

Association rule evaluation for classification purposes

Fernando Berzal, Juan-Carlos Cubero, Nicolás Marín
Daniel Sánchez, Jose-María Serrano, Amparo Vila

Dpto. Ciencias de la Computación e Inteligencia Artificial

ETS Informática. Universidad de Granada. 18071 Granada

{fberzal | JC.Cubero | nicm | daniel | jmserrano | vila }@decsai.ugr.es

Resumen

Association rules have proved to be useful in building both partial and complete classification models. This paper analyzes alternative measures which could replace confidence in order to evaluate the suitability of a given association rule with respect to the classification problem we try to solve when building a classification model.

1. Introduction

Association rules [1] [2] add statistical measures, such as support and confidence, to the classic rules in Propositional Logic. In order to make that difference obvious, their representation uses \Rightarrow instead of the classic \rightarrow .

An association rule is an implication $X \Rightarrow Y$ where X and Y are itemsets with empty intersection (i.e. sets with no items in common). The intuitive meaning of such a rule is that when X appears, Y also tends to appear.

The confidence of an association rule $X \Rightarrow Y$ is the proportion of the transactions containing X which also contain Y . The support of the rule is the fraction of the database which contains both X and Y .

Classification models can be directly built from traditional association rules, using association rules where the class appears in their right-hand side. In the following Section, we describe some of the approaches followed to build classification models from association rules. Afterwards, we will analyze the use of alternative measures to evaluate the suitability

of association rules for classification models. Such a study is important in order to choose the best association rules what will become part of the classification model.

2. Association rules in classification

Some fundamental differences exist between classification and association rule discovery [11]. Association rules do not involve prediction, nor do they provide any mechanism to avoid underfitting and overfitting apart from the crude minimum support user-specified threshold. In classification problems, an inductive bias is also needed to solve classification problems, i.e. a basis for favoring one hypothesis over another (e.g. Occam's razor). This bias, like any other bias, must be domain-dependent. Association rules have, however, been used to solve classification problems directly.

In [3], association rules are used to build partial classification models in domains where conventional classifiers would be ineffective. For example, traditional decision trees are problematic when many values are missing and also when the class distribution is very skewed.

In [23], a tree of rules is built from an arbitrary set of association rules without using an ad-hoc minimum support threshold. The authors have observed that predictivity often depends on high confidence, and rules of high support tend to have low confidence, so the minimum support pruning is not suitable for their classification purposes.

CBA, Classification Based on Associations

[15], builds complete classification models. All “class association rules” are extracted from the available training dataset (i.e. all the association rules containing the class attribute in their consequent), and the most adequate rules are selected to build an “associative classification model”, which uses a default class to make it complete. This classifier builder uses a brute-force exhaustive global search, and yields excellent results when compared to C4.5 [20]. In [18], CBA performance was “improved” allowing multiple minimum support thresholds for the different problem classes and reverting to traditional TDIDT classifiers when no accurate rules are found.

A similar strategy to that of CBA is used to classify text documents into topic hierarchies in [24]. All the generalized association rules with the class attribute in their consequent are extracted, these rules are ranked, and some of them are selected to build a classifier which takes context into account, since class proximity is important when classifying documents into topics.

Other classification models inspired by the association rule mining process have been suggested in the literature.

LB [19], which stands for “Large Bayes”, is an extended Naïve Bayes classifier which uses an Apriori-like frequent pattern mining algorithm to discover frequent itemsets with their class support. This class support is an estimate of the probability of the pattern occurring with a certain class. The proposed algorithm achieves good results. However, it lacks the understandability of symbolic models such as decision trees.

Emerging patterns are itemsets whose support increases significantly from one dataset to another [9]. They have been used to build classifiers following the LB philosophy. For example, CAEP [10] finds all the emerging patterns meeting some support and growth rate thresholds for each class. It then computes an aggregated differentiating score to determine the most suitable class for a given instance. This computation allows the algorithm to perform well when the class populations are unbalanced, although it gives no further insight into

the data.

Liu, Hu and Hsu [16] propose the use of a hierarchical representation consisting of general rules and exceptions in order to replace the usual flat representation model where too many association rules hamper the understanding of the underlying data. The same approach is followed in [17] in order to obtain a good summary of the knowledge contained in an arbitrary set of rules.

CMAR, Classification based on Multiple Association Rules [14], extends the FP-growth [12] frequent pattern mining algorithm to obtain and compactly store rules for classification in a CR-tree. Instead of relying on a single rule to classify data, CMAR considers sets of related rules, taking into account that the most confident rule might not always be the best choice to classify data. Given a data object, CMAR retrieves all the rules matching the object and assigns a label class to the object according to a weighted χ^2 measure which indicates the “combined effect” of the rules.

CPAR, Classification based on Predictive Association Rules [25], avoids generating the large number of candidate rules of CMAR and generates candidate rules directly from the training data using ideas taken from traditional rule-based classification methods such as Quinlan’s FOIL [21].

Finally, ART [5] can be viewed as a degenerate, polythetic decision tree, hence its name, which stands for Association Rule Tree. ART can also be viewed as a decision list learner, since it is a generalized “Separate and Conquer” algorithm, in contrast to the standard “Divide and Conquer” algorithms used to build decision trees. ART is faster than general decision list and rule inducers which need to discover rules one at a time and, as decision tree learners, ART is able to build classification models in an efficient and scalable way. ART classifiers, however, tend to be smaller than the decision trees generated by standard TDIDT approaches.

3. Rule evaluation

Even though accurate classification models can be built using standard association rules [5] [25], it is clear that confidence is not the best measure for a classification rule. Alternative rule evaluation measures might provide better insight into the capability of a given association rule to classify data.

Before delving into the details of existing measures, we should review what we consider a good classification rule. An association rule $A \Rightarrow C$ will be useful in classification problems when the logical implication $A \rightarrow C$ is valid:

- C happens more often when A holds.
- $\neg C$ should be less frequent when A holds.

We might also be interested in verifying the validity of $\neg C \rightarrow \neg A$, which is mathematically equivalent to $A \rightarrow C$. Hence, a potentially useful rule should also verify the following properties:

- $\neg A$ happens more often when $\neg C$.
- A should occur less frequently when C does not hold.

Although the latter two properties might be interesting from a logical point of view, they do not directly affect classification accuracy, where we are interested only in determining the class C given A . Even then, those properties might be desirable to improve the understandability of the classification model obtained.

The following paragraphs discuss some rule evaluation measures and their properties in the context of classification problems.

3.1. Confidence

Even when the rule support might be of interest during the association discovery process from a purely tactical point of view, the rule confidence is what finally determines the rule validity. Its meaning is easy to grasp, and it can be defined as follows:

$$\text{conf}(A \Rightarrow C) = P(C|A) = \frac{P(A \cap C)}{P(A)}$$

Confidence is used to measure the association between the antecedent A and the consequent C : the higher the confidence of $A \Rightarrow C$, the lower the confidence of $A \Rightarrow \neg C$, since $P(C|A) + P(\neg C|A) = 1$.

However, the rule confidence cannot be used to establish a causality relationship between antecedent and consequent. It cannot be used to perform inferences since it does not take into account the validity of $\neg C \Rightarrow \neg A$. A variation of the confidence does, hence its name: Causal confidence [13].

3.2. Causal confidence

The causal confidence measure considers the confidence of the rule $A \rightarrow C$ and the confidence of its counterpart $\neg C \rightarrow \neg A$:

$$\text{conf}_{\text{causal}}(A \Rightarrow C) = \frac{1}{2}(P(C|A) + P(\neg A|\neg C))$$

The average is used just to normalize the measure so that its values are always in the interval $[0, 1]$.

Unfortunately, this measure could have a high value even when the implication $A \rightarrow C$ does not have a representative support but its counterpart $\neg C \rightarrow \neg A$ does. Therefore, causal confidence is not adequate for classification problems.

3.3. Causal support

The idea of causal confidence can also be applied to the support measure in order to obtain the causal support:

$$\begin{aligned} \text{support}_{\text{causal}}(A \Rightarrow C) &= \\ &= P(A \cap C) + P(\neg A \cap \neg C) \end{aligned}$$

As before, even when $P(A \cap C)$ is really low, $P(\neg A \cap \neg C)$ can be high, causing causal support to be high (even when the rule might not be too useful in a classification problem). Consequently, causal support is not a good choice to evaluate rules in classification problems.

3.4. Confirmation

Another measure, called confirmation, takes into account when the antecedent holds

but not its consequent, what might be used to highlight a rule depending on the commonness of its antecedent.

$$K(A \Rightarrow C) = P(A \cap C) - P(A \cap \neg C)$$

Since $P(A \cap \neg C) = P(A) - P(A \cap C)$, confirmation is reduced to $P(A \cap C) - P(A \cap \neg C) = 2P(A \cap C) - P(A)$. Since $P(A \cap C)$ is the rule support and $P(A)$ is the support of the antecedent, confirmation is no more than

$$K(A \Rightarrow C) = \text{support}(A \Rightarrow C) - \text{support}(A)$$

Hence

$$K(A \Rightarrow C) \leq \text{support}(A \Rightarrow C) \quad (1)$$

The second term just penalizes the support of the rule according to the support of the antecedent. In other words, given two equally common rules, confirmation would select as most interesting the rule whose antecedent is less common. That might be useful in the knowledge discovery process, although it does not help in classification problems because it does not take into account relationship between the rule antecedent and the rule consequent.

This measure has been used to obtain three new criteria to evaluate rules:

- **Causal confirmation**, takes into account $\neg C \rightarrow \neg A$ and can be reduced to:

$$\begin{aligned} K_{causal}(A \Rightarrow C) = & \\ & \text{support}_{causal}(A \Rightarrow C) \\ & - \text{support}(A) \\ & + \text{support}(C) \end{aligned} \quad (2)$$

As the standard confirmation, this measure varies depending upon the support of A . On the other hand, the more common the class C is, the higher causal confirmation the rule will have, what certainly makes classification difficult (especially when the class distribution is skewed). Apart from this fact, causal support is not suitable for classification problems.

III Taller de Minería de Datos y Aprendizaje

- **Confirmed confidence** is obtained from the standard confidence when we try to measure the quality of the rule $A \Rightarrow \neg C$ (that is, when the antecedent holds but not the consequent):

$$\begin{aligned} \text{conf}_{confirmed}(A \Rightarrow C) = & \\ = P(C|A) - P(\neg C|A) \end{aligned}$$

However, since $P(\neg C|A) = 1 - P(C|A)$, confirmed confidence can be reduced to

$$\begin{aligned} \text{conf}_{confirmed}(A \Rightarrow C) = & \\ = 2 \cdot \text{conf}(A \Rightarrow C) - 1 \end{aligned}$$

Therefore, if we are using this measure just to rank the candidate rules which might be part of a classification model, the rule confirmed confidence is completely equivalent to the rule confidence (only that the confirmed confidence is defined over the $[-1, 1]$ interval).

- Even a **causal confirmed confidence** can be concocted from the two previous measures, but no interesting results can be expected when dealing with classification problems.

3.5. Conviction

Conviction [6] was introduced as an alternative to confidence to mine association rules in relational databases (implication rules using their authors' nomenclature).

$$\text{conviction}(A \Rightarrow C) = \frac{P(A)P(\neg C)}{P(A \cap \neg C)}$$

Similar in some sense to confirmation, conviction focuses on $A \Rightarrow \neg C$:

$$\text{conviction}(A \Rightarrow C) = \frac{\text{support}(\neg C)}{\text{conf}(A \Rightarrow \neg C)} \quad (3)$$

The lower $P(A \cap \neg C)$, the higher the rule conviction, what makes conviction ideal for discovering rules for uncommon classes.

However, the conviction domain is not bounded and it is hard to compare conviction values for different rules. This constrains the usability of the conviction measure in classification models such as ART [5], since no heuristics to automatically settle a conviction threshold can be devised.

Later in this paper, another bounded measure will be analyzed which is completely equivalent to conviction in the situations of interest for classification problems: the well-known Shortliffe and Buchanan's certainty factors. When there are at least two classes (i.e. $support(C) < 1$) and the rule improves classifier accuracy (i.e. $CF(A \Rightarrow C) > 0$), certainty factors can be defined as

$$CF(A \Rightarrow C) = 1 - \frac{1}{conviction(A \Rightarrow C)}$$

This allows us to substitute conviction for another measure whose domain is bounded. Further properties of the certainty factor will be discussed in 3.9.

3.6. Interest

A rule interest measure was defined in [22] as

$$interest(A \Rightarrow C) = \frac{P(A \cap C)}{P(A)P(C)} = \frac{P(C|A)}{P(C)}$$

The more common A and C , the less interest the rule will have, which is certainly useful to guide the knowledge discovery process.

Among its properties, its symmetry stands out: the interest of $A \Rightarrow C$ equals to the interest of $C \Rightarrow A$.

As happened with conviction, its domain is not bounded, what might make its interpretation harder in a classification model.

In some sense, it can be considered complementary to conviction if we take into account the following equality and compare it with equation 3, although interest focuses on the association $A \Rightarrow C$ while conviction focuses on $A \Rightarrow \neg C$:

$$interest(A \Rightarrow C) = \frac{confidence(A \Rightarrow C)}{support(C)} \quad (4)$$

3.7. Dependency

The following measure is the discrete equivalent of correlation in continuous domains:

$$dependency(A \Rightarrow C) = |P(C|A) - P(C)|$$

In classification problems, it is not suitable since its value is high for common classes even when the corresponding rule confidence is minimum:

$$\begin{aligned} dependency(A \Rightarrow C) &= \\ &= |conf(A \Rightarrow C) - support(C)| \end{aligned}$$

A causal variation [13] of this measure can also be defined as follows, although it is not useful in classification problems either.

Bhandari's attribute focusing measure is also derived from the dependency measure above, as the following expression shows

$$\begin{aligned} Bhandari(A \Rightarrow C) &= \\ &= support(A) \cdot dependency(A \Rightarrow C) \end{aligned}$$

Therefore, it is of no use in classification problems.

3.8. Hellinger's divergence

Hellinger's divergence was devised to measure the amount of information a rule provides [7] and it can be viewed as a distance measure between a priori and a posteriori class distributions:

$$\begin{aligned} H(A \Rightarrow C) &= \sqrt{P(A)} \cdot \\ &[(\sqrt{P(A \cap C)} - \sqrt{P(C)})^2 \\ & - (\sqrt{1 - P(A \cap C)} - \sqrt{1 - P(C)})^2] \end{aligned}$$

This measure has been used in classifiers before and will be evaluated in Section 4.

3.9. Certainty factors

Certainty factors were introduced by Shortliffe and Buchanan to represent uncertainty in the MYCIN expert system. Its use in association rule mining has been proposed in [4].

The certainty factor of a rule $A \Rightarrow C$ is defined as

$$CF(A \Rightarrow C) = \frac{conf(A \Rightarrow C) - support(C)}{1 - support(C)}$$

when $conf(A \Rightarrow C) > support(C)$,

$$CF(A \Rightarrow C) = \frac{conf(A \Rightarrow C) - support(C)}{support(C)}$$

when $conf(A \Rightarrow C) < support(C)$, and

$$CF(A \Rightarrow C) = 0$$

when $conf(A \Rightarrow C) = support(C)$.

This rule evaluation measure can be viewed as the variation degree of the probability of C when A holds. The larger a positive certainty factor, the smaller the decrease of the probability of C not being when A holds.

In extreme situations, the rule confidence determines its certainty factor:

$$conf(A \Rightarrow C) = 1 \Rightarrow CF(A \Rightarrow C) = 1$$

$$conf(A \Rightarrow C) = 0 \Rightarrow CF(A \Rightarrow C) = -1$$

Certainty factor take into account the probability of C apart from the rule confidence. They also verify an interesting property when they are positive (which is when the rules are useful for classification):

$$CF(A \Rightarrow C) = CF(\neg C \Rightarrow \neg A) \quad (5)$$

In classification problems, we could face different situations where certainty factors behavior are at their best:

1. If our problem includes a skewed class distribution, and two candidate rules hold the same confidence value but correspond to classes of different frequency, the rule corresponding to the less common class has a higher certainty factor:

$$conf(A \Rightarrow C) = conf(B \Rightarrow D) \leq 1$$

$$support(C) > support(D)$$

↓

$$CF(A \Rightarrow C) \leq CF(B \Rightarrow D)$$

III Taller de Minería de Datos y Aprendizaje

2. Under some circumstances, comparing certainty factors is equivalent to comparing confidence values:

$$CF(A \Rightarrow C_1) > CF(B \Rightarrow C_2)$$

can be reduced to

$$conf(A \Rightarrow C_1) > conf(B \Rightarrow C_2)$$

when

- Both rules refer to the same class c_k .
- Both rules correspond to classes with the same probability: $support(c_1) = support(c_2)$.

3. Finally, there exist situations where higher certainty factors do not correspond to a higher confidence values. Given two rules so that

$$CF(A \Rightarrow C) > CF(B \Rightarrow D)$$

- When C is more common than D :

$$conf(A \Rightarrow C) > K \cdot conf(B \Rightarrow D)$$

$$K = \frac{support(\neg C)}{support(\neg D)} < 1$$

- When C is less common than D :

$$conf(A \Rightarrow C) > conf(B \Rightarrow D) - \Delta$$

$$\Delta = \frac{support(D) - support(C)}{support(\neg C)} > 0$$

In summary, even though certainty factors are intimately linked to confidence values, a higher certainty factor does not imply a higher confidence value:

$$CF(A \Rightarrow C) > CF(B \Rightarrow D)$$

⇒

$$confianza(A \Rightarrow C) > confianza(B \Rightarrow D)$$

The relative frequency of each class determines the higher certainty factors.

Let us suppose that we have two association rules: $A \Rightarrow c_1$ with 2% support and 50% confidence, and $B \Rightarrow c_2$ with 50% support and 95% confidence. The certainty factors of such rules would be $3/8$ and $3/4$ if c_1 has a 20% support and c_2 has a 80% support. However, if the class distribution varies, being now 10% of c_1 and 90% of c_2 , when certainty factors would be $4/9$ and $1/2$, keeping their relative order. However, if the class distribution is inverted and now c_1 has a 94% support while c_2 only has a 6% support, when certainty factors would become $46/94$ and $1/6$, being the second lower than the first!

Situations such as the one described in the paragraph above should be used as a warning sign when comparing certainty factors to choose the rules to include in a classification model. Sometimes they might be useful, as when they prefer rules corresponding to uncommon classes, although you should also expect shocking difference between classification models when class distributions change.

3.10. Just a matter of usefulness

Certainty factor properties suggest an additional pruning step when considering association rules in order to build classification models. When you build such models, only association rules with a positive certainty factor are really useful, since they increase our knowledge and improve classifier accuracy.

By definition, a positive certainty factor is obtained for a rule $A \Rightarrow C$ when

$$\text{conf}(A \Rightarrow C) > \text{support}(C)$$

This constraint indicates that the use of the rule $A \Rightarrow C$ improves the classification model which would result from using a default class (at least with respect to C in the training set).

That idea can be used to prune association rules which do not verify the above requirement, which can be expressed as

$$P(A \cap C) > P(A) \cap P(C)$$

That rule pruning can be viewed as a ‘usefulness criterion’ which reduces the number of candidate association rules which can become part of the classification model

4. Experimental results

Some experiments have been performed to check the influence of the rule evaluation measure on the process of building a classification model. We have tested different criteria using the ART [5] classification model as a test bed, since ART does not require any specific measure to evaluate the rules the ART classifier is built from. Our experiments try to estimate the suitability of the measures proposed in the previous section, using 10-folded cross validation and the same datasets which were used in [5]. In all our experiments with ART, we used a 5% minimum relative support threshold and the automatic threshold selection heuristics described in [5]. Table 1 summarizes our results.

The following observations can be made from the results we have obtained:

- The usefulness criterion proposed in 3.10 consistently improves ART classification accuracy. Moreover, it improves accuracy without modifying the evaluation measure used during the association rule mining process (that is, the rule confidence). However, this increased accuracy comes at a cost: the increased complexity of the resulting classifier. The resulting ART tree has more leaves and, therefore, training time is somewhat higher since the training dataset must be scanned more times to build the classifier. This result is just an incarnation of the typical trade-off between classifier accuracy and classifier complexity.
- Certainty factors do not improve ART overall performance, probably due to some of their counterintuitive properties (see section 3.9).
- As we could expect from the properties analyzed in the previous section, the use of conviction achieves results which are similar to the results obtained by using certainty factors. From a classification point of view, conviction and certainty factors are equivalent when it is interesting to include a given association rule in the clas-

<i>Dataset</i>	<i>Records</i>	<i>Attr</i>	<i>Classes</i>
AUDIOLOGY	226	70	24
CAR	1728	7	4
CHESSE	3196	36	2
HAYES-ROTH	160	5	3
LENSES	24	6	3
LUNG CANCER	32	57	3
MUSHROOM	8124	23	2
NURSERY	12960	9	5
SOYBEAN	683	36	19
SPLICE	3175	61	3
TICTACTOE	958	10	2
TITANIC	2201	4	2
VOTE	435	17	2

Table 2: Datasets used in our experiments (from the UCI Machine Learning Repository).

sification model, as was mentioned in section 3.5.

- The use of the interest measure (section 3.6) leads to the results we could expect in ART: since this measure is not bounded, the automatic threshold selection criterion in ART does not work properly. Perhaps, the additive tolerance margin might be replaced by a multiplicative tolerance factor. Even then, the definition of the interest measure makes it difficult to establish an initial desirable interest value. Such value might depend on the particular problem and, therefore, the use of the interest measure is not practical in ART. Bounded measures, such as confidence or certainty factors, will be preferred.
- The same rationale applies to Hellinger's divergence (section 3.8). Even when its range is bounded, the interpretation and comparison of Hellinger's divergence values make this measure impractical in classification models such as ART. In fact, no acceptable classification models were built by ART because it is hard to establish an initial desirable value for Hellinger's divergent (a prerequisite to make use of ART automatic parameter setting).

5. Conclusions

In summary, from all the measures we discussed above, only certainty factors and the so-called usefulness criterion are good alternatives to confidence when building ART classification models. Conviction also achieves good results, although certainty factors are preferred since they are equivalent to conviction in the cases the classification process is more interested in (and certainty factors are bounded).

Despite the experimental results, it should be noted that all rule evaluation criteria have their home grounds and their use might be suitable depending on what the user intends to obtain. The availability of a wide variety of rule evaluation measures is a good signal, since it provides us a toolkit to draw on.

Moreover, the use of a measure or another does not affect the computational cost of the rule discovery process. When building classification models such as ART, that cost is proportional to the classification model complexity. Therefore, the study of alternative rule evaluation measures keeps its interest. Such measures are just criteria at our disposal which can be used to guide the knowledge discovery process according to the particular goals and needs of a given problem.

Referencias

- [1] R. Agrawal, T. Imielinski, and A. Swami. Mining association rules between sets of items in large databases. *ACM SIGMOD Record*, 22(2):207–216, June 1993.
- [2] R. Agrawal and R. Srikant. Fast algorithms for mining association rules. In *Proc. 20th Int. Conf. Very Large Data Bases, VLDB*, pages 487–499, September 12–15 1994.
- [3] K. Ali, S. Manganaris, and R. Srikant. Partial classification using association rules. In *Proceedings of the 3rd International Conference on Knowledge Discovery in Databases and Data Mining, Newport Beach, California, USA*, pages 115–118, August 14–17 1997.
- [4] F. Berzal, I. Blanco, D. Sanchez, and M. Vila. Measuring the accuracy and interest of association rules: A new framework. *Intelligent Data Analysis*, 6(3):221–235, 2002.

	Confidence	Usefulness	CF	Conviction	Interest	Hellinger
Accuracy (10-CV)	79.22 %	83.10 %	77.67 %	77.79 %	53.71 %	48.05 %
Training time	18.5s	22.4s	12.7s	10.7s	2.9s	1.8s
Tree topology						
- Leaves	36.9	50.0	28.5	30.0	4.2	1
- Internal nodes	18.3	25.2	16.2	17.3	2.6	0
- Average depth	7.41	8.71	7.39	7.33	2.12	1.00
I/O operations						
- Records	36400	50100	33100	26700	20300	7000
- Scans	63	89	55	59	11	3

Table 1: Experiment summary using different rule evaluation criteria.

- [5] F. Berzal, J. C. Cubero, D. Sánchez, and J. Serrano. Art: A hybrid classification model. *Machine Learning*, 54(1):67–92, January 2004.
- [6] S. Brin, R. Motwani, J. Ullman, and S. Tsur. Dynamic itemset counting and implication rules for market basket data. In *Proceedings of the ACM SIGMOD international conference on Management of Data, Tucson, Arizona, USA*, pages 255–264, May 11-15 1997.
- [7] L. C.-H. and S. D.-G. A multistrategy approach to classification learning in databases. *Data & Knowledge Engineering*, 31:67–93, 1999.
- [8] W. Cohen. Fast effective rule induction. In *Proceedings of the 12th International Conference on Machine Learning, Tahoe City, CA USA*, pages 115–123, July 1995.
- [9] G. Dong and J. Li. Efficient mining of emerging patterns: discovering trends and differences. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge Discovery and Data Mining, San Diego, CA USA*, pages 43–52, August 15–18 1999.
- [10] G. Dong, X. Zhang, L. Wong, and J. Li. Caep: Classification by aggregating emerging patterns. In *Proceedings of the Second International Conference on Discovery Science, Tokyo, Japan*, pages 30–42, 1999.
- [11] A. A. Freitas. Understanding the crucial differences between classification and discovery of association rules - a position paper. *SIGKDD Explorations*, 2(1):65–69, July 2000.
- [12] J. Han, J. Pei, and Y. Yin. Mining frequent patterns without candidate generation. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data, Dallas, TX USA*, pages 1–12, May 15-18 2000.
- [13] Y. Kodratoff. Comparing machine learning and knowledge discovery in databases. *Lecture Notes in Artificial Intelligence, LNAI*, 2049:1–21, 2001.
- [14] W. Li, J. Han, and J. Pei. Cmar: Accurate and efficient classification based on multiple class-association rules. In *Proceedings of the 2001 IEEE International Conference on Data Mining (ICDM'01), San Jose, California, USA*, pages 208–217, November 29 - December 02 2001.
- [15] B. Liu, W. Hsu, and Y. Ma. Integrating classification and association rule mining. In *Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining (KDD-98), New York City, USA*, pages 80–86, August 27–31 1998.
- [16] B. Liu, M. Hu, and W. Hsu. Intuitive representation of decision trees using general rules and exceptions. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000), Austin, Texas*, pages 30–42, July 30 - August 3 2000.
- [17] B. Liu, M. Hu, and W. Hsu. Multi-level organization and summarization of the discovered rule. In *Proceedings of the sixth ACM SIGKDD International Conference on*

- Knowledge Discovery and Data Mining, Boston, MA USA*, pages 208–217, August 20 - 23 2000.
- [18] B. Liu, Y. Ma, and C. Wong. Improving an association rule based classifier. In *Proceedings of the Fourth European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD-2000), Lyon, France*, pages 80–86, September 13–16 2000.
- [19] D. Meretakis and B. Wuthrich. Extending naive bayes classifiers using long itemsets. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge Discovery and Data Mining, San Diego, CA USA*, pages 165–174, August 15–18 1999.
- [20] J. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993.
- [21] J. R. Quinlan and R. M. Cameron-Jones. Foil: A midterm report. In *Proceedings of the 1993 European Conference on Machine Learning, Vienna, Austria*, pages 3–20, 1993.
- [22] C. Silverstein, S. Brin, and R. Motwani. Beyond market baskets: Generalizing association rules to dependence rules. *Data Mining and Knowledge Discovery*, 22(2):39–68, 1998.
- [23] K. Wang, S. Zhou, and Y. He. Growing decision trees on support-less association rules. In *Proceedings of the sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Boston, MA USA*, pages 265–269, August 20–23 2000.
- [24] K. Wang, S. Zhou, and S. Liew. Building hierarchical classifiers using class proximity. In *Proceedings of 25th International Conference on Very Large Data Bases (VLDB'99), Edinburgh, Scotland, UK*, pages 363–374, September 7–10 1999.
- [25] X. Yin and J. Han. Cpar: Classification based on predictive association rules. In *Proceedings of the Third SIAM International Conference on Data Mining, San Francisco, CA, USA*, pages 208–217, May 1-3 2003.