Meta-Association Rules for Fusing Regular Association Rules from Different Databases

M. Dolores Ruiz*, Juan Gómez-Romero*, Maria J. Martin-Bautista*, *Member, IEEE*,
Daniel Sánchez*†, *Senior Member, IEEE*, Miguel Delgado*, *Senior Member, IEEE*
*CITIC-UGR, Dept. Computer Science and A.I. University of Granada (Spain)
†European Centre for Soft Computing, Mieres (Spain)
Email: {mdruiz, jgomez, mbautis, daniel}@decsai.ugr.es, mdelgado@ugr.es

***Abstract*—Association rules have been widely used in many application areas to extract from raw data new and useful information expressed in a comprehensive way for decision makers. Nowadays, with the increase of the volume and the variety of data, the classical data mining workflow is showing insufficient. We can expect in the near future that, more often than not, several mining processes will be carried out over the same or different sources, thus requiring extracted information to be fused in order to provide a unified, not overwhelming view to the user. In this paper we propose a new technique for fusing associations rules. The notion of meta-association rule is introduced for that purpose. Meta-association rules are association rules where the antecedent or the consequent can contain regular rules that have been previously extracted with a high reliability in a high percentage of the source databases. We illustrate out proposal with an example in the domain of crime data analysis.**

I. Introduction

One of the most widely accepted definitions of Information Fusion is due to Llinas and Hall [19], who state that it comprises theories and methods aimed to "combine data from multiple sensors and related information to achieve more specific inferences that could be achieved by using a single, independent sensor." However, it is also acknowledged that in recent years, researchers' focus has shifted from the *lower-level* problems to the *higher-level* problems, and the data sources considered go beyond the common sensor networks. The overarching objective of higher-level fusion is to interpret an observed situation and to evaluate the associated threat, thus providing support to decision-making. The concept of observable situation has been expanded to include any event of interest that can be identified from data sources. Correspondingly, these sources are also more diverse and heterogeneous, ranging from text sources to databases and other contextual information.

At the same time, the exponential growth of available data is posing new challenges to the Data Mining research area. Classical monolithic mining processes are no longer valid, because they have to run over very large datasets, which may even be spatially or temporally distributed. Therefore, datasets must be processed separately, which calls for the implementation of further fusion procedures to combine pieces of information obtained from each dataset into a common view. Our research work settles at the convergence of Data Mining and Information Fusion: Data Mining requires fusion procedures to merge knowledge discovered from different sources to globally understand the underlying relations between data samples, and Information Fusion can benefit from the mining techniques to incorporate refined knowledge extracted from diverse datasets in the higher-level.

Association rules is a well-established technique for mining information from structured databases. They allow the identification of novel, potentially useful and comprehensive knowledge in the form of implications $X \to Y$, which represent the joint co-ocurrence of X and Y in the database. They have been used in several application domains, such as retail market [4], [5], financial data [14], [21], crime and fraud detection [24], [25], sensor networks [3], [28], and elderly care [8], [23].

The motivation of this proposal is to overcome two problems that are usually overlooked by association rules algorithms. First, obtaining rules from very large datasets can be difficult and time-consuming. In that regard, parallel versions of rule mining algorithms can be found in the literature (see Section V). This is particularly necessary in Big Data problems, which involve volumes of data that cannot be managed with a traditional relational database. The MapReduce paradigm has been exploited to reduce the amount of data transferred through the network –one the most common drawback of parallel computing– in the mining process [22]. However, these works assume that the dataset is completely available at the beginning of the process. This is not the case when data is acquired in different time periods, and only the previously extracted rules are available when new data arrives. Second, it is more and more frequent having datasets with similar but not the same meaning, in such a way that they cannot be directly merged.

In both cases, the association rules algorithms work on the assumption that a fully-available uniform dataset is used. We propose in this paper a new technique for fusing association rules obtained from several databases. This is achieved by means of what we have called *meta-association rules*. The meta-association rules are built from the regular association rules previously extracted from each database, and can contain rules in the antecedent and/or in the consequent. The semantics of the meta-association rules are different from the regular rules, since they denote associations between associations. In addition, they can incorporate contextual information related to the original databases. We show that the meta-association rules can be obtained by using a two-stage version of a classical mining algorithm, and that they can convey new relevant information that cannot be obtained by regular rules.

The paper is structured as follows. In Section II we review the concept of association rule and the derived assessment measures. Section III describes our proposal for fusing association rules. Section IV shows the experiments that have been carried out by using open data about crimes in the city of Chicago. Section V presents some related works. Finally, Section VI points out some conclusions and prospective directions for future research.

II. Association Rule Mining

Association rules measure the frequent co-occurrence of attributes in a database. Formally, let D be a database formed by a set of transactions. The columns represent the possible attributes, and the rows represent transactions in which the attributes take values. The attributes can be categorical or quantitative. For categorical attributes we usually have linguistic values associated with the categories, which give us *items* of the form $\langle attribute, value\rangle$. For quantitative attributes we can use the same representation $\langle attribute, numerical\_value\rangle$. However, when attributes are continuous, we usually split them into meaningful intervals, thus obtaining items of the form $\langle attribute, interval\rangle$, which indicates that the value of the attribute lies in the interval. An *itemset* is then defined as a set of different items. For instance, when items i_1 and i_2 are satisfied in transaction t_1 we can say that the itemset $i_1 \wedge i_2$ is fulfilled in t_1.

In Table I we show an example of database, where v_{ij} represents the value of the attribute A_i in the transaction t_j. Any database expressed in this form can be transformed into a boolean database by considering the items as pairs $\langle attribute, value/interval\rangle$, as shown in Table II. In that database, we will have 1 when the item is present in the transaction and 0 when it is not.

TABLE I. Example of Transactional Database

D	A_1	A_2	$\cdots$	A_m
t_1	v_{11}	v_{21}	$\cdots$	v_{m1}
t_2	v_{12}	v_{22}	$\cdots$	v_{m2}
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
t_n	v_{1n}	v_{2n}	$\cdots$	v_{mn}

Accordingly, as defined in [1], an association rule is an "implication" of the form $A \to B$ that relates the presence of itemsets A and B in transactions of D, assuming that I is the set of items, $A, B \subseteq I$, $A \cap B = \emptyset$ and $A, B \neq \emptyset$. The left part of the rule is known as the antecedent of the rule, and the right part as the consequent.

The intensity of the association rule above is usually measured by the ordinary measures of support and confidence, also proposed in [1].

The *support* of a rule is the percentage of transactions satisfying both parts of the rule (i.e. the joint probability $P(A \cup B)$):

$$\mathrm{Supp}(A \to B) = \frac{|\{t \in D \mid A \cup B \subseteq t\}|}{|D|} \tag{1}$$

and the *confidence* of a rule measures the proportion of transactions that satisfying the antecedent, also satisfy the consequent (i.e. the conditional probability $P(B|A)$) :

$$\mathrm{Conf}(A \to B) = \frac{|\{t \in D \mid A \cup B \subseteq t\}|}{|\{t \in D \mid A \subseteq t\}|} \tag{2}$$

The general framework for the extraction of association rules consists in finding those rules whose support and confidence exceed the minimum thresholds $minsupp$ and $minconf$ imposed by the user. In such case, we say that $A \to B$ is *frequent* if $\mathrm{Supp}(A \to B) \geq minsupp$, and *confident* if $\mathrm{Conf}(A \to B) \geq minconf$. Following the notation used in the literature, we define the concept of strong rule:

Definition 1. [2] An association rule $A \to B$ is *strong* if it exceeds the minimum thresholds $minsupp$ and $minconf$ imposed by the user; i.e., if $A \to B$ is frequent and confident.

An alternative framework was proposed in [2], where the accuracy is measured by means of Shortliffe and Buchanan's certainty factors [27] as follows:

Definition 2. [7] Let $\mathrm{supp}(B)$ be the support of the itemset B, and let $\mathrm{Conf}(A \to B)$ be the confidence of the rule. The *certainty factor* of the rule, denoted as $CF(A \to B)$, is defined as

$$\begin{cases} \dfrac{\mathrm{Conf}(A \to B) - \mathrm{supp}(B)}{1 - \mathrm{supp}(B)} & \text{if } \mathrm{Conf}(A \to B) > \mathrm{supp}(B) \\ \dfrac{\mathrm{Conf}(A \to B) - \mathrm{supp}(B)}{\mathrm{supp}(B)} & \text{if } \mathrm{Conf}(A \to B) < \mathrm{supp}(B) \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

where $\mathrm{supp}(B)$ is the proportion of transactions in D satisfying B, i.e. $\mathrm{supp}(B) = |\{t \in D \mid B \subseteq t\}| \,/\, |D|$.

The certainty factor yields a value in the interval [-1, 1] that measures how the belief that B is in a transaction changes when it is known that A is in that transaction. Positive values indicate that the belief increases, negative values mean that the belief decreases, and 0 means no change. The certainty factor measure has better properties than confidence and other quality measures (see [9] for more details), and helps to solve some drawbacks of the confidence measure [2], [7]. In particular, it reduces the number of rules obtained by filtering rules corresponding to statistical independence or negative dependence.

Analogously to the confidence measure, we will say that $A \to B$ is *certain* if $\mathrm{Supp}(A \to B) \geq minCF$, where $minCF$ is the minimum threshold for the certainty factor given by the user. In the literature we can find as well the notion of very strong rule, which is defined as follows:

Definition 3. [2] An association rule $A \to B$ is *very strong* if both rules $A \to B$ and $\neg B \to \neg A$ are strong.

In [2] it is proved that the certainty factor has the following property: $CF(A \to B) = CF(\neg B \to \neg A)$. This means that when using the certainty factor, a strong rule is also very strong. Therefore, the definition of strong rules can be reformulated in such case: a strong rule is a rule that is frequent and certain.

TABLE II. EXAMPLE OF BOOLEAN DATABASE

D	$\langle A_1, v1_1\rangle$	$\langle A_1, v1_2\rangle$	$\cdots$	$\langle A_1, v1_s\rangle$	$\langle A_2, int_1\rangle$	$\cdots$	$\langle A_2, int_l\rangle$	$\cdots$	$\langle A_m, vm_p\rangle$
t_1	1	0	$\cdots$	0	1	$\cdots$	1	$\cdots$	0
t_2	0	1	$\cdots$	1	1	$\cdots$	0	$\cdots$	0
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
t_n	1	1	$\cdots$	0	0	$\cdots$	0	$\cdots$	1

III. FUSING INFORMATION WITH META-ASSOCIATION RULES

In this section we present our proposal for fusing association rules mined from independent databases, and the kind of knowledge represented by the resulting meta-association rules.

A. Architecture

The general process flow of our method is depicted in Fig. 1. It encompasses two sequential steps. First, we obtain association rules from each database by using the framework explained in the previous section. Second, we fuse this information by finding meta-association rules.

Step 1. Let $D_1, D_2, \ldots, D_k$ be k databases that may share some of their attributes. In the example described in Section IV, each dataset includes the crime incidents happened in each district of the city of Chicago. After applying the extraction procedure, we obtain k sets of association rules $R_1, R_2, \ldots, R_k$ from $D_1, D_2, \ldots, D_k$. The number of obtained rules can be different in each case. We note them by $nr_1, nr_2, \ldots, nr_k$ respectively. (To simplify the formulation, we assume that the same thresholds for the support and the confidence/certainty factor are used for each dataset.) Notice that there may be some common rules in the sets $R_1, R_2, \ldots, R_k$.

Step 2. In the second phase, we create a new database including all the rules obtained in the first phase (in binary table format). For that, we consider that each dataset is a transaction, and every rule obtained in the first phase is an item. In addition, if metadata about the original databases is available, it will be incorporated as new items. We refer to this new database as the meta-database, noted by $\mathfrak{D}$. The general composition of $\mathfrak{D}$ is shown in Table III, where $r_1, \ldots, r_n$ are the different rules found in $D_1, \ldots, D_k$ and $at_1, \ldots, at_m$ denote the additional information that can be added in order to obtain more interesting rules. Following the crime data example, where each initial dataset refers to a city district, the additional information can be attributes about district features, such as the number of residents or the security perception index. Afterwards, meta-association rules are extracted as described in the next section.

TABLE III. META-DATABASE OBTAINED AFTER COMPILING THE ASSOCIATION RULES OBTAINED FROM EACH DATASETS.

$\mathfrak{D}$	r_1	r_2	$\cdots$	r_n	at_1	$\cdots$	at_m
D_1	1	1	$\cdots$	0	1	$\cdots$	1
D_2	0	1	$\cdots$	0	0	$\cdots$	1
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
D_k	1	0	$\cdots$	1	1	$\cdots$	0

B. Meta-Association Rule Extraction

Meta-association rules are association rules that represent the frequent co-occurrence of items in the meta-database. Since items in the meta-database are rules r and attributes at, the meta-association rules represent the co-occurrence of rules, rules and attributes, or attributes. As for regular rules, the information given by the meta-rules depends on the interestingness measure used. We will use the certainty factor measure, thus obtaining very strong rules, and consequently, more reliable rules than the ones extracted by using the confidence measure.

Formally, we will obtain three types of meta-association rules:

- $r_i \to r_j$ where r_i, r_j can be rules or a conjunction of rules; for example: $r_i = r_{i1} \wedge \ldots \wedge r_{is}$.
- $at_i \to at_j$ where at_i, at_j can be attributes or a conjunction of attributes.
- $r_i \to at_j$ or $at_j \to r_i$ where r_i, at_j are a conjunction of rules and attributes, and can be mixed; i.e. a rule of the form $r_1 \wedge at_2 \to r_3$ can be found.

The third type of rules are the most interesting, because they relate the co-occurrence of two rules joint with an attribute. This leads to a new piece of information that cannot be obtained by using regular rules on the whole dataset. Some examples and explanations of the semantics of this kind of rules are presented in Section IV.

C. Algorithm and Implementation

The complete algorithm is described in Algorithm 1. It performs the processes of mining the set of rules from the original databases $D_1, \ldots, D_k$, creating the meta-database $\mathfrak{D}$, and mining the meta-association rules.

Similarly to other works in association rules mining, the itemset representation is based on bit strings [10], [20]. Bit strings speed up logical operations with boolean data. Our implementation uses the Java class `java.util.BitSet`. Each dataset is encoded with an array of `BitSet` objects of size equal to the number of attributes, each object corresponding to one transaction. Regarding time requirements, the complexity depends on the total number of transactions and the number of items considered in each step. The first step is $\mathcal{O}(n \times i)$ (being n the number of transactions and i the number of different items), whereas the second step is $\mathcal{O}(k \times (m + r))$ (being k the number of databases, m the number of extended items, and r the number of rules obtained in the first step). In practice, we have experienced quite fast executions of the experimental tests with different choices for the $minsupp$ and $minCF$. Besides, the first step can be performed in various parallel threads. Regarding memory requirements, the footprint is not

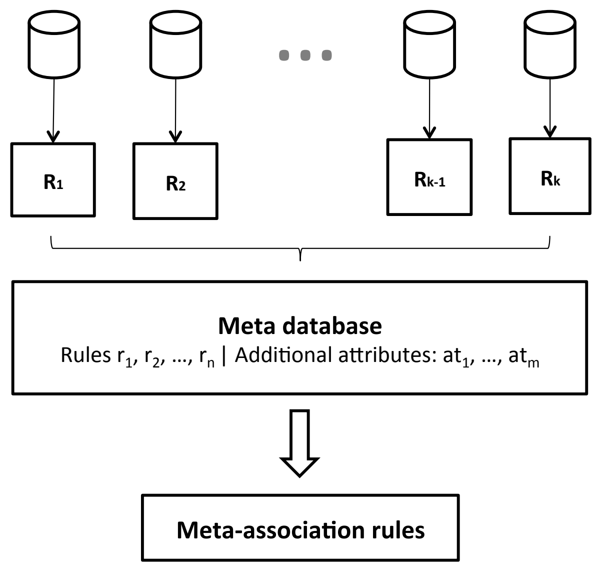

Fig. 1. Process flow: Information fusion by means of meta-association rules

---

Algorithm 1 Meta-association rules mining procedure

Input: $D_1, \ldots, D_k$, $minsupp$, $minCF$
Output: MR (set of meta-association rules)

1: **for all** D_i such that $1 \leq i \leq k$ **do**
2:     # D_i *preprocessing*
3:     Read D_i and store the items I
4:     Transform D_i into a boolean database
5:     Store database into a vector of `BitSets`

6:     *# Mine very strong rules*
7:     Compute the candidate set C of frequent itemsets $Supp(X) \geq minsupp$
8:     Store the `BitSet` vector indexes of $X \in C$ and $Supp(X)$
9:     Compose the rule with $X, Y \in C$

10:     **if** $Supp(X \Rightarrow Y) \geq minsupp$ **and** $CF(X \Rightarrow Y) \geq minCF$ **then**
11:         The rule is a very strong rule
12:     **end if**

13: **end for**

14: # $\mathfrak{D}$ *creation*
15: Compile all different rules from $R_1, \ldots, R_k$
16: Create $\mathfrak{D}$ using compiled rules and additional attributes

17: # *Mining meta-association rules*
18: Repeat steps 1-13 to mine meta-association rules from $\mathfrak{D}$

---

high, since the `BitSet` implementation optimizes bit storage and manipulation.

IV. Experimental Evaluation and Discussion

A. Experiments

For the experimental evaluation of our proposal, we have used a dataset of crime incidents in the city of Chicago in 2012[1]. The selected attributes are the following:

- Quarter of the year in which the incident happened.
- Day period: morning, afternoon, evening, night.
- Crime description according to police standard protocols; e.g., ‘crime of 500$ and under’.
- Location description: street, residence, etc.
- Arrest, if there is an arrest associated to the crime (true or false).
- Domestic, if the crime happened in a domestic environment (true or false).

We have split the dataset in 22 databases according to the districts of the city. The size of the databases is described in Table IV.

Since we want to study the possible relation between crimes in a neighborhood and the educational system, in the meta-database we have considered the following attributes of schools aggregated by district:

- Number of students in the district: low, medium, high, very high.
- Number of misconducts notified in the district: low, very low, medium, high, very high.
- Perceived safety index, obtained by means of surveys: low, medium, high.

[1] https://data.cityofchicago.org/

TABLE IV. DESCRIPTION OF THE DATABASES

District Id	transactions	items
01	12 160	268
02	13 481	286
03	17 730	297
04	19 847	302
05	15 296	290
06	19 301	303
07	20 195	298
08	22 493	333
09	16 704	305
10	15 071	301
11	21 873	311
12	15 930	296
14	12 580	296
15	14 434	285
16	10 851	296
17	9 702	282
18	14 247	272
19	15 689	284
20	5 694	252
22	10 797	275
24	9 538	267
25	19 726	301

For the experiments, we have used a 1.73GHz Intel Core 2 Duo notebook with 4GB of main memory, running Java on Windows 7. We have limited the rules obtained in the first step of the process to have one item in the antecedent and one in the consequent in order to obtain readable rules. For the meta-association rules, we allow two items at most in the antecedent and the consequent. Experiments using different values for the $minsupp$ and $minCF$ thresholds have led us to select $minsupp = 0.1$, $minCF = 0.5$, which have produced 80 meta-association rules.

We show in Table V some non-intuitive rules, which offer the most interesting results. For instance, the second meta-association rule states that there is a co-occurrence relation with high certainty ($CF = 1$) between the rules `(Crime-Description=<$500 under>→ Arrest=false)` and `(Location=RESIDENCE→Arrest=false)`, and the high `safety-index` in a district. That means that it is frequent to have a high perception of security when there are crimes of minor relevance at residences without arrests. This may suggest that, in those districts in which minor crimes are prevalent, the safety perception tends to be high when there is a reduced number of arrests. The last rule is also interesting, although it has a lower certainty factor. In this case, in some districts (13.6%) a higher safety perception (medium) is frequently associated to the fact that crimes are happening in the streets and there are several students in the district.

B. Discussion

This paper is an initial proposal for fusing association rules and additional information from different databases. One of the first limitations of the framework is that, to start the process of meta-association rules extraction, we only consider if a rule has been previously mined from a dataset or not, but not its measures of support and certainty factor. For instance, in the present proposal, a rule found with $CF = 0.5$ in a dataset and the same rule with $CF = 1$ in other dataset are considered as equal when constructing the meta-database.

Therefore, it could be convenient to consider such measures, at least the certainty factor, as a degree of importance of the rule when compiling the meta-database. If we take into account the CF as a degree of importance, we will obtain continuous attributes with values in the interval $[minCF, 1]$. In order to add the rule and the CF as an attribute of the meta-database, one option is to discretize that interval. However, this has some problems due to the crisp boundaries of intervals. For instance, given the values $CF_1 = 0.69$ and $CF_2 = 0.70$ and the intervals [0.6,0.7) and [0.7,0.8), the values would lie in different intervals in spite of they are very similar.

This issue motivates the use of a different representation for these values. Fuzzy sets seem to be an adequate option. There are some proposals in the literature dealing with fuzzy transactional databases, in which the items are satisfied to some extent. This is represented by a degree of fulfillment in the unit interval [7], [11]. Fuzzy databases support the extraction of fuzzy association rules from the continuous representation of values. These rules are also more understandable, because we can divide the intervals into meaningful segments, such as $LowCF, MedCF$ and $HighCF$, which can be represented by fuzzy sets.

Other interesting issue to be addressed is that of considering databases from various sources that will have a different structure. This may lead to the problem of having different specifications of the same item, for instance if we have the item ⟨`Crime-Description,< $500`⟩ in a database and in another database with a different structure ⟨`Crime-quantity,< $500`⟩. They seem to refer to the same item but their descriptions are different since they come from different sources. One interesting way to face this problem could be to use a knowledge repository assisting the algorithm in matching these items.

V. RELATED WORKS

The solution presented in this paper can be of interest in the application areas of association rules mentioned in Section I. To the best of our knowledge, this is the first attempt to develop an algorithm to fuse rules mined from different databases. Most related works are focused on mining association rules in distributed databases; i.e., they extract a single set of rules from the distributed data [6], [15], [16], [26].

Our approach resembles certain aspects of [13], where the author proposes general measures for comparing two sets of rules, namely rule overlapping, average support difference, and average confidence difference. Similarly, there are some notable works aimed at comparing expert users' knowledge with that extracted in the form of rules. In that regard, we can highlight the proposals in [17], [18], where users' knowledge and their impressions are captured in the form of rules (or similar structures), and then, these rules are compared to the rules extracted from data. The authors also give a classification of comparison results that includes conforming rules, rules with unexpected consequences, and rules with unexpected

TABLE V. EXAMPLES OF META-ASSOCIATION RULES FOUND IN THE CITY OF CHICAGO DATASET

Antecedent	Consequent	Supp	CF
(Crime-Description=<$500 under>→Domestic=false) **AND** Safety-Index=High	(Crime-Description=<$500 under>→Arrest=false)	0.227	1
(Crime-Description=<$500 under>→Arrest=false) **AND** (Location-Description=RESIDENCE→Arrest=false)	Safety-Index=High	0.136	1
(Location-Description=STREET→Domestic=false) **AND** Number-of-Students=Low	(Crime-Description=<$500 under>→Arrest=false)	0.136	0.633
(Crime-Description=<$500 under>→Arrest=false)	Safety-Index=High	0.227	0.516
(Location-Description=RESIDENCE→Arrest=false) **AND** Safety-Index=Low	Number-of-Misconducts=High	0.136	0.656
Safety-Index=Medium	(Location-Description=STREET→Domestic=false) **AND** Number-of-Students=Very High	0.136	0.511

reasons. These concepts could be used to compare the rules obtained from each database, and to refine the procedure to combine rules.

More recently, we can find similar contributions in other data mining fields addressing the problem of combining information from users and databases. For example, the proposal in [12] applies a semi-unsupervised algorithm to incorporate knowledge in the clustering process. Similarly, it would be interesting to study how these techniques can be used to improve the fusion of association rules.

VI. CONCLUSIONS AND FUTURE WORK

Fusion techniques are necessary in modern data mining problems, because frequently they need to combine information from different sources. In this paper we have focused on the problem of fusing association rules extracted from different databases. The proposed algorithm is based on a two-pass extension of a classical mining procedure, and gives as a result a set of meta-association rules which may contain regular rules in the antecedent and the consequent. In addition, the meta-rules can also consider available metadata describing the initial databases. The main advantage of the new meta-rules is that they can extract information that is not achieved by simple regular rules. The obtained results are promising, and suggest that the technique can be applied in several areas. More experimentation in that regard is still necessary.

The approach presented here has been tested in a set of coherent databases sharing the same set of attributes. In the near future we plan to enhance the proposal by also considering an heterogeneous set of databases. The algorithm can also be improved by incorporating the theory of fuzzy sets in order to obtain a fuzzy meta-database that considers the importance of the basic rules. This extension, as well as a more comprehensive comparison of results, is another prospective direction for future work.

ACKNOWLEDGEMENT

The research reported in this paper was partially supported by the project PYR-2014-24 from the GENIL program of the CEI BioTic GRANADA; the Andalusian Government (Junta de Andalucía) under projects P11-TIC-7460 and P10-TIC-6109; by the Spanish Ministry for Science and Innovation by the project grants TIN2012-30939; and by the project FP7-SEC-2012-312651, funded from the European Union in the Seventh Framework Programme [FP7/2007-2013] under grant agreement No 312651.

REFERENCES

[1] R. Agrawal, T. Imielinski, and A. Swami. Mining associations between sets of items in massive databases. In *ACM-SIGMOD International Conference on Data*, pages 207–216, 1993.

[2] F. Berzal, M. Delgado, D. Sánchez, and M.A. Vila. Measuring accuracy and interest of association rules: A new framework. *Intelligent Data Analysis*, 6(3):221–235, 2002.

[3] A. Boukerche and S. Samarah. An efficient data extraction mechanism for mining association rules from wireless sensor networks. In *Int. Conference on Communications, ICC'07*, pages 3936–3941, Seattle, 2007.

[4] S. Brin, R. Motwani, J.D. Ullman, and S. Tsur. Dynamic itemset counting and implication rules for market basket data. *SIGMOD Record*, 26(2):255–264, 1997.

[5] M. Chen, A. Chiu, and H. Chang. Mining changes in customer behavior in retail marketing. *Expert Systems with Applications*, 2005.

[6] D.W. Cheung, V.T. Ng, A.W. Fu, and Y. Fu. Efficient mining of association rules in distributed databases. *IEEE Transactions on Knowledge and Data Engineering*, 8(6):911–922, 1996.

[7] M. Delgado, N. Marín, D. Sánchez, and M.A. Vila. Fuzzy association rules: General model and applications. *IEEE Transactions on Fuzzy Systems*, 11(2):214–225, 2003.

[8] M. Delgado, M. Ros, and M.A. Vila. Correct behavior identification system in a tagged world. *Expert Systems with Applications*, 36(6):9899–9906, 2009.

[9] M. Delgado, M.D. Ruiz, and D. Sánchez. Studying interest measures for association rules through a logical model. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 18(1):87–106, 2010.

[10] M. Delgado, M.D. Ruiz, and D. Sánchez. New approaches for discovering exception and anomalous rules. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 19(2):361–399, 2011.

[11] M. Delgado, M.D. Ruiz, D. Sánchez, and J.M. Serrano. A formal model for mining fuzzy rules using the RL representation theory. *Information Sciences*, 181(23):5194–5213, 2011.

[12] I. Diaz-Valenzuela, M.J. Martín-Bautista, and M.A. Vila. Using a semisupervised fuzzy clustering process for identity identification in digital libraries. In *IFSA World Congress and NAFIPS Annual Meeting (IFSA/NAFIPS), 2013 Joint*, pages 831–836, 2013.

[13] D. Dudek. Measures for comparing association rule sets. In *Int. Conference on Artificial Intelligence and Soft Computing, ICAISC 2010, Part I, LNAI 6113*, pages 315–322, 2010.

[14] M. M. A. Ellatif. Association rules technique to diagnosis financial performance for ksa stock market companies. *Working Paper Series, Faculty of Computers and Information. Mansoura University of Egypt*, 2005.

[15] W. Jian and L.Xing ming. An efficient association rule mining algorithm in distributed databases. In *First Int. Workshop on Knowledge Discovery and Data Mining, WKDD 2008*, pages 108–113, 2008.

[16] R. Jin, G. Yang, and G. Agrawal. Shared memory parallelization of data mining algorithms: Techniques, programming interface and performance. *IEEE Transactions on Knowledge and Data Engeneering*, 17(1):71–89, 2005.

[17] B. Liu, W. Hsu, and S. Chen. Using general impressions to analyze discovered classification rules. In *KDD*, pages 31–36, 1997.

[18] B. Liu, W. Hsu, L. Mun, and H. Lee. Finding interesting patterns using user expectations. *IEEE Transactions on Knowledge and Data Engeneering*, 11(6):817–832, 1999.

[19] J. Llinas and D.L. Hall. *Handbook of Multisensor data fusion*, chapter Multisensor data fusion, pages 1–14. CRC Press, 2009.

[20] E. Louie and T.Y. Lin. Finding association rules using fast bit computation: Machine-oriented modeling. *LNAI*, 1932:486–494, 2000.

[21] H. Lu, J. Han, and L. Feng. Stock movement prediction and n-dimensional inter-transaction association rules. In *ACM SIGMOD Workshop on Research Issues on Data Mining and Knowledge Discovery*, pages 121–127, Seattle, 1998.

[22] M. Riondato, J.A. DeBrabant, R. Fonseca, and E. Upfal. Parma: a parallel randomized algorithm for approximate association rules mining in mapreduce. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM'12)*, pages 85–94, 2012.

[23] M. Ros, M.P. Cuéllar, M. Delgado, and M.A. Vila. Online recognition of human activities and adaptation to habit changes by means of learning automata and fuzzy temporal windows. *Information Sciences*, 220:86–101, 2013.

[24] M.D. Ruiz, M.J. Martin-Bautista, D. Sánchez, M.A. Vila, and M. Delgado. Anomaly detection using fuzzy association rules. *Int. J. Electronic Security and Digital Forensics*, In press, 2014.

[25] D. Sánchez, M.A. Vila, L. Cerda, and J.M. Serrano. Association rules applied to credit card fraud detection. *Expert Systems with Applications*, 36:3630–3640, 2009.

[26] N. Saxena, R. Arora, R. Sikarwar, and A. Gupta. An efficient approach of association rule mining on distributed database algorithm. *Int. Journal of Information and Computation Technology*, 3(4):225–234, 2013.

[27] E. Shortliffe and B. Buchanan. A model of inexact reasoning in medicine. *Mathematical Biosciences*, 23:351–379, 1975.

[28] S.K. Tanbeer, C.F. Ahmed, B. Jeong, and Y. Lee. Efficient mining of association rules from wireless sensor networks. In *11th Int. Conference on Advanced Communication Technology, ICACT 2009*, pages 719–724, 2009.