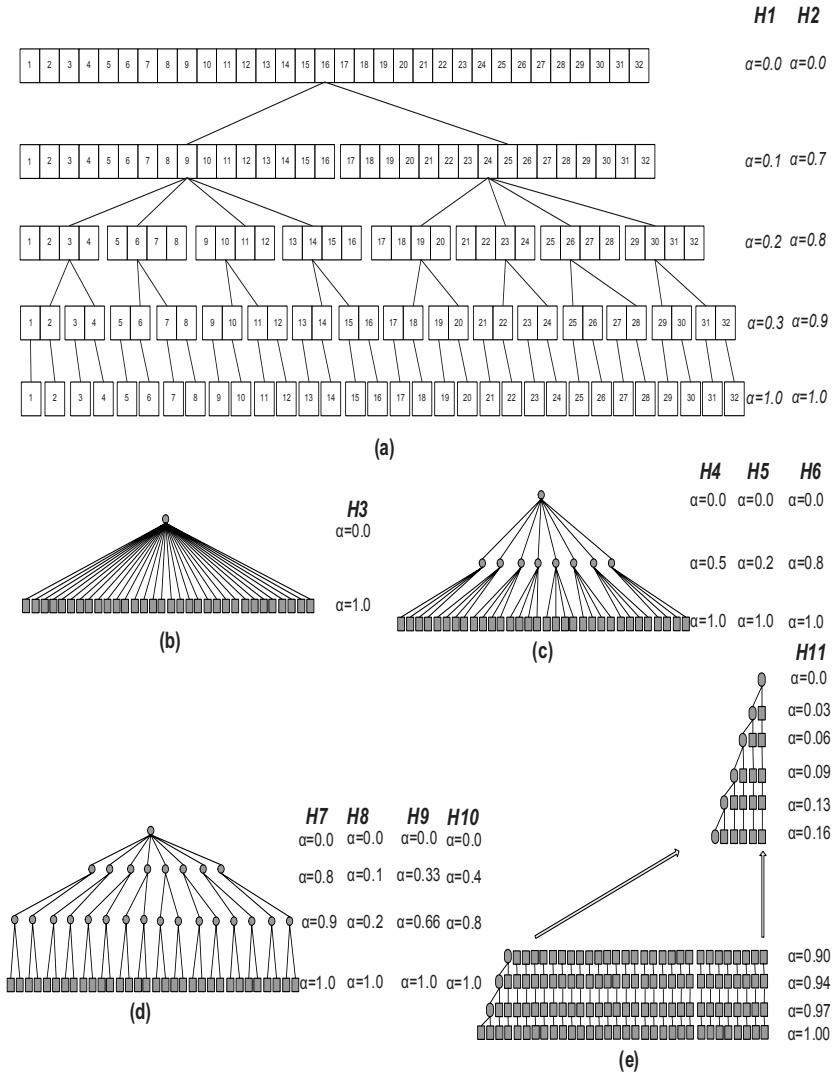


Fig. 4. Hierarchies used for testing. The hierarchies can be grouped as (a) H1, H2, (b) H3, (c) H4, H5, H6, (d) H7, H8, H9, H10 and (e) H11. Each group possesses a look-alike partition tree, but with different α -values.

of values in a fuzzy record (stored during defuzzification in `searchVector`), considered as 75% imprecise, was kept at the 75% the domain size (i.e. 24 values) to allow for the randomness in choice of values, when maintaining consistent number of values. The number of records in each imprecise data file was 30000, so that the random characteristic could be evenly distributed among the partition tree as the `partitionTreeTraversal` algorithm omits the unnecessary branches. Hence the

huge number of records ensures an average case traversal for the comparison purpose in this experiment. Algorithm run times for different percentage of imprecise records are recorded and plotted for different hierarchies.

The numeric values in the partition tree of Fig. 4 (a) are symbols to represent non-numeric descriptive values like all other partition trees of the previous sections. Fig. 4 (a) contains two hierarchies with different α -values ($H1$ and $H2$) although their level-wise construction is common. $H1$ is a distribution where the concepts are nested at the higher levels and $H2$ is a distribution where concepts are nested at the lower levels of the hierarchy according to the α -values. The same domain is used for different types of hierarchies in Fig. 4 (b), (c), (d) and (e) although the numeric symbols are not explicitly presented in the trees. Fig. 4 portrays a total of eleven similarity hierarchies ($H1, H2, \dots, H11$) which can be grouped in five categories and for each of the category only a single similarity tree is drawn. $H1$ and $H2$ have a total of five levels in their hierarchy but they differ in the α -values. Similarly, $H4, H5$ and $H6$ have three levels where $H7$ to $H10$ contain four levels in their hierarchies. The hierarchy $H3$ of (b) has only two levels (the root and the leaves), whereas the hierarchy $H11$ of (e) is a dendrogram with a total of 32 levels. Some levels in the hierarchy of Fig. 4 (e) have been omitted due to the limited space of this presentation. Obviously, Fig. 4 (e) presents the worst case possible among all the cases of Fig. 4 where every hierarchy contains 32 concepts in the lowest level of the hierarchy. The performance behavior is explained in the following subsection. Behavior of `partitionTreeTraversal` Algorithm. The `partitionTreeTraversal` algorithm is applied to different types of hierarchies of Fig. 4 with different percentage of imprecision of data. The time required for the algorithm to traverse the partition tree is proportional to the percentage of imprecision of the data. Fig. 5 illustrates timing with two look-alike hierarchies, $H1$ and $H2$ with different nested characteristic depending on the α -values of Fig. 4 (a). It is evident from Fig. 5 that the structure of these two hierarchies are almost the same because the `partitionTreeTraversal` algorithm traverses the same nodes while storing and updating reported existence of concepts in the hierarchy. This is despite the fact that the generated vote's fractions may be different, due to different propagation of α 's within the tree. The algorithm is tested using nine other hierarchies depicted in Fig. 4. Fig. 6 (a) is a plot of time at different percentage of imprecision of data with the hierarchies $H3$ to $H10$ of Fig. 4(a). It is evident from Fig. 6(a) that the distribution of α -values among the levels does not have significant impact on the `partitionTreeTraversal` algorithm. Because the algorithm traverses the partition tree discretely depending on the number of levels rather than traversing in a continuous domain of α -values. A different algorithm that traverses depending on the continuous domain of α -values (rather than depending on the number of levels) better performs for the hierarchies that have nesting at the top levels for a downward search in the partition tree as it discovers most of the levels at the beginning of the domain of α -values. The timing plots of $H4, H5$ and $H6$ follow almost the same trend in the graph. Similarly, $H7, H8, H9$ and $H10$ also produce similar graphs that are closely aligned in Fig. 6 (a). The trend of $H3$ is separate from the other lines as it is a different hierarchy and performs the best as it has least number of levels in its partition tree. As `partitionTreeTraversal` algorithm is a level dependent depth-first search, it performs the same both for the hierarchies nested at high levels and hierarchies nested at low levels. This behavior is reflected in Fig. 6 (a) because




Fig. 5. Comparison of timing with two hierarchies that have the same distribution of concepts but different α -values. The hierarchies *H1* and *H2* are drawn in Fig. 4(a).

the hierarchies with the same tree are clearly grouped in the plot. The plot has three groups because it includes three types of partition trees, (b), (c) and (d) of Fig. 4.

The three types of hierarchies are well separated in the plot of Fig. 6 (a).

Fig. 6 (b) shows that the hierarchy *H11* takes the worst possible time compared to all other hierarchies (*H3* to *H10*). From the experiment, it becomes apparent that `partitionTreeTraversal` is a number of levels-dependent search, where the performance decreases with increase of the number of levels in the partition tree which suggests flattening the partition trees (by making them more bushy), or inclusion of new concepts without increment in the number of levels.

Impact of Conceptualization at Higher Levels and Lower Levels. Let us consider the partition trees of Fig. 7 (a) and (b). The hierarchy of (a) starts splitting at high levels of the tree generating low α -values whereas (b) splits at the comparatively low




Fig. 6. (a) Plot for hierarchy *H3* to *H10* of Fig. 4 (b). Plot for hierarchy *H3* to *H11* of Fig. 4.

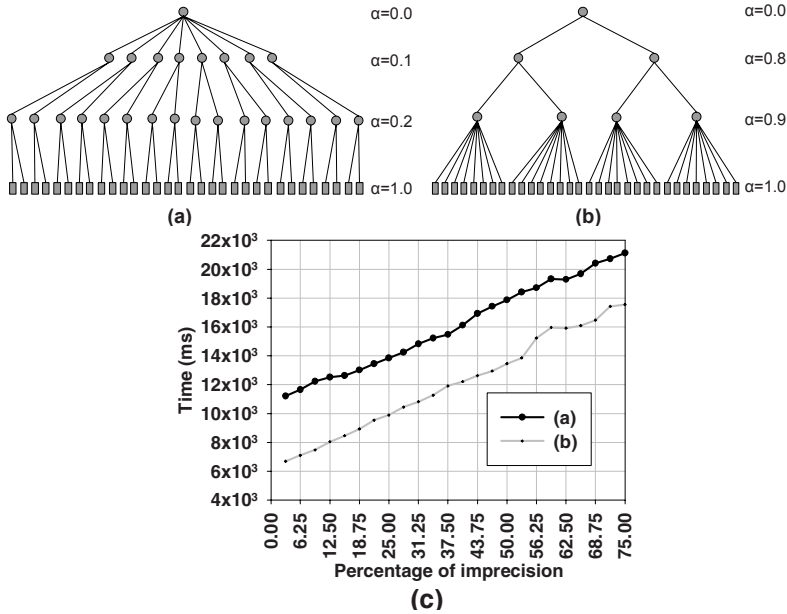


Fig. 7. (a) A partition tree where the abstracts start to split at high levels. (b) A partition tree where the abstracts split at low level. (c) Plot of runtime behaviors of `partitionTreeTraversal` algorithm with the hierarchies of (a) and (b) at different percentage of imprecision of data.

level. Although both the hierarchies have 32 leaf nodes at the bottom of the tree, in (b), they are concentrated fast at the bottom of the tree. In contrary, the other hierarchy possesses more concepts relatively at the top of the tree resulting in more intermediate nodes at every level. The `partitionTreeTraversal` algorithm would take more time in the case of Fig. 7 (a) compared to (b) as the algorithm needs to traverse more nodes for traversing the hierarchy of (a). The resulting runtime behaviors with these two hierarchies with respect to percentage of imprecision of data are plotted in Fig. 7 (c). The figure shows that the `partitionTreeTraversal` algorithm performs better with the hierarchy of Fig.7 (b) than that of (a). This suggests low branching factor near the root of the partition tree and comparatively higher branching factor at the bottom.

4 Conclusions

The work presented in this paper shows that the fuzzy collections can be transferred to the atomic values in the efficient way. The heuristic algorithm presented here allows for transfer of fuzzy records that reflect uncertainty to the form that allows analysis of such data via majority of regular data mining algorithms. Although, the presented heuristic for counting fractions of records as frequencies of observed atomic values might be questioned by statisticians, the technique in our opinion can be

treated as a “proof of concept” that imprecise data does not have to be disregarded, but can be mined in a regular manner. The task of discovery of the most appropriate data defuzzification algorithms remains as the topic for future research.

References

- [1] J.C. Cubero, N. Marín, J.M. Medina, O. Pons, M.A. Vila, “Fuzzy object Management in an Object-Relational Framework”, Proc. of 10th Intl. Conf. of Information Processing and Management of Uncertainty in Knowledge-Based Systems, 2004, pp.1767—1774.
- [2] Z. Ma (Ed.), “Advances In Fuzzy Object-oriented Databases: Modeling and Applications”, Idea Group Publishing, Hershey, PA, USA, 2004.
- [3] F. Berzal, N. Marín, O. Pons, M.A. Vila, “Development of applications with fuzzy objects in modern programming platforms”, International Journal of Intelligent Systems, 20 (11), 2005, pp. 1117 – 1136.
- [4] B.P. Buckles & F.E. Petry, “A fuzzy representation of data for relational databases”, Fuzzy Sets and Systems, 7(3), 1982, pp. 213-226.
- [5] F.E. Petry, “Fuzzy Databases: Principles and Applications”, Kluwer Academic Publishers, Boston, MA, 1996.
- [6] S. Shenoi and A. Melton, “Proximity Relations in the Fuzzy Relational Database Model”, International Journal of Fuzzy Sets and Systems, 31(3), 1989, pp. 285-296.
- [7] S. Shenoi, A. Melton, and L. T. Fan, “Functional Dependencies and Normal Forms in the Fuzzy Relational Database Model”, Information Sciences, 60(1-2), 1992, pp. 1-28.
- [8] R. Angryk, “Similarity-driven Defuzzification of Fuzzy Tuples for Entropy-based Data Classification Purposes”, Proc. of 15th IEEE Int. Conf. on Fuzzy Systems (FUZZ-IEEE '06), Vancouver, Canada, July 2006, pp. 1490-1498.
- [9] R. Angryk, “On Interpretation of Non-Atomic Values and Induction of Decision Rules in Fuzzy Relational Databases”, Proc. of 8th Intl. Conf. on Artificial Intelligence and Soft Computing, Zakopane, Poland, June 2006, in Series: Lecture Notes in Artificial Intelligence (LNAI), Vol. 4029, Springer-Verlag, 2006, pp. 170-181.
- [10] R. Angryk, F. Petry, R. Ladner, “Mining Generalized Knowledge from Imperfect Data”, Proc. of 10th Intl. Conf. on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU '04), Perugia, Italy, July 2004, pp. 739-746.
- [11] R. Angryk, F. Petry, “Discovery of Abstract Knowledge from Non-Atomic Attribute Values in Fuzzy Relational Databases,” in: B. Bouchon-Meunier, G. Coletti, R. Yager (Eds.), Modern Information Processing, From Theory to Applications, Elsevier, 2006, pp. 171-182.
- [12] M. Kantardzic, J. Zurada, “New Generation of Data Mining Applications”, IEEE Press and John Wiley, 2005.
- [13] L.A. Zadeh, “Similarity relations and fuzzy orderings”, Information Sciences, 3(2), 1970, pp. 177-200.
- [14] H.-H. Bock, E. Diday (Eds.), “Analysis of Symbolic Data, Exploratory Methods for Extracting Statistical Information from Complex Data”, Springer, 2000.
- [15] Bertrand, P. & Goupil, F. (2000), “Descriptive statistics for symbolic data”, in Analysis of Symbolic Data: Exploratory Methods for Extracting Statistical Information from Complex Data (eds. H.-H. Bock and E. Diday), Berlin, Springer-Verlag, pp 103-124.