Importance of Association Rules in Data Mining: A Review

¹L. Arockiam, ¹S.S. Baskar and ²L. Jeyasimman

¹Department of Computer Science, St. Joseph's College, 620002 Trichirappalli, India

²Department of Computer Applications, JJ College of Engineering and Technology,
620009 Trichirappalli, India

Abstract: Discovering association rules are a core topic of data mining. This survey study presents an overview to some of the previous researches on this subject, evaluating the current status of the field and revealing the future trends in this area. The idea behind association rules is presented and comparison of different algorithms is provided as part of the evaluation.

Key words: Association rules, Apriori, AprioriTid, AprioriHybrid, evaluation, subject

INTRODUCTION

Discovering association rules is at the heart of data mining. It detects hidden linkages of unrelated data. These linkages can provide rules. Those that exceed a certain threshold are deemed to be interesting. These rules allow actions to be taken based upon data pattern. They can also help making and justifying decisions.

The best example for mining association rules is the market-basket problem. As Data Mining Lecture Notes describes, the market-basket problem is about finding out what items people buy together without knowing the person so that marketers can position items accordingly in the store to generate higher volumes of sales and making other kinds of sales decisions. Some of the rules discovered may be trivial e.g., people who buy bread tend to buy butter. It is the extraordinary rules that are interesting e.g., people who buy diapers also buy beers. It is the ability to discover the interesting rules that makes association rules valuable and contributes to knowledge discovery.

The problem of discovering association rules can be generalized into two steps:

- Fining large itemsets
- Generating rules from these itemsets

Previous researches are mostly along these lines and have grown into various dimensions. This study aims at reviewing some of the earlier researches in the field and evaluating the current status and future trends. Background theories behind discovering association rules are being discussed in this study. Algorithms with respect to association rule are being discussed. The last section concludes the contribution of the study.

THEORIES IN ASSOCIATION RULES

Association rules are defined as statements of the form $\{X_1, X_2, \dots, X_n\} \rightarrow Y$ which means that Y may present in the transaction if $X_1, X_2, \dots, X_n$ are all in the transaction. The association rule will be useful for discovering interesting relationships hidden in the large data base. The rule is only probable, not identical. There can be a set of items, not just a single item. The probability of finding Y in a transaction with all $X_1, X_2, \dots, X_n$ is called confidence. The threshold that a rule holds in all transactions is called support. The level of confidence that a rule must exceed is called interestingness (Chen *et al.*, 1996).

There are different types of association rules like Boolean association rule and generalized association rule. The simplest form is the type that only shows valid or invalid association. This Boolean nature of the rule gives the name Boolean Association Rules. In the market-basket example, people who buy skim milk also buy low fat oil is a Boolean association rule.

Rules that are aggregating several association rules together are called Multilevel or generalized association rules (Chen *et al.*, 1996). These rules usually involve a hierarchy and mining is done at a higher concept level. For example, people who buy milk also buy bread. In this example, milk and bread each contains a hierarchy of different types and brands but mining at the lowest level may not produce very interesting rules.

A more complicated type of rules is the quantitative association rules. This type of rules mines over quantitative (price) or categorical (gender) attributes and is denoted by Srikant and Agrawal (1996) $\{ \langle \text{attribute: value} \rangle, \langle \text{attribute: value} \rangle, \dots, \langle \text{attribute: value} \rangle \} \rightarrow \langle \text{attribute: value} \rangle$.

For example, people whose age is between 30 and 35 with income >75000 per year buy cars >20000. However, the above types do not address the fact that transactions are temporal in nature. For example, mining before a product is introduced to or after a product is discontinued from the market will both adversely affect the support threshold. In view of this (Ale and Rossi, 2000) introduced the concept of an attribute's lifetime into the mining algorithm of temporal association rules.

In spite of the various kinds of rules, the algorithm to discover association rules can generally be broken down into two steps:

- Find all large (frequent) itemsets-A large itemset is a set of items that exceeds the minimum support
- Generate rules from the large itemsets

Since, its introduction by Srikant and Agrawal (1996) the Apriori algorithm has been considered the most used algorithm for finding frequent item sets. Many improvements (Zaki *et al.*, 1996; Han *et al.*, 2000a) are made in the Apriori algorithms such as speed up and scale up for finding frequent item sets. Improvements in Apriori are made by addressing its erroneous belief of generating too many candidate itemsets. Few algorithms have been developed that are not based on Apriori (Han *et al.*, 2000b; Laks *et al.*, 2000; Park *et al.*, 1997) but it addresses the issues of speed of Apriori.

Generate rules from the large item sets are mostly characterized by confidence and interestingness. Some investigation on different way of generating rules (Mata *et al.*, 2002) and alternative measure to interestingness (Chan and Au, 1997; Natarajan and Shekar, 2005) are explored. There are also explorations about generating different types of rules (Srikant and Agrawal, 1996; Ale and Rossi, 2000).

FINDING LARGE ITEM SETS

Most of the earlier researches in Mining Association Rules were actually done in this topic. Lot of research was done to improve the performance and scalability of Apriori included using parallel computing. There were also studies to improve the speed of finding large itemsets with hash table, map and tree data structures.

In the past, AIS (Agrawal *et al.*, 1993) and SETM (Houtsma and Swami, 1993) were used for finding large item sets before Apriori. AIS and SETM generate candidate itemsets on-the-fly during the pass of database scan. Large itemsets from previous pass are checked if they are present in the current transaction. Hence, new

itemsets are formed by extending existing itemsets. These algorithms turn out to be ineffective because they generate and count too many candidate item sets that turn to be small (infrequent) (Srikant and Agrawal, 1996).

To overcome the problem, Apriori, AprioriTid and AprioriHybrid were proposed by Srikant and Agrawal (1996). Apriori and AprioriTid generate itemsets by using only the large itemsets found in the previous pass without considering the transactions. AprioriTid improves Apriori by only using the database at the first pass. Counting in subsequent passes is done using encodings created in the first pass which is much smaller than the database. This leads to a dramatic performance improvement of three times faster than AIS and four times faster than SETM in one of their experiments by Srikant and Agrawal (1996). A further improvement, called AprioriHybrid can be achieved when Apriori is used in the initial passes and switches to AprioriTid in the later passes if the candidate k-itemset is expected to fit into the main memory.

The problem with Apriori is that it generates too many 2-itemsets that are not frequent. A Direct Hashing and Pruning (DHP) algorithm (Park *et al.*, 1997) that reduced the size of candidate set by filtering any k-itemset out of the hash table if the hash entry does not have minimum support. This powerful filtering capability allows DHP to complete execution when Apriori is still at its second pass (Park *et al.*, 1997).

Algorithm Apriori:

```

L1 = {large 1-itemsets};
FOR (k = 2; Lk-1 != 0; k++) DO BEGIN
    Ck = apriori-gen (Lk-1);
    FORALL transactions t in D DO BEGIN
        Ct = subset (Ck, t);
        FORALL candidates c in Ct DO
            c.count++;
    END
    Lk = {c in Ck | c.count >= minsup}
END
Answer = Sum Lk;

```

```

FUNC apriori-gen (set Lk-1) BEGIN
    INSERT INTO Ck
    SELECT p.item1, p.item2, ..., p.itemk-1, q.itemk
    FROM Lk-1 p, Lk-1 q
    WHERE p.item1 = q.item1, ..., p.itemk-2 = q.itemk-2, p.itemk-1 < q.itemk;

    FORALL itemset c in Ck DO
        FORALL (k-1)-subsets s of c DO
            IF (s not in Lk-1) THEN
                DELETE c from Ck;
    END

```

Algorithm AprioriTid:

```

L1 = {large 1-itemsets};
C1' = database D;
FOR (k = 2; Lk-1 != 0; k++) DO BEGIN
    Ck = apriori-gen (Lk-1);
    Ck' = 0;

```

```

FORALL entries t in Ck-1 DO BEGIN
  Ck = {c in Ck | (c-c [k]) in t.set-of-itemssets ^ (c-c [k-1]) in t.set-of-itemssets};
  FORALL candidates c in Ck DO
    c.count++;
    IF (Ck != 0) THEN Ck' += <t.TID,Ck>;
  END
  Lk = {c in Ck | c.count >= minsup}
END
Answer = Sum Lk;

```

Some further efforts to improve Apriori algorithm utilize parallel algorithm (Zaki *et al.*, 1996) proposed 3 parallel algorithms based on Apriori to speed up mining of frequent itemsets. The Count Distribution (CD) algorithm minimizes communication at the expense of carrying out duplicate computations. The Data Distribution (DD) algorithm uses the main memory of the system to broadcast local data to all other nodes in the system. The Candidate Distribution algorithm is a load balancing algorithm that reduces synchronization between the processors and segments the database based upon different transaction patterns. These parallel algorithms were tested among each other and CD had the best performance against the Apriori algorithm. Its overhead is <7.5% when compared with Apriori by Zaki *et al.* (1996).

Scalability is another important research in data mining because data in databases are getting bigger day by day. Hence, algorithms should be able to scale up to handle large number of data.

Considering the above idea, the research was carried (Zaki *et al.*, 1996). According to Han *et al.* (2000a, b) tried to make DD and CD scalable by the Intelligent Data Distribution (IDD) algorithm and Hybrid Distribution (HD) algorithm, respectively. IDD addresses the issues of communication overhead and redundant computation by using aggregate memory to partition candidates and move data efficiently (Zaki *et al.*, 1996). HD improves over IDD by dynamically partitions the candidate set to maintain good load balance. Experiment results show that the response time of IDD is 4.4 times less than DD on a 32-processors System and HD is 9.5% better than CD on 128 processors (Han *et al.*, 2000b).

Another scalability study of data mining was done by introducing a light-weight data structure called Segment Support Map (SSM) that reduces the number of candidate itemsets needed for counting (Laks *et al.*, 2000). SSM contains the support count for the 1-itemset. The individual support counts are added together as the upper bound for k-itemsets. Applying this to Apriori, the effort to generate 1-itemset is saved by simply inspecting those SSM support counts that exceed the support threshold.

Furthermore, those 1-itemsets that do not meet the threshold will be discarded to reduce the number of higher level itemsets to be counted.

Another study to improve Apriori by using a novel data structure is the frequent pattern tree or FP-tree. It stores information about the frequent patterns. A FP-Tree-Based Method, called FP-growth is also proposed to mine frequent patterns and does not involve candidate generation. Not only did (Han *et al.*, 2000a) show that FP-growth performs an order of magnitude better than Apriori, it also showed that it is more scalable.

From iterative methods like Apriori and DHP to innovative use of data structures like SSM and FP-tree, research in mining large itemsets has made it more scalable and efficient. The need for faster and more scalable algorithms will continue because databases are getting bigger and bigger everyday. Research in distributed algorithms for finding large itemsets will gain more and more attention as more databases are integrated together. This presents a new level of challenge that demands more flexible algorithms in view of different representation of similar data, e.g., zip code maybe represented by string or integer, so that data do not need to be normalized before mining. More researches in parallel algorithms are also expected as grid computing is gaining interest.

Compare to the numerous researches done to search for better algorithms to mine large itemsets, the qualifying criterion-support threshold and the mechanism behind it counting have received much less attention. The problem with support threshold is that it requires expert knowledge which is subjective at best, to set this parameter in the system. Setting it arbitrarily low will qualify itemsets that should be left out, vice versa. Moreover, as database size increases, the support threshold may need to be adjusted.

GENERATING RULES

Research on rule generations mainly focus on newer algorithms, deriving more types of association rules and interestingness of the rules. Newer algorithms mostly employ innovative strategy like parallel computing (Zaki *et al.*, 1996) and Evolutionary Algorithm (EA) (Mata *et al.*, 2002). Newer rules types add dimension and quality to the traditional Boolean rule type. Examples are quantitative association rules (Srikant and Agrawal, 1996) and temporal association rules (Ale and Rossi, 2000). Newer criteria (Chan and Au, 1997; Natarajan and Shekar, 2005) on interestingness tend to be more objective than support and confidence.

Most of the association rules are generated by counting the number of occurrence of the rule in the database the confidence. Therefore, it is intuitive to partition the set of all frequent itemsets and count in parallel. Together with the 3 parallel algorithms to mine frequent itemsets, a parallel implementation for rule generation using the previously-stated approach is presented by Zaki *et al.* (1996).

In recent years, EA has been widely adopted in many scientific areas. EA borrows mechanisms of biological evolution and applies them in problem-solving. It is especially suitable for searching and optimization problems. Hence, the problem of mining association rules is a natural fit. EA is used to generate association rules (Mata *et al.*, 2002). It takes a population frequent itemsets as the initial population. Using EA that includes crossover and mutation of these itemsets, the population will evolve into one that contains itemsets with better and better fitness function. When the desired number of frequent itemsets is left in the population, the algorithm will stop. This novel way of generating/searching association rules allows for overlapping intervals in different itemsets. For example, one frequent itemset can have interval from (10 and 20) and another one can be (15 and 25). This is a sharp contrast to other techniques that divided the attributes into non-overlapping intervals. Rules that fall across two intervals may not be possible for discovery hence, a loss of information. This algorithm allows new rules to be discovered.

Research is also done in the types of association rules. At the beginning of data mining research, Boolean association rules dominated the field. Later on, more focus was put on quantitative association rules. Quantitative association rules are rules over quantitative and categorical attributes like age and marital status. Mining these rules involve partitioning the values of the attribute but may loss information as a result of the division. Srikant and Agrawal (1996) introduced an algorithm based on Apriori to mine quantitative rules. It also introduced a partial completeness to measure information loss due to partitioning and greater than expected interest as interestingness measure. Partial completeness is directly proportional to information loss. Given the minimum support and partial completeness by the user, the system can figure out the number of partitions needed. A quantitative rule is interesting only if it has greater than expected support and/or confidence specified by the user. The algorithm scales up linearly with experimental data (Srikant and Agrawal, 1996).

The time dimension is one that exists in all transaction. Therefore, it should be included in finding large itemsets, especially when not all items exist throughout the entire data gathering period. The temporal concept by limiting the search for frequent itemsets to the lifetime of the item set members is used by Ale and Rossi (2000). It also introduced the concept of temporal support, in addition to the normal support and confidence. The lifetime of an item is defined by the first and last time an item appears in the database. The temporal support is the minimum interval width. Thus, a rule is considered as long as there is enough support or temporal support. A byproduct of this approach is that old, obsolete itemsets can be deleted.

Any associations discovered by the above algorithms are eligible to be rules. The quality of those rules is measured by confidence. However, only those rules with confidence above a certain level are interesting and deserve attention. Most algorithms define interestingness in terms of user-supply thresholds for support and confidence. The problem is that these algorithms rely on the users to give suitable values. A new algorithm called APACS2 is proposed by Chan and Au (1997) that does not require such guesswork but makes use of an objective interestingness measure called adjusted difference. Moreover, APACS2 can discover both positive and negative association rules. A new concept of relatedness as an alternative approach to determine interestingness presented by Natarajan and Shekar (2005).

APACS2 uses adjusted difference as an objective interestingness measure. Adjusted difference is defined in terms of standardized difference and maximum likelihood estimate. More details can be found in Chan and Wong (1990) regarding the statistical theories. If the magnitude of the adjusted difference is >1.96 , i.e., 95 percentiles of the normal distribution, the association is regarded as significantly different and hence, interesting. If the adjusted difference is positive, it means the rule is likely, vice versa. The directional nature of adjusted difference gives association rules discovery a new dimension.

Interestingness can be subjective using support and confidence or objective using adjusted difference (Chan and Au, 1997). An opposite concept relatedness was introduced by Natarajan and Shekar (2005) to examine relationship between two items based on their co-occurrence frequency and context. Relatedness is meant to be used in lieu with interestingness to quantify association rules. Relatedness has three components:

- Average predictive ability of the presence of one item given the presence of the other
- The intensity of the occurrence of an item-pair with respect to other items
- The substitutability of another item for the items in the item-pair

These three measures give the strength of the relationship in terms of the frequency of the rule in relation to other items.

The study of rule generation started from a single Boolean type using subjective measures of support and confidence. It has grown to include various rule types with objective measures like adjusted difference and relatedness. The new researches have added both quality and quantity to rule generation. The quality is improved by using more objective measures to qualify rules. The quantity is increased by novel methodologies that enable mining rules in overlapping intervals and negative associations.

It is surprising to see that this topic is not researched more. The quality of the rules a system determined interesting is equally if not more, important than the speed and scale to find these rules because the real goal of data mining is to mine interesting rules. Only interesting rules are useful to help decision making. Uninteresting or trivial rules, albeit from larger database quicker do not. Therefore, more effort should continue to pour in to investigate more objective measures so that data mining can be parameter-free and operational, thus less subjective with higher quality. This way, domain experts can focus on interpreting the rules as oppose to worrying about how to tune the mining parameters to produce meaningful rules.

As more and more new data types are created, multimedia data for instance, more researches can also be done to define more association rule types. This may allow newer behavioral pattern to be analyzed and outcome predicted. The ability to profile multimedia data in its raw data format for data mining is particularly useful in medicine (Simoff *et al.*, 2002) or even homeland security.

CONCLUSION

Discovering association rules has been studied over a decade. A more thrust was focused on the performance and scalability of the algorithms but not enough attention was given to the quality (interestingness) of the rules generated. In future, the trend will be to turn the attention to the application of these researches in various areas of the lives, e.g., agriculture, genetic research, medicine, homeland security, etc.

Since, the more databases are integrated and the data in the databases are getting bigger and bigger, the algorithmic study on how to scan faster and more will receive less attention. Distributed algorithms that allow sharing of workload in a grid computing environment should gain more awareness. With more and more data is created outside of traditional database, data mining and rule discovery will grow out of scanning database tables into accessing data in its raw format, e.g., video clips. Performance and scalability issues will become more prominent. Newer rule types may be necessary to facilitate new data analysis. More objective measures of interestingness may also be required to avoid manipulation of rule discovery criteria by domain experts to produce desired results.

Discovery association rules will be a more thrust area in future. Based on the research and development in the past decade, application of association rule mining is growing in different sectors.

REFERENCES

- Agrawal, R., T. Imielinski and A. Swami, 1993. Mining association rules between sets of items in very large databases[C]. Proceedings of the ACM SIGMOD Conference on Management of Data, May 25-28, 1993, Washington, DC., USA., pp: 207-216.
- Ale, J.M. and G.H. Rossi, 2000. An approach to discovering temporal association rules. Proc. ACM Symp. Applied Comput., 1: 294-300.
- Chan, K.C.C. and A.K.C. Wong, 1990. APACS: A system for the automatic analysis and classification of conceptual patterns. Comput. Intell., 6: 119-131.
- Chan, K.C.C. and W.H. Au, 1997. An effective algorithm for mining interesting quantitative association rules. Proceedings of the ACM Symposium on Applied Computing, February 28-March 1, 1997, San Jose, CA., USA., pp: 88-90.
- Chen, M.S., J. Han and P.S. Yu, 1996. Data mining: An overview from a database perspective. IEEE Trans. Knowledge Data Eng., 8: 866-883.
- Han, E.H.S., G. Karypis and V. Kumar, 2000b. Scalable parallel data mining for association rules. Trans. Knowl. Data Eng., 12: 728-737.
- Han, J., J. Pei and Y. Yin, 2000a. Mining frequent patterns without candidate generation. Proceedings of ACM SIGMOD International Conference on Management of Data, May 15-18, 2000, Dallas, TX., USA., pp: 1-12.
- Houtsma, M. and A. Swami, 1993. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center.

- Mata, J., J.L. Alvarez and J.C. Riquelme, 2002. Evolutionary computing and optimization: An evolutionary algorithm to discover numeric association rules. Proceedings of the ACM Symposium on Applied Computing, March 2002, Madrid, Spain.
- Natarajan, R. and B. Shekar, 2005. Data Mining (DM): Poster papers: A relatedness-based data-driven approach to determination of interestingness of association rules. Proceedings of the ACM Symposium on Applied Computing, March 13-17, 2005, Santa Fe, NM., USA.
- Park, J.S., M.S. Chen and P.S. Yu, 1997. Using a hash-based method with transaction trimming for mining association rules. *IEEE Trans. Knowl. Data Eng.*, 9: 813-825.
- Simoff, S.J., C. Djeraba and O.R. Zaiane, 2002. MDM/KDD2002: Multimedia data mining between promises and problems. *ACM SIGKDD Explorations Newsletter*, 4: 118-121.
- Srikant, R. and R. Agrawal, 1996. Mining quantitative association rules in large relational tables. Proceedings of the International Conference on Management of Data, June, 1996, New York, USA., pp: 1-12.
- Zaki, M.J., M. Ogihara, S. Parthasarathy and W. Li, 1996. Parallel data mining for association rules on shared-memory multi-processors. Proceedings of the ACM/IEEE Conference on Supercomputing, November 17-22, 1996, Pittsburgh, PA., USA.