Discovering market basket patterns using hierarchical association rules

Marijana Zekić-Sušac[1],* and Adela Has[1]

[1] *Faculty of Economics in Osijek, University of Josip Juraj Strossmayer in Osijek,*
Trg Lj. Gaja 7, 31000 Osijek, Croatia
E-mail: ⟨{*marijana, adela.has*}*@efos.hr*⟩

Abstract. Association rules are a data mining method for discovering patterns of frequent item sets, such as products in a store that are frequently purchased at the same time by a customer (market basket analysis). A number of interestingness measures for association rules have been developed to date, but research has shown that there a dominant measure does not exist. Authors have mostly used objective measures, whereas subjective measures have rarely been investigated. This paper aims to combine objective measures such as support, confidence and lift with a subjective approach based on human expert selection in order to extract interesting rules from a real dataset collected from a large Croatian retail chain. Hierarchical association rules were used to enhance the efficiency of the extraction rule. The results show that rules that are more interesting were extracted using the hierarchical method, and that a hybrid approach of combining objective and subjective measures succeeds in extracting certain unexpected and actionable rules. The research can be useful for retail and marketing managers in planning marketing strategies, as well as for researchers investigating this field.

Key words: association rules, data mining, market basket analysis

1. Introduction

The aim of data mining is to discover relationships hidden in large amounts of data sourced primarily from transactional databases [8]. Association rules are used in data mining to discover patterns of frequent sets of items that appear together. Besides association rules, data mining includes machine learning methods such as neural networks, decision trees, support vector machines, and also various statistical methods that assist in multivariate analysis, prediction, classification, and pattern recognition. One of the problems to which association rules can be applied is the market basket problem, which assumes a number of products that a customer can purchase, either as a single transaction or over time as a sequence of transactions. These products can be items displayed in a supermarket, insurance packages that customers might be willing to purchase,

---

* Corresponding author.

or other offered products and services. Association rules have shown to be useful in discovering which products are frequently purchased together [7], [11], and [12], and which in turn can be used to define efficient marketing strategies for advertising and promotional activities.

Most of the research in the area of association rules uses the apriori algorithm for discovering relationships among items, while other algorithms, such as the tree-building technique, have not been investigated adequately. In this paper, the tree-building technique is used and association rules with and without hierarchical item grouping are compared and analyzed.

2. Previous research

Research on association rules focuses primarily on finding criteria and measures for extracting interesting rules that can help managers in defining marketing strategies. Most rule extraction criteria are based on minimum support and confidence measures, which are probability-based objective measures that evaluate the generality and reliability of association rules [8]. Tan et al. [14] compare twenty-one measures of association rule interestingness based on eight properties, and conclude that none of the measures satisfy all the properties.

Lee et al. [7] base their approach on defining different minimum support for different item sets. They use a measure of maximum constraint, and then propose an algorithm based on the apriori approach to find large item sets and association rules subject to this constraint. However, they suggest considering the requirements of mining problems when using the algorithm. Zhou and Yau [15] suggest a special algorithm for mining infrequent items. Liao et al. [8] suggest an "extracted probability" measure that modifies the confidence of rules such that it computes extracted probability based on linear, quadratic and logarithmic models, and helps to extract some unapparent rules with low support, but with a high confidence. When comparing their measure with a standard confidence-support measure, they discovered certain advantages of their approach in that it generated fewer rules that businesses would find interesting. Besides the above objective measures, association rules can be evaluated using subjective measures, such as unexpectedness and actionability, as suggested by Silberschatz and Tuzhilin [13]. However, subjective measures are user-dependent and are not well defined yet. Abdullah et al. [1] uses association rules in education and proposes a new measurement called Critical Relative Support (CRS) to mine critical least association rules from educational data. Their research has shown that the measure reveals significant rules on student examination results and reduces uninterested association rules by up to 98%. Panchon Alvarez et al. (2012) suggests an evolutionary-based approach for dealing with numeric and discrete attribute values in association rules. They comprise quantitative and categorical attributes without the need for an apriori discretization of the domain of numeric attributes.

The majority of the above research was based on the apriori algorithm which focuses mostly on positive association rules such as "if a customer buys

item *X*, he/she will also buy item *Y*" which can sometimes be misleading. Chen et al. [4] suggest using properties of propositional logic and propose an algorithm for mining highly coherent rules without a minimum support threshold. Their algorithm produced association rules more reliable for managers. Lin et al. [9] suggested an improved frequent pattern growth technique for mining association rules, providing some advantages over apriori algorithm, such as faster rule generation and a scalable approach.

Previous research has shown that association rules have relevant application in many areas, that there is a variety of objective and subjective measures proposed and tested with different datasets, while most authors use the apriori algorithm modified to some degree as a basis for producing rules. In this paper, we also use the frequent pattern tree technique, testing it on a real dataset and observing certain benefits and limitations in situations in which original and hierarchically grouped items exist.

3. Methodology

Association rules as a method of data mining was first introduced by Agrawal et al. [2] who proposed the method for market basket analysis. Association rules can be very useful in discovering unknown data relationships, providing the basis for decision-making in marketing, retail, education, and other areas.

According to Liu et al. [10], association rules are described in a specific manner. If $I = \{i_1, i_2, ..., i_m\}$ is a set of items, and D is a set of transactions (a dataset), where each transaction d is a set of items such that $d \subseteq I$, an association rule is an implication of the form $X \rightarrow Y$, where $X \subset I$, $Y \subset I$, and $X \cap Y = \emptyset$. There are two basic measures of the interestingness of an association rule: a support *s* and a confidence *c*. The rule has the support *s* in *D* if *s* percent of transactions in *D* contains $X \cup Y$. The rule $X \rightarrow Y$ holds true in the transaction set *D* with the confidence *c* if *c* percent of transactions in *D* that supports *X* also supports *Y*. If *k* is the number of transactions containing *X*, *l* is the number of transactions containing *Y*, *m* is the number of transactions containing $X \cup Y$, and *n* is the total number of transactions in D, then *s* can be expressed as a probability by (1):

$$s(X \rightarrow Y) = \frac{m}{n}, \quad s(Y \rightarrow X) = s(X \rightarrow Y) \tag{1}$$

and *c* can be expressed as a conditional probability by (2):

$$c(X \rightarrow Y) = \frac{m}{k}, \quad c(Y \rightarrow X) = \frac{m}{l} \tag{2}$$

According to [8] for an association rule, $X \rightarrow Y$, *s(X)* or *s(X∪Y)* is used to represent the generality of the rule, and *c(X→Y)* is used to represent the reliability of the rule. Although in general, a rule with high generality and reliability is considered interesting, many authors emphasize that rules with low generality can sometimes have a very high reliability and therefore, can be very interesting. Another frequently used measure of interestingness is the lift value

or the interest factor. It is used in data mining to measure deviation from statistical independence [14], and can be defined as:

$$lift\,(X \rightarrow Y) = \frac{c(X \rightarrow Y)}{s(Y)} \tag{3}$$

The lift $X \rightarrow Y$ measures an increase in the possibility of purchasing Y after previously purchasing X, rather than purchasing Y without X [6], i.e. to what extent does the likelihood of purchasing Y increase after purchasing X. A lift value of 1 indicates that the items exist independently in the database, as expected. When the value is greater than one, it indicates that the items are associated [5].

Given a set of transactions D (a dataset), the problem of mining association rules involves discovering all the relevant rules. Any item can appear on the left-hand side (called body or antecedent) or the right-hand side of an association rule (called head or consequent).

The standard algorithm used in association rules is the apriori algorithm introduced by Agrawal and Srikant [3]. This algorithm scans the entire dataset, evaluates all possible rules, and then retains only those rules that have a support s and confidence c greater than predefined minimum values as relevant ones. The apriori algorithm for establishing association rules from transactional data consists of the following two steps: (1) finding all frequent itemsets from a transactional database that satisfy the minimum support s, and (2) generating association rules from the frequent itemsets that satisfy the minimum confidence c [1]. In the first stage of the algorithm, the support s is computed for each itemset of size k ($k=1,2,...m$, where m is the maximum number of simultaneously items purchased in an itemset) in order to find the most frequent itemsets. This is done in one pass over the data, and when all frequent itemsets are found, the algorithm moves to the second stage - generating rules that satisfy the minimum confidence c. The procedure of mining frequent itemsets is multi-leveled, where the level k denotes the number of items in a set. Therefore, at level k, all potentially frequent k-itemsets generated from frequent itemsets at level $k-1$ are found.

The apriori algorithm has certain limitations, such as being a time consuming procedure that examines all combinations or the possible scenario of retaining only trivial rules but discarding some of the interesting rules. Association rules in our experiments are generated by an improved algorithm called the tree-building technique, which compresses a large database into a compact, Frequent-Pattern tree (FP-tree) structure [9]. The advantage of this algorithm is its speed, i.e. it scans the entire database only once. A divide-and-conquer approach is used, by first computing the frequent items and characterizing them into a frequent-pattern tree. The FP-tree is a compressed database, on which the association rule mining is performed. In addition, the algorithm does not require generation of the candidate itemset. Hence, it is more efficient than the apriori algorithm [9]. The disadvantage of the FP-tree algorithm is the recursive generation of a large number of conditional FP trees as part of the mining procedure.

Association rules are evaluated using different measures, and are usually divided into objective and subjective measures [13]. The basic objective measures are the support *s* and the confidence *c*. According to [13], the rule (i.e. pattern) is interesting if it satisfies two criteria: (a) unexpectedness (if it is surprising to the user), and/or (b) actionability (if the user can do something with it to his or her advantage). These are subjective measures but have some advantage over objective measures. Silberschatz and Tuzhilin [13] proposed a method for measuring unexpectedness, which is based on a user's belief system. According to these authors, a pattern (i.e. a rule) is more interesting if it is less expected i.e. in greater contradiction to a user's previous beliefs. Their approach is based on the Dempster-Shafer theory, according to which a belief function is assigned to beliefs and, in general, this does not satisfy the axioms of probability theory [13].

In this paper, we combined an objective and subjective approach, such that three measures of rule interestingness were used: heuristical unexpectedness and heuristical actionability as subjective measures which had been suggested in [13], and also the confidence *c* as the objective measure, where we consider a rule with a minimal confidence of 51% as being potentially interesting (as it is more likely than a random choice). Although Silberschatz and Tuzhilin [13] suggest a computational method for measuring unexpectedness based on the frequency of items (i.e. support), in this paper we use a sale's manager's subjective heuristic estimation, where the expert sales manager estimates unexpectedness and actionability as binary values such that:

$$unexpectedness = \begin{cases} 0, & \textit{if the rule was not expected based on heuristics} \\ 1, & \textit{if the rule was expected based on heuristics} \end{cases} \tag{4}$$

$$actionability = \begin{cases} 0, & \textit{if the rule cannot be used to create a useful action} \\ 1, & \textit{if the rule can be used to create a useful action} \end{cases} \tag{5}$$

where the term "heuristics" refers to the assumed heuristic knowledge of the expert sales manager and team, the term "user" refers to a sales or marketing department belonging to the observed company, and the term "useful action" implies any marketing activity, such as discount, paired advertising, performed in order to increase the sale of an itemset that constructs a certain rule. In our experiments, the support *s* was selected in a cross-validation procedure such that different values of *s* were tested in order to find the most suitable value for each of the tested models. Minimum *s*, used to produce the rules, ranged from 3% to 50%. The size of the itemsets varied from 1 to 4, but only itemsets with a size of 2 were found to be significant for the analysis.

4. Data

This research utilized data from a transactional database of a Croatian retail store chain. The dataset contained 14012 transactions, carried out in the period

from 3 - 21 January 2011. Each transaction included the following variables: purchase date, account number, code and article name. In all, the data included 1230 different items, and they appeared on 6363 different accounts. Each of these accounts represents a single customer. Due to the large number of different dataset items, transactions containing very rare items were excluded, leaving 7006 transactions with 278 different items appearing on 3158 different accounts. The acquired data was processed using two strategies, the generation of (1) association rules at the first-level data hierarchy and (2) association rules at the second-level data hierarchy. To provide item groups for the second strategy used in generating association rules, item were grouped together that represented the same type of product, but from either a different manufacturer, different packaging, brand, net weight, or volume. The grouping procedure resulted in 38 large groups of items. Given that the results of the first strategy indicated that the majority of association rules consisted of individual milk products as association rule bodies and heads (because of the high support of these items in the transactions), we decided to perform a second-level hierarchical grouping by additionally adding individual milk products into a single group of items named "milk products". The aim was to create a higher probability of other, less supported items to appear in the extracted association rules.

5. Results

The results are described separately for each tested strategy, and interesting association rules are extracted and discussed.

5.1. Generating association rules from the first-level hierarchical grouping

In this strategy, the first-level hierarchical grouping of items is used prior to generating the association rules. After cross-validating with different values of the support s and the confidence c (values ranged from 3% to 50%), the following parameters were selected as best-suited to the various generated rules: the minimal support coefficient $s = 10\%$, the minimal confidence coefficient $c = 10\%$. These low s and c values were used due to the large number of items in the dataset. The procedure showed that higher values of s produced a small number of extracted rules that contained only a few of the most frequent items. The minimal value of $s=10\%$ produced diverse rules, with some having a very high confidence c. In all, 36 association rules were extracted as significant, where the first 8 rules with a confidence c of at least 50% were sorted according to the confidence value and presented in Table 1. The table also contains two subjective measures of rule interestingness, expressed as unexpectedness and actionability, estimated in conjunction with the sales manager of the observed

store. The symbolic representation of rules $X \rightarrow Y$ is used in all tables, and can be translated into a rule in the form „*if X then Y*".

Rule ID	Rule $X \rightarrow Y$	Objective measures		Subjective measures	
		Support (%)	Confidence (%)	Heuristical unexpectedness	Heuristical actionability
1	TOWEL → BAG	11.11111	80.00000	0	1
2	TUNA → DOUGH*	9.25926	71.42857	1	1
3	TOWEL → TP	9.25926	66.66667	0	1
4	CREAM → CHEESE	15.74074	60.71429	0	1
5	TP → TOWEL	9.25926	58.82353	0	1
6	MILK → YOGURT	20.37037	51.16279	0	1
7	YOGURT → MILK	20.37037	51.16279	0	1
8	CREAM → YOGURT*	12.96296	50.00000	1	1

*Association rules that satisfy all three suggested criteria of interestingness (confidence min. 50%, unexpectedness, and actionability)

Table 1: *Association rules generated by the 1ˢᵗ level grouping strategy (min. s=10%)*

As is evident from Table 1, only two rules TUNA → DOUGH (if TUNA then DOUGH), and CREAM → YOGURT (if CREAM then YOGURT) satisfy all three criteria used in this research (confidence >= 50%, positive unexpectedness, and actionability). Therefore, these rules are considered highly interesting ones. The rule TUNA → DOUGH is the most interesting one, since it has a very high confidence (71.43%), revealing a high probability that a customer will buys dough if buying tuna, and is positively valued by the expert in terms of unexpectedness and actionability. The support of this rule is not high (9.26%), revealing that item pairs with a low frequency can have a high probability and therefore be of interest for decision makers. Furthermore, sales management at the observed store was surprised by this rule. The rule CREAM → YOGURT is also an interesting one due to its high confidence (50%) and a support value of 12.96%. It was deemed by the expert as both unexpected and actionable, and therefore, exceptionally useful for carrying out marketing activities, such as discounts. As is evident in Table 1, the highest confidence (80%) with a support of 11.11% is obtained for the rule TOWEL → BAG (if TOWEL then BAG), which was not surprising, because "bag" refers to a sort of plastic bag sold in some stores to carry purchased items, and the expectation was that it would be bought along with any purchased item at the store. Interesting enough, various milk products that were the most frequently purchased items in the observed stores are not contained in the first three rules attributed with a highest confidence. To find the reason as to Table 1 contains a large number of rules estimated to be not interesting to a user, more attention is directed to the support for each individual dataset item. The results showed that the individual support s is greater than 25% for only 7 items. Table 2 presents these items and their associated support coefficients.

ID	Frequent item	Support (%)
1	Yogurt	39.81481
2	Milk	39.81481
3	Cheese	33.33333
4	Bread	30.55556
5	Cigarettes	30.55556
6	Bag	28.70370
7	Cream	25.92593

Table 2: *Support coefficient s for individual items that have s ¿ 25%*

The items presented in Table 2 have the highest frequency in the dataset. Generating a new set of association rules based only on data relating to these items is recommended. A total of 25 association rules were generated from these selected transactions, where only one association rule was selected as unexpected and actionable: CREAM → YOGURT (if CREAM then YOGURT) with a support of 12.96% and confidence of 50.00%. The graphical representation of rules below is based on a selection of items with a minimal support of 25% for rule generation. The size of a node is the relative support of each rule and its color intensity illustrates relative confidence (light-colored nodes have a smaller relative confidence, while dark-colored nodes have a higher relative confidence). As can be seen from Figure 1, the highest confidence (60.71%) is obtained for the rule CREAM → CHEESE (if CREAM then CHEESE) with the support of 15.74%. It becomes obvious that milk products dominate in the association rules, due to its high support of individual items.
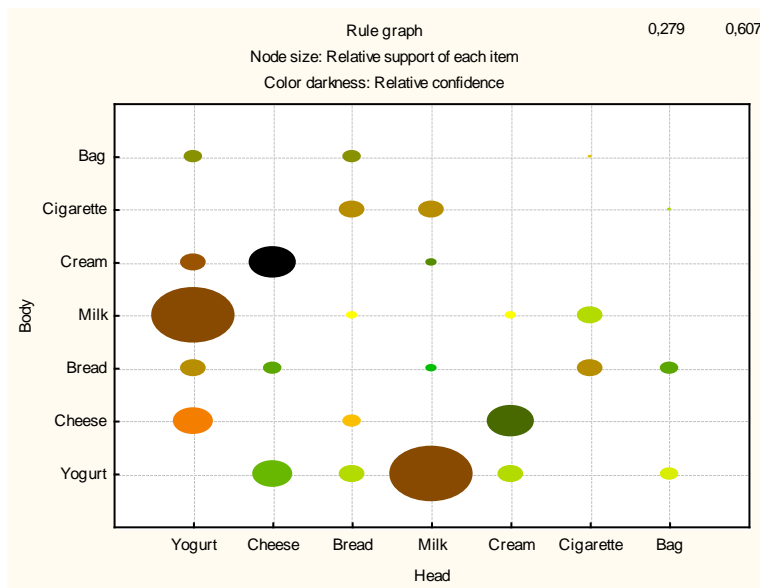

Figure 1: *Graph of association rules at the 1st level of grouping (items with min s = 25%)*

The web graph of the above rules (Figure 2) presents the relative support for items, and relative joint support for items, and the lift value of an itemset. The size of a node represents relative support for each item (a small node represents smaller relative individual support for an item, while a larger node represents a greater relative individual support). The thickness of a line which connects two items shows relative joint support for the rule (the thicker the line, the greater the joint support); whereas color intensity of a line indicates the lift value of a rule (a light colored line has a smaller relative lift, while a dark colored line has a greater relative lift). Lift value is symmetric, which means that the lift value $X \rightarrow Y$ is equivalent to the lift value $Y \rightarrow X$ that also applies to support for rules, such that $s(X \rightarrow Y)$ has the same value as $s(Y \rightarrow X)$.
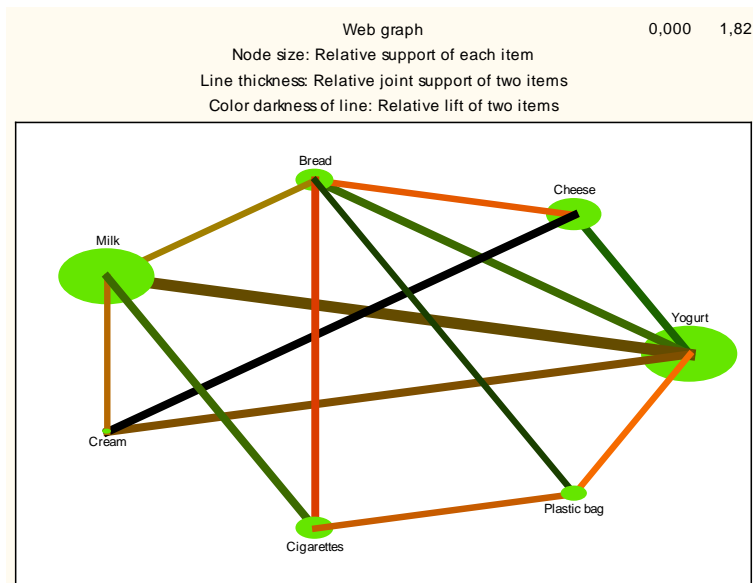

Figure 2: *Web graph at the 1st hierarchical level (items with min s= 25%)*

From Figure 2 it is obvious that milk and yogurt have the highest relative support and the highest relative joint support, whereas a dark color for their connection shows a high lift for the rule. The highest lift is obtained for the cream-cheese pair, but these items exhibit a low individual support, especially in regards to the cream.

Since the rules that contain milk products are mostly expected and rarely surprising, in order to find more interesting associations it would be useful to group the individual milk products (except the milk itself) together and observe the generated rules.

5.2. Generating association rules on the second-level hierarchical grouping

When the second-level grouping of items is performed according to the second strategy described above, a total of 30 rules was generated with the minimum support for individual item equivalent to $s = 10\%$. However, 15 rules had a confidence of c >= 50%. Table 3 shows the rules, the support, confidence and subjective measures of interestingness for the respective grouping stage. For reasons of clarity, only rules with a confidence >= 50% are shown in the table.

Rule ID	Rule $X \rightarrow Y$	Objective measures		Subjective measures	
		Rule Support (%)	Rule Confidence (%)	Heuristical unexpectedness	Heuristical actionability
1	MILK → MILK_PRODUCTS	33.82353	88.46154	0	1
2	JUICE → MILK_PRODUCTS*	20.58824	82.35294	1	1
3	FRUIT → MILK_PRODUCTS*	35.29412	80.00000	1	1
4	DETERGENT → MILK_PRODUCTS*	20.58824	73.68421	1	1
5	PLASTIC_BAG → MILK_PRODUCTS	27.94118	70.37037	0	1
6	DOUGH → MILK_PRODUCTS	19.11765	68.42105	0	1
7	DOUGH → VEGETABLES*	19.11765	68.42105	1	1
8	VEGETABLES → MILK_PRODUCTS*	32.35294	64.70588	1	1
9	BREAD → MILK_PRODUCTS	23.52941	64.00000	0	1
10	CIGARETTES → MILK_PRODUCTS	20.58824	58.33333	1	0
11	CIGARETTES → BREAD	20.58824	58.33333	1	0
12	FRUIT → VEGETABLES	25.00000	56.66667	0	1
13	BREAD → CIGARETTES	20.58824	56.00000	0	1
14	BREAD → PLASTIC_BAG	19.11765	52.00000	0	1
15	VEGETABLES → FRUIT	25.00000	50.00000	0	1

Table 3: *Association rules generated by the 2nd level grouping strategy (min. $s=10\%$)*

As can be seen from Table 3, 15 rules had a confidence greater than 50%, which is a substantial progress in comparison to the results in previous tables. As expectedly, the pair milk and milk products have a very high confidence (88.46%) and a high support (33.82%). By grouping milk products together, other items also appeared just as frequently and with a confident that not been expected. For example, the rule JUICE → MILK_PRODUCTS (if JUICE then MILK_PRODUCTS) has a confidence of 82.35% and support of 20.58%, which

the user considers surprising and actionable. Similar values are obtained for the rule (if FRUIT then MILK_PRODUCTS), with a confidence of 80% and support of 35.29%. Besides these two, interesting rules are also: DETERGENT → MILK_PRODUCTS (if DETERGENT then MILK_PRODUCTS), DOUGH → VEGETABLES (if DOUGH then VEGETABLES), and VEGETABLES → MILK_PRODUCTS (if VEGETABLES then MILK_PRODUCTS). These all have high confidences (above 50%), are unexpected and actionable. In order to graphically represent the most important rules, items with an individual support greater than 25% were one again selected and Figure 3 shows these rules in the form of a web graph.
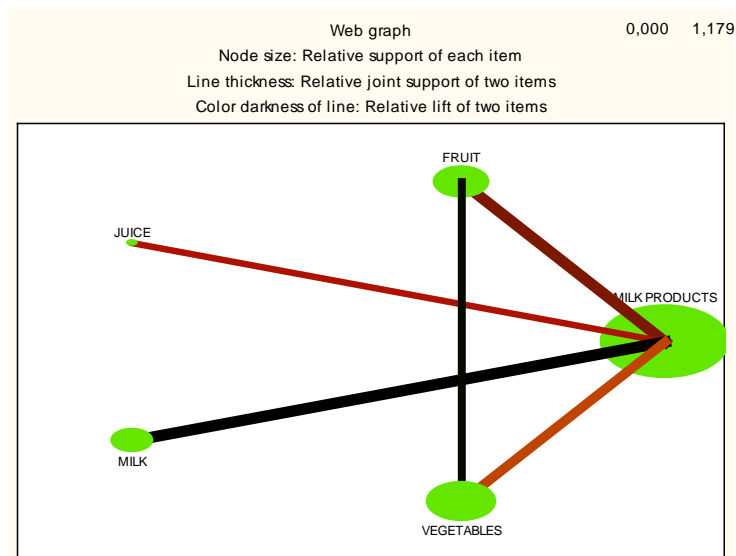

Figure 3: *Web graph at the 2nd hierarchical level (items with min s= 25%)*

The web graph shown in Figure 3 has a very interesting form, where milk products dominate in terms of node size and show a high purchase frequency for these items, while milk products are frequently purchased together with all other items (milk, fruit, vegetables and juice). The purchase with the highest relative joint support is the fruit - milk products pair and the milk- milk products pair. The greatest lift value is obtained for the milk – milk products pair, and for the fruit – vegetables pair (black colored line), whereas juice and milk are frequently bought only with milk products, not with fruit or vegetables. The smallest relative support is visible for the juice – milk products pair.

As can be seen from Figure 3, milk products, vegetables, fruit, milk, and juice are the most frequent items on the second-level item grouping. Their individual support is greater than 25%, but when observed in pairs, the highest relative support (the largest sized nodes) is obtained for the fruit – milk products pair, which is very surprising and actionable.

The above results indicate that significantly more distinctive rules are created when the second level of grouping is performed, i.e. when the most frequent items (milk products) are grouped together, thus allowing other items to appear in the association rules more clearly. At the first level of item grouping, the quantity of extracted rules was large, but the quality of the rules was not high, give the many rules associated with very small support in the dataset. Therefore, a very small number of rules were selected as being interesting at the first level of item grouping. At a higher grouping level, the support of the item sets increased and a larger number of rules had confidences $>= 50\%$, and were also evaluated as being interesting.

The evident reason for efficiency of higher-level item grouping could be due to the specific characteristics of the milk products group, their important role in Croatian nutrition habits, but could also be generalized for any group of such products that are frequently consumed in a country. Besides the fact that higher-level grouping degrades analytical precision, the results of this research stress the necessary role of item grouping in analyzing the market basket, given that extracting clearer and more confident rules for use by decision makers is important.

6. Discussion and conclusion

The paper investigates the efficiency of association rules based on a technique to construct a frequent-pattern tree using a real dataset from a Croatian retail store. Two strategies for generating association rules were used: first-level and second-level hierarchical grouping of items. Choosing interesting rules is a challenging task when analyzing the market basket. Previous research has suggested using either objective or subjective approaches. In this paper, the objective and subjective approach are combined using support and confidence coefficients in conjunction with a subjective assessment of rule for unexpectedness and actionability based on the heuristics of a human expert. The results have shown that first-level grouping generated a large number of rules leading to very few interesting rules, whereas second-level grouping generated rules that were more supported, more significant and more interesting in terms of heuristical unexpectedness and actionability. Although the results are limited to the observed dataset, the paper can assist researchers and practitioners involved in association rules and analysis of the market basket for the purpose of supporting marketing decisions on advertising and discounting paired items, with the final goal being increasing sales and profit.

References

[1] Abdullah, Z., Herawan, T., Ahmad, N. and Mat Deris, M. (2011). Mining significant association rules from educational data using critical relative support approach. Procedia - Social and Behavioral Sciences, 28, 97–101. doi:10.1016/j.sbspro.2011.11.020

[2] Agrawal, R., Imielinski, T. and Swami, A. (1993). Database mining: A performance perspective. IEEE Transactions on Knowledge and Data Engineering, 5, 914–925. doi:10.1109/69.250074.

[3] Agrawal, R. and Srikant, R. (1994). Fast algorithms for mining association rules in large databases, in J.B. Bocca, M. Jarke, C. Zaniolo (Eds.): Proceedings of 20th International Conference on Very Large Data Bases VLDB '94, September 12-15, 1994, Santiago de Chile, Chile. Morgan Kaufmann, 487–499.

[4] Chen, C.-H., Lan, G.-C., Hong, T.-P. and Lin, Y.-K. (2013). Mining high coherent association rules with consideration of support measure. Expert Systems with Applications, `http://dx.doi.org/10.1016/j.eswa.2013.06.002` [Accessed on July, 18, 2013].

[5] Hahsler, M. and Hornik, K. (2007). New probabilistic interest measures for association rules. Intelligent Data Analysis, 11(5), 437–455.

[6] Kwang-Il A., (2012). Effective product assignment based on association rule mining in retail. Expert Systems with Applications, 39(16), 12551–12556. doi:10.1016/j.eswa.2012.04.086.

[7] Lee, Y.-C., Hong, T.-P. and Lin, W.-Y., 2005. Mining association rules with multiple minimum supports using maximum constraints. International Journal of Approximate Reasoning, 40, 44–54. doi:10.1016/j.ijar.2004.11.006.

[8] Liao, C.-W., Perng, Y.-H. and Chiang, T.-L., (2009). Discovery of unapparent association rules based on extracted probability. Decision Support Systems, 47, 354–363. doi:10.1016/j.dss.2009.04.006.

[9] Lin, K.-C. , Liao, I.-E. and Chen, Z.-S. (2011). An improved frequent pattern growth method for mining association rules. Expert Systems with Applications, 38, 5154–5161. doi:10.1016/j.eswa.2010.10.047.

[10] Liu, B., Ma, Y. and Wong, C.-K. (2001). Classification using association rules: Weaknesses and enhancements, in R. L. Grossman et al. (eds), Data Mining for Scientific and Engineering Applications. Springer, `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.8.7943&rep=rep1&type=pdf` [Accessed on July, 19, 2013].

[11] Pachón Álvarez, V. and Mata Vázquez, J. (2012). An evolutionary algorithm to discover quantitative association rules from huge databases without the need for an a priori discretization. Expert Systems with Applications, 39, 585–593. doi:10.1016/j.eswa.2011.07.049.

[12] Pasquier, N., Bastide, Y., Taouil R. and Lakhal, L. (1999). Efficient mining of association rules using closed itemset lattices. Information System, 24, 25–46.

[13] Silberschatz, A. and Tuzhilin, A. (1995). On subjective measures of interestingness in knowledge discovery. Proceedings from the First International Conference on Knowledge Discovery and Data mining (KDD-95), U. M. Fayyad, R. Uthurusamy (Eds.): Montreal, Canada, August 20-21, 275–281.

[14] Tan, P.N., Kumar V. and Srivastava J. (2004). Selecting the right interestingness measure for association patterns. Information Systems, 29, 293–313. doi:10.1016/s0306-4379(03)00072-3.

[15] Zhou, L. and Yau, S. (2007). Efficient association rule mining among both frequent and infrequent items. Computers and Mathematics with Applications, 54, 737–749. doi:10.1016/j.camwa.2007.02.010.