

Resolving Rule Conflicts with Double Induction

Tony Lindgren and Henrik Boström

Department of Computer and Systems Sciences,
Stockholm University and Royal Institute of Technology,
Forum 100,
164 40 Kista, Sweden
tony.henke@dsv.su.se
<http://www.dsv.su.se>

Abstract. When applying an unordered set of classification rules, the rules may assign more than one class to a particular example. Previous methods of resolving such conflicts between rules include using the most frequent class in the conflicting rules (as done in CN2) and using naïve Bayes to calculate the most probable class. An alternative way of solving this problem is presented in this paper: by generating new rules from the examples covered by the conflicting rules. These newly induced rules are then used for classification. Experiments on a number of domains show that this method significantly outperforms both the CN2 approach and naïve Bayes.

1 Introduction

Two major induction strategies are used in top-down rule induction: Divide-and-Conquer (DAC) [15] and Separate-and-Conquer (SAC) [7]. The former strategy does not generate overlapping rules and hence no conflict between rules can occur, while the latter strategy can give rise to overlapping rules. The rules can be ordered to avoid conflicts (which results in so called decision lists) [16, 3] or they may be used without any ordering [2]. In the latter case one has to deal with conflicts that can occur when two or more rules cover the same example but assign different classes to it.

This work addresses the problem of handling conflicting rules and how to solve the conflicts in a more effective way than using existing methods. Previous methods that deals with conflicts among rules either returns the most frequent class covered by the conflicting rules (as done in CN2 [2]) here referred to as *frequency-based classification*, or uses naïve Bayes [13] to calculate the most probable class, as done in the rule induction system RDS [1].

1.1 Motivation

Why bother solving rule conflicts when they can be avoided? When comparing SAC with unordered rules to DAC it seems to the case that SAC induce fewer rules which also are less complex than the rules induced by DAC. The complexity

of the DAC rules is caused by irrelevant attributes. These attributes also makes it hard to interpret DAC rules, because there is an uncertainty if the conditions of the rule are relevant or not. SAC rules does not use conditions that are irrelevant for the classification. The use of irrelevant attribute is related to the replication problem [14] that DAC suffers from. The replication problem is best illustrated by an example: the smallest decision tree for the boolean function $(X_1 \wedge x_2) \vee (x_3 \wedge x_4)$ is shown in Fig. 1. In that figure we see two similar structures that are repeated, this repetition could be avoided using SAC. A comparison of DAC, SAC with ordered rules and SAC with unordered rules on this problem is shown in Fig. 2.




Fig. 1. The smallest decision tree for the boolean function $(x_1 \wedge x_2) \vee (x_3 \wedge x_4)$. Two similar structures are replicated.

There exist at least two strong reasons of why to use SAC with unordered rules instead of SAC with ordered rules. The first reason is that classification using unordered rules compared to using ordered rules generally has a higher prediction accuracy. The second reason is that with ordered rules, each rule is dependent on the rules before it in the ordering. This means that the rules can not be interpreted by them self, something which unordered rules can. Another nice property that unordered rules has is that each class has their own set of rules. Often in ordered rules one class is regarded as a default class this means that the last rule in the ordering has no conditions and cover the default class. This line of reasoning is known as closed world assumption. As a consequence it is not easy to get understanding of why a particular example was classified as the default class.

```

DAC:
IF  $X_1 = 0$  AND  $X_3 = 0$  THEN  $C = 0$ 
IF  $X_1 = 0$  AND  $X_3 = 1$  AND  $X_4 = 0$  THEN  $C = 0$ 
IF  $X_1 = 0$  AND  $X_3 = 1$  AND  $X_4 = 1$  THEN  $C = 1$ 
IF  $X_1 = 1$  AND  $X_2 = 1$  THEN  $C = 1$ 
IF  $X_1 = 1$  AND  $X_2 = 0$  AND  $X_3 = 0$  THEN  $C = 0$ 
IF  $X_1 = 1$  AND  $X_2 = 0$  AND  $X_3 = 1$  AND  $X_4 = 0$  THEN  $C = 0$ 
IF  $X_1 = 1$  AND  $X_2 = 0$  AND  $X_3 = 1$  AND  $X_4 = 1$  THEN  $C = 1$ 

SAC ordered:
IF  $X_1 = 1$  AND  $X_2 = 1$  THEN  $C = 1$  ELSE
IF  $X_3 = 1$  AND  $X_4 = 1$  THEN  $C = 1$  ELSE
(Default rule):  $C = 0$ 

SAC unordered:
IF  $X_1 = 1$  AND  $X_2 = 1$  THEN  $C = 1$ 
IF  $X_3 = 1$  AND  $X_4 = 1$  THEN  $C = 1$ 
IF  $X_1 = 0$  AND  $X_3 = 0$  THEN  $C = 0$ 
IF  $X_1 = 0$  AND  $X_4 = 0$  THEN  $C = 0$ 
IF  $X_2 = 0$  AND  $X_3 = 0$  THEN  $C = 0$ 
IF  $X_2 = 0$  AND  $X_4 = 0$  THEN  $C = 0$ 

```

Fig. 2. Rules describing the boolean function from DAC, SAC with and without ordering

In this work we propose a novel way of resolving such conflicts: by applying rule induction on the examples covered by the conflicting rules, in order to generate a new set of rules that can be used instead of the conflicting rules. The motivation for this method is that by focusing on the examples within the region of interest (i.e. where the example to be classified resides), one is likely to obtain rules that better separate classes in this region compared to having to consider all examples, since examples outside the region may dominate the rule induction process, making the separation of classes within the region of marginal importance. This novel method is called *Double Induction*.

The paper is organised as follows. In section 2, we first review frequency-based classification and naïve Bayes to resolve rule conflicts, and then describe Double Induction. In section 3, the three different methods are compared empirically and the results are presented. In the discussion (Section 4) the results are analysed and relations to other work is made. Finally, in section 5, conclusions are made and pointers to future work is given.

2 Ways of Resolving Classification Conflicts

In this section, we first recall two previous methods for resolving classification conflicts among overlapping rules and then introduce the novel method, Double Induction.

2.1 Frequency-Based Classification

The system CN2 [2] resolves classification conflicts between rules in the following way. Given the examples in Fig. 3, the class frequencies of the rules that covers the example to be classified (marked with '?') are calculated:

$$C(+) = covers(R_1, +) + covers(R_2, +) + covers(R_3, +) = 32$$

and

$$C(-) = covers(R_1, -) + covers(R_2, -) + covers(R_3, -) = 33$$

where $covers(R, C)$ gives the number of examples of class C that are covered by rule R . This means that CN2 would classify the example as belonging to the negative class (-). More generally:

$$FreqBasedClassification = \underset{Class_i \in Classes}{argmax} \sum_{j=1}^{|CovRules|} covers(R_j, C_i)$$

where $CovRules$ is the set of rules that cover the example to be classified.




Fig. 3. Three rules covering an example to be classified (marked with ?). The training examples are labelled with their respective classes (+ and -).

2.2 Naïve Bayes Classification

Bayes theorem is as follows:

$$P(C|R_1 \wedge \dots \wedge R_n) = P(C) \frac{P(R_1 \wedge \dots \wedge R_n|C)}{P(R_1 \wedge \dots \wedge R_n)}$$

where C is a class label for the example to be classified and $R_1 \dots R_n$ are the rules that cover the example. As usual, since $P(R_1 \wedge \dots \wedge R_n)$ does not affect the relative order of different hypotheses according to probability, it is ignored. Assuming (naively) that $P(R_1 \wedge \dots \wedge R_n|C) = P(R_1|C) \dots P(R_n|C)$, the maximum a posteriori probable hypothesis (MAP) is:

$$h_{MAP} = \underset{Class_i \in Classes}{\operatorname{argmax}} P(Class_i) \prod_{R_j \in Rules}^{|\text{Rules}|} P(R_j|Class_i)$$

where $Rules$ is the set of rules that cover the example to be classified.

If we again consider the example shown in Fig. 3, we get:

$$\begin{aligned} P(+|R_1 \wedge R_2 \wedge R_3) &= P(+)*P(R_1|+)*P(R_2|+)*P(R_3|+) = \\ &40/80 * 12/40 * 14/40 * 6/40 = 0.0079 \end{aligned}$$

$$\begin{aligned} P(-|R_1 \wedge R_2 \wedge R_3) &= P(-)*P(R_1|-)*P(R_2|-)*P(R_3|-) = \\ &40/80 * 3/40 * 3/40 * 27/40 = 0.0019 \end{aligned}$$

Hence naïve Bayes assigns the positive (+) class to the example. Note that if a rule involved in a conflict does not cover any examples of a particular class, this would eliminate the chances for that class to be selected, even if there are several other rules that cover the example with a high probability for that class. To overcome this problem, we use Laplace- m correction (described in [8]) in the experiments.

2.3 Double Induction

The idea of Double Induction is to induce new rules based on the examples that are covered by the rules in conflict. By doing this we obtain a completely fresh set of rules that are tailor made to separate the examples in this subset of the whole domain. By concentrating on a small subspace of the example space there is a higher chance to find rules that separate the classes better. The training set is randomly divided into two halves in order to create a grow set (to induce rules) and a prune set (to prune rules). This division is likely to further reduce the probability that we will find the same rules as before. The Double Induction algorithm is given in Fig. 4.

Input: R_1 = rules from the first induction round, e = example to be classified,
 E_1 = training examples
Output: C = a class assigned to e

```

collect all rules  $R_{1,e} \subseteq R_1$  that cover  $e$ 
if conflictingRules( $R_{1,e}$ ) then
    collect all training examples  $E_2 \subseteq E_1$  covered by  $R_{1,e}$ 
    induce new rules  $R_2$  from  $E_2$ 
    collect all rules  $R_{2,e} \subseteq R_2$  that cover  $e$ 
    if conflictingRules( $R_{2,e}$ ) then
        let  $C = \text{naiveBayes}(R_{2,e}, E_2)$ 
    else let  $C = \text{majorityClass}(R_{2,e})$ 
else let  $C = \text{majorityClass}(R_{1,e})$ 

```

Fig. 4. The Double Induction algorithm.

It is worth noting that when solving conflicts with naïve Bayes on the newly induced rules, the a priori probability is computed from the examples covered by the previously conflicting rules. This means that the a priori probability reflects the probability distribution of this subspace (contained by the rules in conflict), rather than the a priori probability of the whole domain.

Consider again the scenario shown in Fig. 3. Given the examples covered by R_1 , R_2 and R_3 , Separate-and-Conquer may come up with the three new rules shown in Fig. 5. The unlabelled example is then classified using these newly induced rules.

In our hypothetical scenario, the example is covered by R_6 resulting in that the positive class is assigned to the example.

In our experiments we use two different sets of examples as input in the second rule induction round: one which has no duplicates of examples and one with the concatenation of all the covered examples of every rule, in which some examples may be present more than once (i.e., a multi-set). The latter is in a sense a weighting scheme of the examples in the example set. One reason for using such a weighting scheme is that it has been empirically demonstrated that the examples in the intersection are more important (and hence should have more weight) than other examples when doing the classification, see [9]. Note that this type of weighting is implicitly used in both frequency-based classification and naïve Bayes.

3 Empirical Evaluation

Double Induction has been implemented as a component to the Rule Discovery System (RDS) [1], which is a rule based machine learning system that sup-

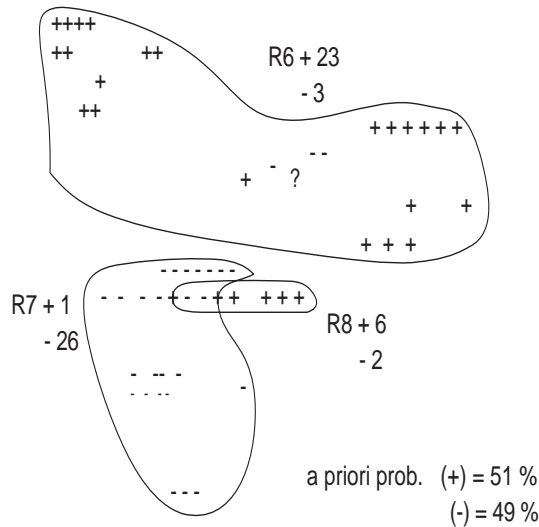


Fig. 5. Three new rules from a second induction round.

ports various rule induction techniques, e.g., separate-and-conquer, divide-and-conquer, and ensemble methods like bagging and boosting.

Double Induction has been implemented both with and without weighting of the examples according to the coverage of previous learned rules. Both of these variants are compared to frequency-based classification and naïve Bayes.

3.1 Experimental Setting

The employed search strategy was Separate-and-Conquer together with information gain to greedily choose what condition to add when refining a rule. One half of the training set was used as grow set and the second half was used as pruning set. The rules were pruned using incremental reduced error pruning (IREP) [6]. These settings were used in both induction rounds, i.e. both when inducing the initial set of rules and when resolving rule conflicts.

The experiments were performed using 10 fold cross-validation. In N fold cross-validation the data is partitioned in N disjunctive folds, running N separate experiments on each domain. Using N-1 out of the N folds to train on and the last fold to test on. Hence all N sets will be used as test set once and all sets will be used for training N-1 times.

All datasets used were taken from the UCI Machine Learning Repository except the King-Rook-King-Illegal (KRKI) database which comes from the Machine Learning group at the University of York. In Table 1, the domains used in the experiment are shown, as well as their main characteristics.

Table 1. The domains used in the experiment

| Domain | Classes | Class distribution | Examples |
|-------------------------|---------|--------------------------------------------------------------------------------------------------------------------|----------|
| The Glass | 2 | 24.1, 75.9 | 112 |
| Balance Scale | 3 | 8, 46, 46 | 625 |
| Breast Cancer | 2 | 29.7, 70.3 | 286 |
| Liver-disorders | 2 | 42, 58 | 345 |
| Car Evaluation | 4 | 4, 4, 22, 70 | 1728 |
| Dermatology | 6 | 5.5, 13.3, 14.2, 16.7, 19.7, 30.6 | 366 |
| Congressional Voting | 2 | 45.2, 54.8 | 435 |
| Ionosphere | 2 | 36, 64 | 351 |
| Lymphography | 4 | 1.4, 2.7, 41.2, 54.7 | 148 |
| New thyroid | 3 | 69.8, 16.3, 13.9 | 215 |
| Primary tumor | 22 | 24.8, 5.9, 2.7, 4.1, 11.5, 0.3, 4.1, 1.8, 0, 0.6, 8.3, 4.7, 2.1, 7.1, 0.6, 0.3, 2.9, 8.6, 1.8, 0.6, 0.3, 7.1 | 339 |
| Sonar | 2 | 53, 47 | 208 |
| Nursery | 5 | 33.3, 0.0, 2.5, 32.9, 31.2 | 12960 |
| Shuttle Landing Control | 2 | 47.8, 52.2 | 278 |
| KRKI | 2 | 34, 66 | 1000 |

3.2 Experimental Results

The rule conflict resolution methods were tested on fifteen domains. In four of these domains, the rules learned were without conflict. These domains were: Glass, Liver-disorders, Ionosphere and Sonar.

The results from the domains with rule conflicts are shown in Table 2, where the result for each domain has been obtained by ten-fold cross-validation. Exactly the same folds and generated rules are used by the four classification methods.

In Table 2, the first column gives the name of the domain. The second column shows the accuracy of Double Induction with weighted examples, while the third column gives the accuracy of Double Induction with no weighting. The fourth column shows the accuracy of frequency-based classification and the fifth column shows the accuracy of naïve Bayes. The sixth column shows the percentage of all classifications that involve conflicting rules. This gives an upper-bound of how much the accuracy can be improved.

Table 3 gives numbers that are used for the significance test. Inside each parenthesis, two numbers are given: the number of domains in which the method in the leftmost column has a higher respectively lower accuracy compared to the method in the upper row (wins, losses). The value after each parenthesis shows the p-value according to an exact version of McNemar’s test. An asterisk (*) is used to signal that the result is statistically significant, using the threshold $p < 0.05$.

Table 2. Accuracy of Double Induction (with and without weights), frequency based classification and naïve Bayes.

| Domain | D. Ind. w weight | D. Ind. | frequency-b. | naïve B. | conflicts |
|---------------|------------------|---------|--------------|----------|-----------|
| Balance scale | 83.30 | 82.43 | 81.20 | 82.78 | 50.8 |
| Breast cancer | 75.00 | 74.60 | 73.81 | 75.00 | 21.4 |
| Car | 86.82 | 87.21 | 78.04 | 77.59 | 34.1 |
| Dermatology | 91.69 | 91.69 | 91.69 | 91.08 | 5.2 |
| C. Votes | 95.20 | 94.95 | 94.44 | 93.94 | 6.6 |
| Lymphography | 80.00 | 80.00 | 79.26 | 80.00 | 26 |
| New thyroid | 85.71 | 87.24 | 85.71 | 85.71 | 5.6 |
| Primary tumor | 38.33 | 38.33 | 37.50 | 37.50 | 49.2 |
| Nursery | 85.63 | 85.42 | 79.91 | 84.32 | 26.0 |
| Shuttle | 94.86 | 94.07 | 93.68 | 94.86 | 7.5 |
| KRKI | 98.35 | 98.24 | 92.20 | 98.02 | 11.2 |

Table 3. Result of McNemar’s test.

| | D. Ind. w. weight | D. Induction | frequency-b. | naïve Bayes |
|-------------------|-------------------|--------------------|--------------------|------------------|
| D. Ind. w. weight | (0, 0), 1 | (6, 2), 0.289 | (9, 0), 3.91e-3* | (7, 0), 1.56e-2* |
| D. Induction | (2, 6), 0.289 | (0, 0), 1 | (10, 0), 1.953e-3* | (7, 3), 0.344 |
| frequency-b. | (0, 9), 3.91e-3* | (0, 10), 1.953e-3* | (0, 0), 1 | (3, 7), 0.344 |
| naïve Bayes | (0, 7), 1.56e-2* | (3, 7), 0.344 | (7, 3), 0.344 | (0, 0), 1 |

To minimise the amount of work within Double Induction, the conflicts already seen (for a particular fold) are saved to allow re-use of the generated rules in the second round whenever the same conflict is observed again (within the same fold). This can greatly speed up the classification process. For the current domains, this reuse of classifications varies from 2 percent to 60 percent of all classifications with conflicts.

4 Discussion and related work

It is worth noting that Double Induction with weighting is significantly more accurate than both naïve Bayes and frequency-based classification. It is not significantly more accurate than Double Induction without weighting, but is still better (6 wins and 2 losses).

Double Induction without weighting performs significantly better than frequency-based classification, but not significantly better than naïve Bayes. There is no significant difference in accuracy between naïve Bayes and frequency-based classification.

4.1 Rule stretching

Rule stretching (RS)[5] addresses the problem of how to classify examples which is not covered by any rule. The usual approach to this problem is to assign

the majority class of the domain to the uncovered example(s). RS deals with this problem by “stretching” the rules, i.e. relaxing the conditions in the rules (as little as possible), so that they cover the uncovered example that is to be classified. Then the class coverage distributions for each rule is updated and the rule with the highest class coverage distribution is used to classify the example. Results from empirical comparison with this method compared to using the majority class showed that RS method outperformed the use of the majority class in all domains. The methods were compared in total on seven domains.

In a later article [4] RS was used for classification of all test examples in a domain. In a comparison between RS and using the most accurate rule, RS proved to perform better than using the most accurate rule in six out of seven domains.

RS and Double Induction (DI) have similarities. Both use existing rules as a basis for further investigation. DI only considers the coverage of the rule while RS considers the rules conditions. But they aim at solving different problems. Both methods are computationally more expensive than their standard counterpart but there is a pay back of all the computing in the form of better performance.

4.2 Analogical Prediction

Analogical Prediction (AP) [11] is a general method for inducing rules that are tailor made for the prediction of a specific example. The method can be regarded as a step stone between standard Inductive Logic Programming (ILP) [12] learning setting and that of Instance-Based Learning (IBL). In ILP rules are induced from the training examples possibly together with background knowledge. In IBL no rules are created, instead are the training examples used with a proximity measure to classify test examples. AP uses the training examples together with background knowledge and for each test example induce the most compressive rule that is best for classifying that particular example. The positive aspects of AP compared to IBL is that it uses rules which are easy to interpret, while IBL use a proximity measure which can be hard to interpret. And it has the same nice properties that IBL has but ILP lacks namely that it is easy to update the knowledge base (examples).

When comparing DI with AP they are quite similar, both methods induce rules when a test example is to be classified. But DI only considers the examples covered by the rules in conflict when inducing new rules while AP considers all the training data. Also AP is restricted to find one rule such restrictions are not applied to DI. The aim of both methods is quite similar, to find rules which best predict a particular example.

4.3 Instance-Based Learning

As mentioned above IBL [10] uses a proximity measure which is applied to the training examples for each example that is to be classified. In a sense DI is an IBL type of algorithm, because when it induces new rules it only considers a subset of the original training set (the examples covered by the rules in conflict). So

DI excludes examples that is not central to solving the rule conflict, similarly to IBL which only considers training examples that are in the proximity of the example to be classified.

4.4 Intersecting Rules

Intersecting rules (IR) [9] also considers solving the same problem as Double Induction. The difference is that IR tries to use bayes rule if it can, i.e. use the probability distribution in the intersection of the rules in conflict together with the a priori probability of a domain to classify an example. The problem is that often there are not any examples present in the intersection of the conflicting rules. In such a case IR search for a partition of the rules in so few nonempty-partitions as possible, i.e. it tries to preserve the dependencies between the rules as long as possible. This has show to be successful when solving rule conflicts, in 11 domains IR has higher accuracy in all domains compared to frequency-based classification and naïve Bayes.

DI and IR both aim at solving the same problem while IR is based on bayes rule and uses the original induced rules to predict the class of an example DI solves the same problem by inducing new rules. It should be noted that the weighting scheme that DI uses comes from the empirical evidence that examples in the intersection of conflicting rules is more important that rules which lie outside the intersection. This is one of the conclusions which can be drawn from IR experiments. This notion is further supported in the experiments done in this paper which show better performance when using DI with weighting than without it.

5 Conclusion

Instead of using the information conveyed in a set of conflicting rules (and the class distributions of the whole training set, if using naïve Bayes) and making a qualified guess of which class is most probable given the rules, we suggest doing a second round of rule induction in hope of finding a better separation of the examples. This is often possible due to that the induction algorithm only has to consider the examples that are relevant to the region in which the example to be classified resides. This idea has been empirically demonstrated to significantly outperform two previous methods for resolving rule conflicts: frequency-based classification and naïve Bayes.

Furthermore, as seen by the experiments, it is worthwhile to use the weighting of the examples that is given by the first set of rules, because it helps in the second round to produce a better set of rules. This confirms the earlier observation in [9], that examples in the intersection of conflicting rules are more important than other examples.

One of the major drawbacks of using Double Induction is of course the computational cost. However, if accuracy is of uttermost importance and quick response time is not, then it is a useful technique.

One issue that needs further investigation is the use of other weighting schemes than the one used in our experiments which is quite conservative. Another issue is to resolve new conflicts that occur between rules obtained in the second round, by continuing recursively instead of using naïve Bayes after the second round. This recursive step could be done until no further conflicts are present in the resulting rules or until no new rules can be found. A more delicate measure for rule conflicts could prove to be useful. As the case is now there exists no measurement of the degree of the conflict. For example comparing two different conflicts; in the first conflict 9 rules advocate the pos class and 1 rule the neg class. All 9 rules that advocate the pos class are “good” rules (they cover a lot of examples and are quite clean (few neg examples is covered)) while the single rule advocate the neg class is a “bad” rule (cover few examples, and just barley have a majority class of neg covered examples). The second conflict involves two rules that both are good but have different majority classes. Clearly the second conflict is a “harder” conflict than the first one, this should be captured by a conflict measure.

References

1. Henrik Boström. Rule discovery system user manual. *Compumine AB*, 2003.
2. P. Clark and R. Boswell. Rule induction with CN2: Some recent improvements. In *Proceedings Fifth European Working Session on Learning*, pages 151–163, Berlin, 1991. Springer-Verlag.
3. P. Clark and T. Niblett. The CN2 induction algorithm. *Machine Learning*, 3, 261–283, 1989.
4. M. Eineborg. Fuzzifying hyperplanes in the hypothesis space. In *Proceedings of the First International Workshop on Hybrid Intelligent Systems*, pages 313–322. Springer-Verlag, 2001.
5. M. Eineborg and H. Boström. Classifying uncovered examples using rule stretching. In *Proceedings of the Eleventh International Conference on Inductive Logic Programming*, pages 41–50. Springer-Verlag, 2001.
6. J. Fürnkranz and G. Widmer. Incremental reduced error pruning. In W.W. Cohen and H. Hirsh, editors, *Proceedings of the 11th International Conference on Machine Learning*, pages 70–77. Morgan Kaufmann, 1994.
7. Johannes Fürnkranz. Separate-and-conquer rule learning. *Artificial Intelligence Review*, 1999.
8. R. Kohavi, B. Becker, and D. Sommerfield. Improving simple bayes. In *Proceedings of the European Conference on Machine Learning*, 1997.
9. Tony Lindgren and Henrik Boström. Classification with intersecting rules. In *Proceedings of the 13th International Conference on Algorithmic Learning Theory (ALT'02)*, pages 395–402. Springer-Verlag, 2002.
10. Tom M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.
11. Stephen Muggleton and Michael Bain. Analogical prediction. In *Proceedings of the 9th International Conference on Inductive Logic Programming*, pages 234–244. Springer-Verlag, 1999.
12. Stephen Muggleton and Luc de Raedt. Inductive logic programming: Theory and methods. *The Journal of Logic Programming*, 1994.
13. Duda R. O. and Hart P. E. *Pattern Classification and Scene Analysis*. Wiley, 1973.

14. Giulia Pagallo and David Haussler. Boolean feature discovery in empirical learning. *Machine Learning*, 5(1):71–99, 1990.
15. J. R. Quinlan. Induction of decision trees. *Machine Learning*, 1:81–106, 1986.
16. R. Rivest. Learning decision lists. *Machine Learning*, 2(3), 229-246, 1987.