Attribute dependencies, understandability and split selection in tree based models

Marko Robnik - Šikonja and Igor Kononenko
University of Ljubljana,
Faculty of Computer and Information Science,
Tržaška 25, 1001 Ljubljana, Slovenia,
Marko.Robnik@fri.uni-lj.si Igor.Kononenko@fri.uni-lj.si

Abstract

The attributes' interdependencies have strong effect on understandability of tree based models. If strong dependencies between the attributes are not recognized and these attributes are not used as splits near the root of the tree this causes node replications in lower levels of the tree, blurs the description of dependencies and also might cause drop of accuracy. If Relief family of algorithms which is capable of estimating the attributes' dependencies is used for split selectors we can partly overcome the problem. However, typically we still want to optimize accuracy of the tree and therefore use accuracy as the split selector measure near the fringe of the tree. We present a technique which helps us select a split criterion during tree growing based on some theoretical properties of Relief's estimate. We support our claims with empirical results.

1 Introduction

The problem of estimating the quality of attributes (features) is an important issue in machine learning and has received much attention in the literature. There are several measures for estimating the attribute's quality. If the target concept is a discrete variable (classification problem) these are e.g., information gain (Hunt et al., 1966), Gini index (Breiman et al., 1984), distance measure (Mantaras, 1989), j-measure (Smyth and Goodman, 1990), Relief (Kira and Rendell, 1992), ReliefF (Kononenko, 1994), MDL (Kononenko, 1995), Contextual Merit (Hong, 1997), and also χ^2 and G statistics are used. If the target concept is presented as a real valued function (continuous class, regression problem) then the estimation heuristics are e.g., the mean squared and the mean absolute error (Breiman et al., 1984) and RReliefF (Robnik Šikonja and Kononenko, 1997). There are several important tasks in the process of machine learning, e.g., feature selection, constructive induction, decision and regression tree building, which contain attribute estimation procedure as their (crucial) ingredient.

Most of the heuristic measures for estimating the quality of the attributes assume the independence of the attributes and indeed the experimental results (Kononenko et al., 1997) indicate that this myopia has no or only marginal effect on the majority of real world problems. However when we face a new problem it is unreasonable to try only these restricted measures unless we know in advance that there are no strong conditional dependencies between the attributes. Relief algorithms (Relief, ReliefF and RReliefF) do not make this assumption and are non-myopic in this sense. They are aware of the contextual information, efficient and can correctly estimate the quality of attributes in problems with strong dependencies between the attributes as well. While Relief algorithms have commonly been viewed as a feature selection methods that are applied in a preprocessing step before the model is learned (Kira and Rendell, 1992) and are one of the most successful preprocessing algorithms to date (Wettschereck et al., 1997; Dietterich, 1997), they are actually general feature estimators and have been used successfully in a variety of settings: to select splits in the building phase of decision tree learning (Kononenko et al., 1997), to select splits and guide the constructive induction in learning of the regression trees (Robnik Šikonja and Kononenko, 1997), as attribute weighting method in instance based learning (Wettschereck et al., 1997), and also in inductive logic programming (Pompe and Kononenko, 1995).

Decision and regression trees are popular description languages for representing knowledge in machine learning. While constructing a tree, a learning algorithm at each interior node selects a splitting rule (a feature) which divides

the problem space into two separate subspaces. To select the appropriate splitting rule the learning algorithm has to evaluate several possibilities and decide which would partition the given (sub)problem most appropriately. The estimation of the quality of splitting rules seems to be of principal importance.

The focus of this paper is on the effect the attributes' interdependencies have on the understandability of decision and regression trees. Impurity based measures try to select as pure splits as possible i.e., separate instances with different class values into different subtrees. Strong dependencies between the attributes are not properly detected by these measures, therefore dependent attributes are not selected as splits near the root of the tree which would guarantee compact representation of dependencies. Instead they are selected in subsequent levels of the tree, mostly near the fringe (if at all) which causes node replication, blurs the description of dependencies and also might cause drop of accuracy.

One possible solution to this is to use multivariate splits (Brodley and Utgoff, 1995) but this can be computationally expensive. Another solution to this problem is to use non-myopic attribute estimator as split selection heuristics. ReliefF was used for that purpose within a decision trees learning system (Kononenko et al., 1997) and achieved significantly better results in terms of accuracy and complexity of decision trees in domains with highly dependent features compared to impurity based measures. It achieved the same performance in domains with independent features. RReliefF was used as split selector in regression tree learning system and also achieved better results in terms of error and complexity of the tree compared to the mean squared error (Robnik Šikonja and Kononenko, 1997). The tree structure chosen by Relief algorithms is closer to human partition of the problem and improves on the understandability and extracted knowledge (Kononenko et al., 1997; Dalaka et al., 1999).

Additional difficulty lies in the fact that if we want to improve accuracy we have to take into consideration that test nodes closer to the leaves of the tree should be chosen to maximize accuracy on the training set. This effect was investigated by (Brodley, 1995; Lubinsky, 1995) and showed empirically that this can improve classification accuracy, however the authors were working with myopic split selection measures and could not detect and address the effect the attribute dependencies have on the structure of the tree.

In this paper we are trying to look at both problems from the same point of view and to suggest some heuristic measures which would help us to automate the selection of the appropriate split.

In the next Section we describe the problem of attribute dependencies, explain how Relief algorithms detect these dependencies and establish relation between Relief's estimations and myopic measures. Section 3 describes how these relation can be used to determine when to switch from Relief to accuracy as split selection measure, and Section 4 describes empirical evaluation of our method. The last Section summarizes and gives guidelines for further work.

2 Attribute dependencies

When we are talking about attribute dependencies in this paper we have in mind the conditional dependencies (conditioned upon class values). Let us illustrate by three examples defined in Boolean space.

1. Let class value be defined as XOR (parity) of two attributes $C = X_1 \oplus X_2$. Knowing the values of the class and one of the attribute we have all the information about the other attribute, while the knowledge of the class value alone does not contain any information about the attributes' values. In this case we can say that X_1 and X_2 are (maximally) conditionally dependent.

2. Let class value be defined as conjunction of two attributes $C = X_1 \wedge X_2$. The knowledge about the class and one of the attributes is sometimes enough to determine the value of the other attribute (e.g., if $C = 0$ and $X_1 = 1$, the value of X_2 must be 0) but not always (e.g., if $C = 0$ and $X_1 = 0$, the value of X_2 is still undetermined), while the class value alone contains even less information (e.g., if $C = 0$ we cannot determine the values of X_1 and X_2). We can say that X_1 and X_2 are somehow conditionally dependent.

3. Let class be an arbitrary non-constant Boolean function, both values of the attributes are equiprobable ($P(X_1 = 1) = P(X_2 = 1) = 0.5$) and attributes are defined as $P(X_1 = 1|C = 1) = p$, $P(X_1 = 1|C = 0) = 1 - p$ and $P(X_2 = 1|C = 1) = q$, $P(X_2 = 1|C = 0) = 1 - q$. Each attribute contains some knowledge about the class value, but they contain this knowledge independently from each other. The knowledge of the class value conveys certain information about values of each of attribute, but this information does not increase even if we know the value of one of the attributes. So in this case we can say that X_1 and X_2 are conditionally independent. Another example of conditional independence are random attributes.

Relief algorithms estimate the quality of attributes in problems with strong dependencies between the attributes. How do they do that? The basic Relief algorithm (Kira and Rendell, 1992) is presented on Figure 1 and shortly described below.

The key idea is to estimate the quality of attributes according to how well their values distinguish between the instances that are near to each other. For that purpose, given a randomly selected instance R (line 3), Relief searches for its two nearest neighbors: one from the same class, called *nearest hit* H, and the other from a different class, called *nearest miss* M (line 4). It updates the quality estimation $W[A]$ for all the attributes A depending on their values for R, M, and H (lines 5 and 6). If the values of attribute A at instances R and H are different then attribute A separates two instances of the same class which is not desirable so we add negative update to the quality estimation $W[A]$. If the values of attribute A at instances R and M are different then attribute A separates two instances with different class values which is desirable so we add positive update to the quality estimation $W[A]$. The process is repeated for m times, where m is a user-defined parameter. Function $diff(Attribute, Instance1, Instance2)$ calculates the difference between the values of Attribute for two instances. For discrete attributes it was originally defined as:

$$diff(A, I_1, I_2) = \begin{cases} 0 & ; value(A, I_1) = value(A, I_2) \\ 1 & ; otherwise \end{cases} \quad (1)$$

and for continuous attributes as:

$$diff(A, I_1, I_2) = \frac{|value(A, I_1) - value(A, I_2)|}{max(A) - min(A)} \quad (2)$$

The function $diff$ is used also for calculating the distance between instances to find the nearest neighbors. The total distance is the sum of distances over all attributes.

Relief's estimate $W[A]$ of the quality of attribute A is an approximation of the following difference of probabilities (Kononenko, 1994):

$$\begin{aligned} W[A] = & \\ & P(\text{diff. value of A}|\text{nearest inst. from diff. class}) \\ & -P(\text{diff. value of A}|\text{nearest inst. from same class}) \end{aligned} \quad (3)$$

The original Relief can deal with discrete and continuous attributes. However, it can not deal with incomplete data and is limited to two-class problems. Its extension which solves these and some other problems is called ReliefF (Kononenko, 1994). One of the differences interesting for our paper is that ReliefF uses several nearest neighbors in the approximation of probabilities. The adaptation of ReliefF for regressional problems (continuous class) is called RReliefF (Robnik Šikonja and Kononenko, 1997).

The estimations of Relief algorithms are strongly related to impurity functions (Kononenko, 1994). As this relation is the basis for this paper we will repeat it below. When the number of nearest neighbors increases i.e., when we eliminate the requirement that the selected instance is the nearest, the Equation (3) becomes

$$\begin{aligned} W'[A] &= P(\text{different value of A}|\text{different class}) \\ &- P(\text{different value of A}|\text{same class}) \end{aligned} \quad (4)$$

If we rewrite

$$P_{eqval} = P(\text{equal value of A})$$

$$P_{samecl} = P(\text{same class})$$

$$P_{samecl|eqval} = P(\text{same class}|\text{equal value of A})$$

we obtain using Bayes' rule:

$$W'[A] = \frac{P_{samecl|eqval} P_{eqval}}{P_{samecl}} - \frac{(1 - P_{samecl|eqval}) P_{eqval}}{1 - P_{samecl}}$$

For sampling with replacement in strict sense the following equalities hold:

$$P_{samecl} = \sum_C P(C)^2$$

$$P_{samecl|eqval} = \sum_V \left(\frac{P(V)^2}{\sum_V P(V)^2} \times \sum_C P(C|V)^2 \right)$$

Using the above equalities we obtain:

$$W'[A] = \frac{P_{eqval} \times Ginigain'(A)}{P_{samecl}(1 - P_{samecl})} \quad (5)$$

where

$$\begin{aligned} Ginigain'(A) &= \sum_V \left(\frac{P(V)^2}{\sum_V P(V)^2} \times \sum_C P(C|V)^2 \right) \\ &- \sum_C P(C)^2 \end{aligned} \quad (6)$$

is highly correlated with the Gini-index gain (Breiman et al., 1984) for classes C and values V of attribute A. The difference is that instead of factor

$$\frac{P(V)^2}{\sum_V P(V)^2}$$

Algorithm Relief
Input: for each training instance a vector of attribute values and the class value
Output: the vector W of estimations of the qualities of attributes

1. set all weights $W[A] := 0.0$;
2. **for** i := 1 **to** m **do begin**
3. randomly select an instance R;
4. find nearest hit H and nearest miss M;
5. **for** $A := 1$ **to** #all_attributes **do**
6. $W[A] := W[A] - diff(\text{A,R,H})/m + diff(A, R, M)/m$;
7. **end**;

Figure 1: Pseudo code of the basic Relief algorithm

the Gini-index gain uses

$$\frac{P(V)}{\sum_V P(V)} = P(V)$$

Equation (5) (which we call myopic ReliefF), shows strong correlation of Relief's weights with the Gini-index gain. The probability $\sum_V P(V)^2$ that two instances have the same value of attribute A in Eq. (5) is a kind of normalization factor for multi-valued attributes. Impurity functions tend to overestimate multi-valued attributes and various normalization heuristics are needed to avoid this tendency (e.g. gain ratio (Quinlan, 1986), distance measure (Mantaras, 1989), and binarization of attributes (Cestnik et al., 1987)). Equation (5) shows that Relief exhibits an implicit normalization effect. Another deficiency of Gini-index gain is that its values tend to decrease with the increasing number of classes. Denominator which is constant factor in Equation (5) for a given attribute again serves as a kind of normalization and therefore Relief's estimates do not exhibit such strange behavior as Gini-index gain does. This normalization effect remains even if Equation (5) is used as (myopic) attribute estimator what was empirically confirmed in (Kononenko, 1995).

The above derivation eliminated the "nearest instance" condition from the probabilities. If we put it back we can interpret Relief's estimates as the average over local estimates in smaller parts of the instance space. This enables Relief to take into account the context of other attributes, i.e. the conditional dependencies between attributes given the class value which can be detected in the context of locality. From the global point of view, these dependencies are hidden due to the effect of averaging over all training instances, and exactly this makes impurity functions myopic. Impurity functions use correlation between the attribute and the class disregarding the context of other attributes. This is the same as using the global point of view and disregarding the local peculiarities. We illustrate this in Figure 2 which shows dependency of ReliefF's estimate to number of nearest neighbors taken into account. The estimates are for a parity domain with two informative and 10 random attributes and 200 examples. The dotted line shows how the ReliefF's estimate of one of informative attributes is becoming more and more myopic with increasing number of nearest neighbors and how it eventually becomes indistinguishable from random attribute (Note that it is indistinguishable for all myopic estimators).

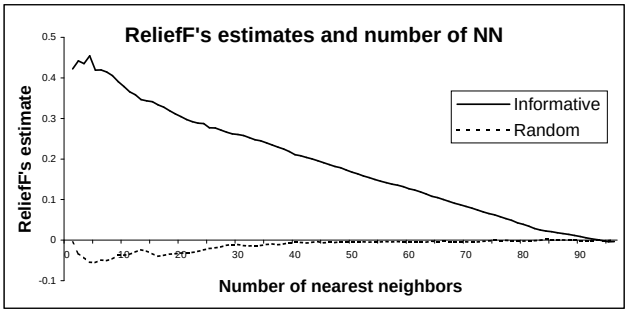

Figure 2: ReliefF's estimates of informative attribute are deteriorating with increasing number of nearest neighbors in parity domain.

3 Global vs. local view, switching estimators

In the process of building a (decision, regression) tree using Relief as the attribute selection heuristics we use all important dependent attributes in the upper part of the tree and come to a point where there are no more dependencies, i.e., the local view becomes the global view. This is the point when there are no more hidden dependencies and we should start optimizing the tree for accuracy. How can we detect this point? We argue that it can be detected by measuring the correlation between the global estimate (Equation (5)) and the average over local estimates (Re-

lief's weights, Equation (3)) for various attributes.

The (linear) correlation coefficient $\rho_{x,y}$ between two vectors of observations x and y is defined as

$$\rho_{x,y} = \frac{\frac{1}{N}\sum_{i=1}^{N}(x_i - \mu_x)(y_i - \mu_y)}{\sigma_x \sigma_y} \quad (7)$$

where $-1 \leq \rho_{x,y} \leq 1$. This is a well known statistical measure which measures the degree of linear relationship between the x_i and y_i values. In our case x_i is the ReliefF's estimate of i-th attribute and y_i is the estimate of the i-th attribute by another estimator.

Accuracy as attribute estimator has some major drawbacks (Brodley, 1995) in terms of the structure of the tree but once we have extracted the dependencies it can be very useful if we want to optimize the tree for accuracy. If we classify with the majority class the accuracy as the attribute estimator it is defined as

$$Acc = \sum_{V} P(V) \max_{C} P(C|V) \quad (8)$$

If we want to compute the accuracy for a binary split of a multivalued attribute we group the attribute's values into two groups (left and right) and use the group probabilities in Equation 8.

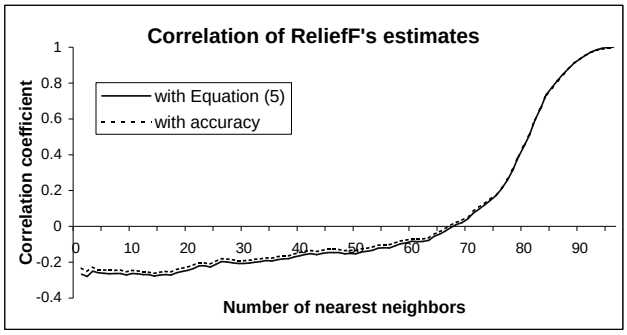

Figure 3: Correlation between ReliefF and Myopic Relief and ReliefF and accuracy in parity domain.

Figure 3 shows the correlation coefficient between the ReliefF's estimate and Equation (5)(solid line) and ReliefF's estimate and accuracy (dotted line) varying over the number of nearest neighbors taken into account for ReliefF's estimates. As before the estimates are for a parity domain with two informative and 10 random attributes and 200 examples. We see that the correlation coefficients are very similar and as it is indicative in other domains as well this leads to further simplification of our heuristic for selecting split criteria: by monitoring the correlation coefficient between ReliefF's estimate and accuracy we determine when there are no more dependencies in a local subspace (tree node). When the correlation coefficient is high enough (e.g., ≥ 0.8) we switch to accuracy optimization and select Equation (8) as the split selector measure in the subtree of a current node.

4 Empirical evaluation

Some recent studies e.g., (Dietterich, 1998; Rassmusen, 1996), have pointed out that the assumption of independence of training sets is violated in k-fold cross-validated paired Student's t-test, however (Dietterich, 1998) observes that this test is very powerful and has only somewhat elevated probability of Type I error. As most of the datasets we are using contain relatively small number of instances and 10-fold cross-validation is economical with the data we have selected 10-fold cross-validated paired Student's t-test for computing the significance of the differences. The significance of the differences were computed at 0.05 level. As we are comparing over several datasets we are also using Bonferroni adjustment (see e.g., (Salzberg, 1997)).

4.1 Decision trees

We have developed a learning system which builds binary decision trees by splitting training examples based on the values of attributes. The attribute in each node is selected according to the estimates of the attributes' quality. These estimates are computed on the subset of the examples that reach current node.

We run our system with three different estimators: ReliefF alone, accuracy alone, and combined ReliefF and accuracy (depending on the threshold for the correlation coefficient, which was set to 0.8). All other parameters for growing and pruning the tree were the same and set to their default values. We used two stopping criteria, namely the minimal number of cases in the leaf (=5) or the minimal purity of a leaf (97% of cases are from the same class) and postpruned the tree with the m-estimate ($m = 2$).

The datasets used are mostly from UCI (Murphy and Aha, 1995), those which are not are shortly described below. Their characteristics are described in Table 1.

Parity-2: Parity problem with two significant binary attributes and 10 random binary attributes. 5% of randomly selected instances were labeled with wrong class.

Parity-3: Same as Parity-2 except that there were three significant attributes for the parity relation which makes the problem harder.

Table 1: Some numerical characteristics of categorical datasets used in our experiments (number of classes, number of all, discrete and continuous attributes, number of instances, the percentage of majority class and the class entropy (in bits)).

domain	#class	#atts.	#disc.	#cont.	#inst.	maj.class (%)	entropy
Credit	2	15	9	6	690	56	0.85
Diabetes	2	8	0	8	768	65	0.93
Glass	7	9	0	9	214	36	2.18
Heart	2	13	6	7	270	56	0.99
Iris	3	4	0	4	150	33	1.59
KRK1	2	18	18	0	1000	67	0.92
KRK2	2	6	0	6	1000	67	0.92
LED	10	7	7	0	1000	11	3.33
Lymphography	4	18	18	0	148	55	1.28
Parity-2	2	12	12	0	200	54	0.99
Parity-3	2	13	13	0	200	54	0.99
Primary	22	17	17	0	339	25	3.89
Promoter	2	35	35	0	106	59	1.00
Rheumatic	6	32	11	21	355	66	1.73
Soya	15	35	35	0	630	15	3.62
Tic-tac-toe	2	9	9	0	958	65	0.93
Vote	2	16	16	0	435	61	0.96

KRK1: The problem of legality of King-Rook-King chess endgame positions. The attributes describe the relevant relations between pieces, such as "same_rank" and "adjacent_file".

KRK2: Same as KRK1 except that the only available attributes are the coordinates of pieces.

For each domain we collected the results as the average of 10 fold cross-validation. We present results in Table 2. Besides the accuracy we included also the number of leaves in the tree as a measure of complexity of the tree.

There are 7 datasets where ReliefF is significantly more accurate than accuracy (KRK1, KRK2, Parity-3, Primary, Rheumatic, Soya, Tic-tac-toe). The combined approach was significantly better than accuracy alone in 6 datasets (KRK2, Parity-3, Primary, Rheumatic, Soya, Tic-tac-toe). There were no significant differences in accuracy between ReliefF and combined approach. Accuracy was never significantly better than the other two methods. The significance of the differences are summarized in Table 3.

Table 3: Number of significant differences in favor of the estimators in the first column over the estimators in the first row.

estimator	ReliefF	accuracy	combination
ReliefF	-	7	0
accuracy	0	-	0
combination	0	6	-

The combined approach produced the trees which are on average slightly larger than those produced solely by ReliefF and slightly smaller than those produced solely by the accuracy.

Visual inspection of the trees confirmed that mostly the splitting criteria was switched when there were no more dependencies between the attributes. In average 37% of splits at the bottom of the trees were selected by the accuracy.

We also experimented with the correlation threshold and it turned out that the switching method is quite robust against this parameter. By setting this parameter in the interval from 0.7 to 0.95 we get very similar results. Above 0.95 there are very few nodes selected by accuracy and below 0.7 there are majority of nodes selected by the accuracy.

What can we conclude from the above experiments? Although in average 37% of splits at the bottom of the tree were selected by the accuracy this switch did not not help ReliefF in improving the score in accuracy. Comparing that with significant improvement which was observed when switching from gain ratio to accuracy (Brodley, 1995; Lubinsky, 1995) it seems that ReliefF does not need such switch as its self-normalization factors (see Equation (5)) and locality of estimates already take that into account.

To confirm this hypothesis we built this switching crite-

Table 2: Accuracy and complexity of the decision tree produced with RReliefF, accuracy and combination as the split selectors.

	ReliefF		accuracy		combination	
domain	#leaves	acc.	#leaves	acc.	#leaves	acc.
Credit	30	0.84	37	0.82	32	0.84
Diabetes	78	0.72	103	0.69	81	0.72
Glass	26	0.66	26	0.70	28	0.67
Heart	21	0.74	23	0.75	21	0.77
Iris	5	0.95	5	0.93	4	0.93
KRK1	11	0.99	6	0.91	11	0.92
KRK2	100	0.72	167	0.60	108	0.70
LED	31	0.73	31	0.71	31	0.71
Lymphography	14	0.74	12	0.72	14	0.72
Parity-2	4	0.96	16	0.91	4	0.96
Parity-3	8	0.95	12	0.58	8	0.95
Primary	48	0.42	6	0.28	45	0.42
Promoter	9	0.72	9	0.77	9	0.75
Rheumatic	53	0.65	55	0.56	53	0.65
Soya	37	0.92	46	0.87	38	0.92
Tic-tac-toe	43	0.94	29	0.82	41	0.94
Vote	6	0.96	2	0.96	6	0.96

rion into our regression tree learning system CORE (Robnik Šikonja, 1997) and test it with two different models in the leaves of the tree (mean prediction value and linear models). This extended definition of the accuracy puts Relief's abilities to a much harder test than in decision trees where the model in the leaf is a majority class.

4.2 Regression trees

Similarly to decision trees we have built the correlation based switching criterion into our regression tree learning system CORE (Robnik Šikonja, 1997) which uses regressional version of ReliefF called RReliefF (Robnik Šikonja and Kononenko, 1997) or Mean Squared Error (MSE) as the attribute estimator. In the leaves of the tree we used two different predictors to forecast the function value: mean prediction value of the training instances in the leaf or linear model which fits the instances in the leaf computed with singular value decomposition method (Press et al., 1988). MSE estimate of the attribute A is minimum over all possible splits of attribute A of the weighted mean squared error of the predictor ϕ:

$$MSE(\phi, A) = \min_{split\ of\ A} p_L \cdot R_{t_L}(\phi) + p_R \cdot R_{t_R}(\phi), \quad (9)$$

where t_L and t_R are the subsets of cases that go left and right, respectively, by the split based on A, and p_L and p_R are the proportions of cases that go left and right. $R_t(\phi)$ is the mean squared error of the predictor ϕ compared to real prediction values c_i of instances x_i in the subset t :

$$R_t(\phi) = \frac{1}{N_t} \sum_{i=1}^{N_t} (c_i - \phi(x_i))^2. \quad (10)$$

$\phi(x_i)$ is the value predicted by ϕ.

MSE is used as attribute estimator and as the success criterion (counterpart of accuracy in decision trees). Previous study (Robnik Šikonja and Kononenko, 1997) already compared RReliefF and MSE on regression trees and it turned out that on the problems which contain dependent attributes RReliefF produced smaller trees with significantly lower error, while on problems without strong dependencies RReliefF and MSE produce comparable trees.

We run our system with three different estimators: ReliefF alone, accuracy alone, and combined ReliefF and accuracy (with the correlation coefficient set to 0.8). All other parameters for growing and pruning the tree were the same and set to their default values. We used the minimal number of instances in the leaf (=5) as the stopping criterion and postpruned the tree with the m-estimate ($m = 2$).

We ran our system on the artificial data sets and on data sets with continuous prediction value from UCI (Murphy and Aha, 1995). Artificial data sets used were:

Fraction is floating point generalization of the parity con-

cepts of order $I = \{2, 3\}$. Each domain contains continuous attributes with values from 0 to 1. The predicted value is the fractional part of the sum of I important attributes: $C = \sum_{j=1}^{I} A_j - \lfloor \sum_{j=1}^{I} A_j \rfloor$

Modulo-8 is integer generalization of the parity of order $I = \{2, 3\}$. Value of each attribute is an integer value in the range 0-7. Half of the attributes are treated as discrete and half as continuous; each continuous attribute is exact match of one of the discrete attributes. The predicted value is the sum of the I important attributes by modulo 8: $C = (\sum_{j=1}^{I} A_j) \bmod 8$.

ParityC is continuously randomized parity of order $I = \{2, 3\}$. Each data set consists of discrete, Boolean attributes. The I informative attributes define parity concept: if their parity bit is 0, the predicted value is set to a random number between 0 and 0.5, otherwise the predicted value is randomly chosen to be between 0.5 and 1.

Linear presents simple linear dependency with 4 important attributes: $f = A_1 - 2A_2 + 3A_3 - 3A_4$. The attributes are continuous with values chosen randomly between 0 and 1.

Cosinus is non-linear dependency with cosine multiplied by the linear combination of two attributes: $f = \cos(4\pi A_1) \cdot (-2A_2 + 3A_3)$. The attributes are continuous with values from 0 to 1.

Altogether there were 8 artificial data sets, each consisting of 10 attributes - 2 or 3 important, the rest are random, and containing 1000 examples. UCI datasets used were:

Abalone: predicting the age of the abalone, 1 nominal and 7 continuous attributes, 4177 instances.

Autompg: city-cycle fuel consumption, 1 nominal, 6 continuous attributes 398 instances.

Autoprice: prices of the vehicles, 10 nominal , 15 continuous attributes, 201 instances.

Cpu: relative CPU performance, 6 continuous attributes, 209 instances.

Housing: housing values in Boston area, 1 nominal, 12 continuous attributes, 506 instances.

Servo: rise time of a servo mechanism, 2 nominal, 2 integer attributes, 167 instances.

Wisconsin: time to recur in Wisconsin breast cancer database, 32 continuous attributes, 198 instances.

For each experiment we collected the results as the average of 10 fold cross-validation. In Table 4 we give relative mean squared error:

$$RE_t(\phi) = \frac{R_t(\phi)}{R_t(\mu)} \tag{11}$$

ϕ is the evaluated predictor (regression tree) and μ is a predictor which always returns the mean of the training instances. Sensible predictors have $RE(\phi) < 1$. Besides the relative error we included also the number of leaves in the tree as a measure of complexity of the tree. The significance of the differences were computed by a two-sided paired Student's t-test at a 0.05 level using the Bonferroni adjustment and are presented in Table 5.

With linear models in the leaves there are 5 datasets where RReliefF has significantly lower error than MSE (Abalone, Fraction-3, Parity-2, Parity-3, and Wisconsin). The combined approach is significantly better than MSE in 4 datasets (Abalone, Fraction-3, Parity-2, and Parity-3). There are no significant differences in accuracy between RReliefF and combined approach. MSE never produces significantly lower error than the other two methods. With mean prediction value in the leaves we obtained similar results. Both RReliefF and the combined approach have significantly lower error than MSE in 4 datasets (Fraction-2, Fraction-3, Modulo8-2, and Parity-3)

Table 5: Number of significant differences in RE in favor of the estimators in the first column over the estimators in the first row.

estimator	RReliefF	MSE	combination
	linear models in leaves		
RReliefF	-	5	0
MSE	0	-	0
combination	0	4	-
	mean prediction value in leaves		
RReliefF	-	4	0
MSE	0	-	0
combination	0	4	-

When the switching criterion was used on average 25% and 36% of the splits at the bottom of the trees were selected by the MSE with linear models and mean prediction value, respectively.

We consider the above results as a confirmation of our hypothesis that (R)ReliefF does not need such a switch as it is already self-normalizing.

Table 4: RE and complexity of the regression trees produced with RReliefF, MSE and combination as the split selectors.

	linear models in leaves						mean prediction value in leaves					
	RReliefF		MSE		combination		RReliefF		MSE		combination	
domain	#leaf	RE	#leaf	RE	#leaf	RE	#leaf	RE	#leaf	RE	#leaf	RE
Abalone	133	0.62	170	0.80	141	0.66	186	0.58	166	0.56	191	0.60
Autompg	7	0.14	7	0.17	7	0.14	13	0.21	11	0.22	13	0.23
Autoprice	4	0.20	2	0.14	3	0.20	8	0.26	5	0.14	9	0.25
Cosinus	25	0.11	15	0.05	23	0.06	53	0.30	46	0.22	49	0.29
CPU	3	0.14	3	0.16	3	0.14	6	0.30	6	0.29	6	0.30
Fraction-2	22	0.35	20	0.52	21	0.58	40	0.29	107	1.12	48	0.49
Fraction-3	52	0.88	63	2.06	43	1.15	96	0.91	152	1.73	99	0.98
Housing	7	0.23	7	0.19	7	0.24	16	0.28	14	0.22	16	0.27
Linear	1	0.00	1	0.00	1	0.00	27	0.18	29	0.19	28	0.17
Modulo8-2	21	0.00	10	0.00	16	0.00	40	0.04	115	0.96	37	0.05
Modulo8-3	94	1.49	66	1.18	93	1.50	113	1.40	148	1.70	113	1.40
ParityC-2	4	0.27	48	0.35	4	0.27	4	0.27	48	0.35	4	0.27
ParityC-3	8	0.27	103	0.50	8	0.27	8	0.27	103	0.50	8	0.27
Servo	4	0.22	6	0.21	4	0.22	4	0.22	4	0.23	4	0.23
Wisconsin	14	1.35	10	2.06	16	1.52	29	1.51	25	1.85	32	1.76

5 Conclusions

We present a method for switching the split criterion from ReliefF to accuracy in the process of building the decision tree and from RReliefF to MSE in the process of building regression tree. Our method is based on theoretical properties of Relief estimations. It allows ReliefF (RReliefF) to recognize dependencies between the attributes and use them near the root of the tree which guarantees their compact representation and switch to accuracy (error) optimization when there are no more such dependencies.

Contrary to impurity based methods which produce significant lower error when combined with the error optimization near the fringe of the tree (Brodley, 1995; Lubinsky, 1995) it seems that ReliefF does not need this switch as its self-normalization factors in Equation (5) and global view combined with local estimations already contains it. However this hypothesis needs a further theoretical study.

Acknowledgments

The results in this paper are part of the first author's work on PhD thesis at University of Ljubljana, Faculty of Computer and Information Science. Both authors were supported by Slovenian Ministry of Science and Technology.

References

Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984). *Classification and regression trees*. Wadsworth Inc., Belmont, California.

Brodley, C. E. (1995). Automatic selection of split criterion during tree growing based on node location. In *Proceedings of the XII International Conference on Machine Learning*. Morgan Kaufmann.

Brodley, C. E. and Utgoff, P. E. (1995). Multivariate decision trees. *Machine Learning*, 19(1):45–77.

Cestnik, B., Kononenko, I., and Bratko, I. (1987). Assistant 86: A knowledge-elicitation tool for sophisticated users. In Bratko, I. and Lavrač, N., editors, *Progress in Machine Learning, Proceedings of EWSL 87*, Wilmslow. Sigma Press.

Dalaka, A., Kompare, B., Robnik Šikonja, M., and Sgardelis, S. (1999). Modeling by artificial intelligence tools the effects of environmental conditions on apparent photosynthesis of stipa bromoides. (submitted).

Dietterich, T. G. (1997). Machine learning research: Four current directions. *AI Magazine*, 18(4):97–136.

Dietterich, T. G. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10(7):1895–1924.

Hong, S. J. (1997). Use of contextual information for feature ranking and discretization. *IEEE transactions on knowledge and data engineering*, 9(5):718–730.

Hunt, E., Martin, J., and Stone, P. (1966). *Experiments in Induction*. Academic Press, New York.

Kira, K. and Rendell, L. A. (1992). A practical approach to feature selection. In D.Sleeman and P.Edwards, editors, *Proceedings of International Conference on Machine Learning*, pages 249–256. Morgan Kaufmann.

Kononenko, I. (1994). Estimating attributes: analysis and extensions of Relief. In De Raedt, L. and Bergadano, F., editors, *Machine Learning: ECML-94*, pages 171–182. Springer Verlag.

Kononenko, I. (1995). On biases in estimating multi-valued attributes. In *Proceedings of the International Joint Conference on Aartificial Intelligence IJCAI'95*, pages 1034–1040. Morgan Kaufmann.

Kononenko, I., Šimec, E., and Robnik-Šikonja, M. (1997). Overcoming the myopia of inductive learning algorithms with RELIEFF. *Applied Intelligence*, 7:39–55.

Lubinsky, D. J. (1995). Increasing the performance and consistency of classification trees by using the accuracy criterion at the leaves. In *Proceedings of the XII International Conference on Machine Learning*. Morgan Kaufmann.

Mantaras, R. (1989). ID3 revisited: A distance based criterion for attribute selection. In *Proceedings of Int. Symp. Methodologies for Intelligent Systems*, Charlotte, North Carolina, USA.

Murphy, P. and Aha, D. (1995). UCI repository of machine learning databases. (http://www.ics.uci.edu/ mlearn/-MLRepository.html).

Pompe, U. and Kononenko, I. (1995). Linear space induction in first order logic with ReliefF. In Della Riccia, G., Kruse, R., and Viertl, R., editors, *Mathematical and Statistical Methods in Artificial Intelligenc, CISM Courses and Lectures No.363*. Springer Verlag.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1988). *Numerical recipes in C*. Cambridge University Press.

Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1(1):81–106.

Rassmusen, C. E. (1996). *Evaluation of Gaussian processes and other methods for non-linear regression*. PhD thesis, University of Toronto, Department of Computer Science.

Robnik Šikonja, M. (1997). CORE - a system that predicts continuous variables. In *Proceedings of Electrotehnical and Computer Science Conference*, pages B145–148, Portorož, Slovenia.

Robnik Šikonja, M. and Kononenko, I. (1997). An adaptation of Relief for attribute estimation in regression. In Fisher, D., editor, *Machine Learning: Proceedings of the Fourteenth International Conference (ICML'97)*, pages 296–304. Morgan Kaufmann Publishers.

Salzberg, S. L. (1997). On comparing classifiers: Pitfalls to avoid and a recommended approach. *Data Mining and Knowledge Discovery*, 1:317–328.

Smyth, P. and Goodman, R. (1990). Rule induction using information theory. In Piatetsky-Shapiro, G. and Frawley, W., editors, *Knowledge Discovery in Databases*. MIT Press.

Wettschereck, D., Aha, D. W., and Mohri, T. (1997). A review and empirical evaluation of feature weighting methods for a class of lazy learning algorithms. *Artificial Intelligence Review*, 11:273–314.