



Munich Personal RePEc Archive

Credit scoring with boosted decision trees

Bastos, Joao

CEMAPRE, School of Economics and Management (ISEG),
Technical University of Lisbon

1 April 2008

Online at <https://mpra.ub.uni-muenchen.de/8156/>
MPRA Paper No. 8156, posted 08 Apr 2008 12:11 UTC

Credit scoring with boosted decision trees

João A. Bastos*

CEMAPRE, School of Economics and Management (ISEG)

Technical University of Lisbon, Portugal

Abstract

The enormous growth experienced by the credit industry has led researchers to develop sophisticated credit scoring models that help lenders decide whether to grant or reject credit to applicants. This paper proposes a credit scoring model based on boosted decision trees, a powerful learning technique that aggregates several decision trees to form a classifier given by a weighted majority vote of classifications predicted by individual decision trees. The performance of boosted decision trees is evaluated using two publicly available credit card application datasets. The prediction accuracy of boosted decision trees is benchmarked against two alternative data mining techniques: the multilayer perceptron and support vector machines. The results show that boosted decision trees are a competitive technique for implementing credit scoring models.

1 Introduction

The accurate assessment of consumer credit risk is of uttermost importance for lending organizations. Credit scoring is a widely used technique that helps financial institutions evaluate the likelihood for a credit applicant to default on the financial obligation and decide whether to grant credit or not. The precise judgment of the creditworthiness of applicants allows financial institutions to increase the volume of granted credit while minimizing possible losses. The credit industry has experienced a tremendous growth in the past few decades (Crook et al., 2007). The increased number of potential applicants impelled the development of sophisticated techniques that automate the credit approval procedure and supervise the financial health of the borrower. The large volume of loan portfolios also imply that modest improvements in scoring accuracy may result in significant savings for financial institutions (West, 2000).

The goal of a credit scoring model is to classify credit applicants into two classes: the “good credit” class that is liable to reimburse the financial obligation and the “bad credit” class that should be denied credit due to the high probability of defaulting on the financial obligation. The classification is contingent on sociodemographic characteristics of the borrower (such as age, education level, occupation and income), the repayment

*bastos@lipc.fis.uc.pt

performance on previous loans and the type of loan. These models are also applicable to small businesses since these may be regarded as extensions of an individual customer. In the last few decades, various quantitative methods were proposed in the literature to evaluate consumer loans and improve the credit scoring accuracy (for a review, see e.g. Crook et al., 2007). These models can be grouped into parametric and non-parametric or data mining models. The most popular parametric models are the linear discriminant analysis and the logistic regression. Linear discriminant analysis was the first parametric technique suggested for credit scoring purposes (Reichert et al., 1983). This approach has attracted criticism due to the categorical nature of the data and the fact that the covariance matrices of the good credit and bad credit groups are typically distinct. The logistic regression (Wiginton, 1980) allows to overcome these deficiencies and became a common credit scoring tool of practitioners in financial institutions. Non-parametric techniques applied to credit scoring include the k -nearest neighbor (Henley and Hand, 1996), decision trees (Frydman et al., 1985; Davis et al., 1992), artificial neural networks (Jensen, 1992), genetic programming (Ong et al., 2005) and support vector machines (Baesens et al., 2003). More recently, research on hybrid data mining approaches has shown promising results (Lee et al., 2002; Hsieh, 2005; Lee and Chen, 2002).

While the pursuit of better classifiers for credit scoring applications is a crucial research effort, improved accuracies can be easily achieved by aggregating scores predicted by an ensemble of individual classifiers. West et al. (2005) found that the accuracy of an ensemble of neural networks is superior to that of a single neural network in credit scoring and bankruptcy prediction applications. This paper proposes a credit scoring model of consumer loans based on boosted decision trees, a powerful learning technique in which an ensemble of decision trees is developed to form a classifier given by a weighted majority vote of classifications predicted by the individual trees. The decision trees are grown sequentially using reweighted training sets. If an instance is misclassified by a tree its weight is increased. Consequently, the predominance of “hard-to-classify” instances in the training sample increases with the number of grown trees. The performance of boosted decision trees is evaluated using two real world credit datasets from the UC Irvine Machine Learning Repository (Asuncion and Newman, 2007) and compared to that of a multilayer perceptron and a support vector machine. The rest of this paper is organized as follows. In the next section, boosted decision trees are introduced. This is followed by a description of the data sets and a comparison of the predictive accuracy of the models. A discussion of the relative contribution of the attributes to separate the good credit and bad credit classes is also given. Section 4 concludes the paper.

2 Boosted decision trees

2.1 Decision trees

Suppose one has a database of several credit applicants described by n attributes or characteristics: x_1, x_2, \dots, x_n . These applicants belong to two classes which will be denoted by “good credits” and “bad credits”. The goal of a credit scoring model is to find a classifier that separates the good credit sample from the bad credit sample. A decision tree consists of a set of sequential binary splits of the data. The algorithm begins with

a root node containing a sample of good and bad credit applicants. Then, the algorithm loops over all possible binary splits in order to find the attribute x and corresponding cut-off value c which gives the best separation into one side having mostly good credits and the other mostly bad credits. For example, in Figure 1 the figure of merit is optimized when the data in the root node is split between instances with attribute $x_i \geq c_i$ and those with $x_i < c_i$. This procedure is then repeated for the new daughter nodes until a stopping criterion is satisfied. Defining the purity p of a node as the fraction of good credit instances in it, the splitting attribute and cut-off value are those that minimize the sum of the Gini indices $p(1 - p)$ of the created daughter nodes. If, for any attribute or cut-off value, the sum of the Gini indices of the daughter nodes is higher than the Gini index of the parent node, the parent node is not split. Since the Gini index is a measure of the statistical dispersion or diversity of the population in a node, minimizing the Gini index results in daughter nodes that are more homogeneous than the parent nodes.

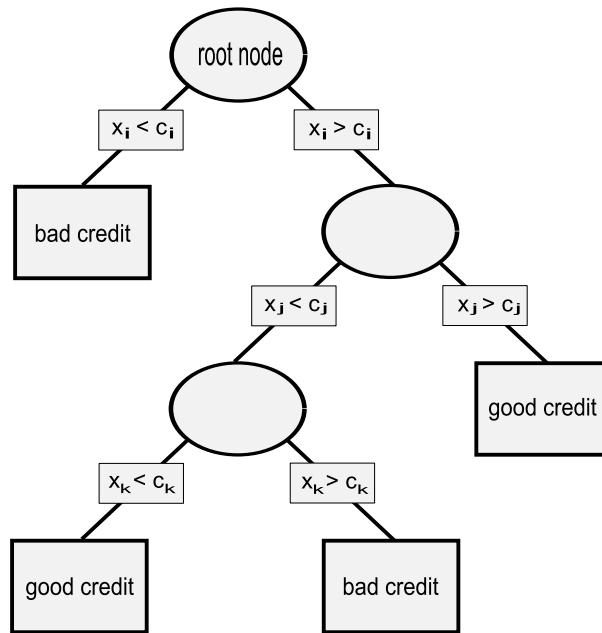


Figure 1: Illustration of a decision tree.

Unsplit nodes are denoted by “leaves” and are depicted by rectangles in Figure 1. The leaves are classified according to the most prevalent class in them. A leaf is called “good credit leaf” if it contains a number of good credit applicants larger than the number of bad credit applicants. Otherwise, it is called “bad credit leaf”. A good (bad) credit is correctly classified if it lands on a good (bad) credit leaf. Very frequently the resulting trees are quite large. Note that, in principle, a decision tree could be grown until all leaves contain only good credit instances or only bad credit instances. However, such tree would be highly overtrained. In these circumstances, the generalization performance may be improved if the tree is “pruned”. Pruning consists in cutting back the tree in order to get rid of statistically insignificant nodes (Breiman et al., 1984).

Decision trees have been available since the 1980’s and have been applied to the development of credit scoring models (Frydman et al., 1985; Davis et al., 1992). They are a powerful and flexible classifier. However, a well known limitation of decision trees is their

instability, since small fluctuations in the data sample may result in large variations in the classifications assigned to the instances. For example, if there are two attributes having similar discriminating power, a small fluctuation in one of these attributes may cause the algorithm to split a given node using the other attribute, while the former would have been selected without the fluctuation. Since the whole tree structure is modified below this node, the fluctuation may produce a completely different classifier response. This difficulty is overcome by growing a forest of decision trees and classifying the instances with the majority vote of the classifications given by individual trees.

2.2 Boosting

Boosting (Freund and Schapire, 1991; Schapire, 2002; Friedman, 2003) is a procedure that aggregates many “weak” classifiers in order to achieve a high classification performance. Additionally, boosting helps stabilizing the response of classifiers with respect to changes in the training sample. The boosting algorithm initiates by giving all credit applicants the same weight $w^{(0)}$. After a classifier is built, the weight of each applicant is changed according to the classification given by that classifier. Then, a second classifier is built using the reweighted training sample. This procedure is typically repeated several hundreds of times. The final classification of a credit applicant is a weighted average of the individual classifications over all classifiers. There are several methods to update the weights and combine the individual classifiers. The most popular boosting algorithm is AdaBoost (Freund and Schapire, 1996) which is adopted in this study. After the k th decision tree is built, the total misclassification error ε_k of the tree, defined as the sum of the weights of misclassified credits over the sum of the weights of all credits, is calculated:

$$\varepsilon_k = \frac{\sum_{i \text{ mis}} w_i^{(k)}}{\sum_i w_i^{(k)}}, \quad (1)$$

where i loops over all instances in the data sample. Then, the weights of misclassified credit applicants are increased (*boosted*)

$$w_i^{(k+1)} = \frac{1 - \varepsilon_k}{\varepsilon_k} w_i^{(k)}. \quad (2)$$

Finally, the new weights are renormalized, $w_i^{(k+1)} \rightarrow w_i^{(k+1)} / \sum_i w_i^{(k+1)}$ and the tree $k+1$ is constructed. Note that, as the algorithm progresses, the predominance of hard-to-classify instances in the training set is increased. The final classification or “score” of credit applicant i is a weighted sum of the classifications over the individual trees

$$F_i = \sum_{k=1}^N \log \left(\frac{1 - \varepsilon_k}{\varepsilon_k} \right) f_i^{(k)}, \quad (3)$$

where $f_i^{(k)} = 1(-1)$ if the k th tree makes the instance land on a good (bad) credit leaf and N is the number of grown trees. Therefore, good credits will tend to have large positive scores, while bad credits will tend to have large negative scores. Furthermore, trees with lower misclassification errors ε_k are given more weight when the final classification is computed.

3 Empirical analysis

3.1 Data sample

In this study, the credit scoring models were developed using two popular credit card application datasets from the UC Irvine Machine Learning Repository (Asuncion and Newman, 2007). The German credit dataset consists of 1000 instances, of which 700 instances correspond to creditworthy applicants and 300 instances correspond to applicants to whom credit should not be extended. Each applicant is described by 24 attributes describing the status of existing accounts, credit history records, loan amount and purpose, employment status and an assortment of personal information such as age, sex and marital status. Three attributes are continuous and the remaining are categorical. The Australian credit dataset contains 690 instances, of which 307 correspond to creditworthy applicants and 383 correspond to applicants to whom credit should be refused. Each instance is described by 14 attributes. Six attributes are continuous while the remaining are categorical. In order to preserve the confidentiality of the data, the names and values of the attributes were replaced by meaningless identifiers. This dataset has the appealing feature of containing attributes that are continuous, nominal with small number of values and nominal with large number of values. A few instances had attributes with missing values. These were replaced by the mode and mean of the attribute for categorical and continuous variables, respectively. Note that, because in the node splitting procedure only the best discriminating variable is selected, boosted decision trees are insensitive to the inclusion of attributes with weak discriminating power, while the training time only scales linearly with the dimensionality of the input patterns.

3.2 Performance tuning

In a pattern classification problem, the data sample is usually divided into a training set and an independent (out-of-sample) test set. The classifier learns the features of the population with the training set and its predictive power is estimated using the test set. In order to train classifiers with a large fraction of the available data and evaluate the generalization accuracy with the complete dataset a 10-fold cross-validation was implemented. This technique consists of randomly dividing the dataset into ten mutually exclusive subsets of equal size and, sequentially, testing each of these subsets using the classifier trained on the remaining subsets.

There is no formal theory specifying how to select the optimal topology and parameters for a given classifier. In practice, the selection of the best set of parameters is accomplished either by heuristic rules or by “grid-search”. In this approach, different parameter values are scanned and the set with best predictive performance is selected. Since the predictive performance of the algorithms may be a multimodal function of the parameters, large parameter ranges should be considered in order to minimize the likelihood of encountering local optima. The performance of boosted decision trees (BDT) is optimized by adjusting two parameters: the number of decision trees that are aggregated to form the final classifier and the minimum number of credit applicants that a tree node must contain in order to be split. When the number of applicants in a node reaches this threshold value, the growth of the branch is terminated. The multilayer perceptron (MLP) contained a single hidden

layer.¹ The input layer contained a number of nodes equal to the number of attributes in the samples (24 nodes for the German dataset and 14 nodes for the Australian dataset) while the output layer contained a single node. The activation function of the neurons in the hidden layer was a sigmoid, while a linear activation function was used in the output layer. The network was trained by error back-propagation using the steepest descent algorithm. Three parameters were optimized: the number of neurons in the hidden layer, the number of epochs and the learning rate. The support vector machine (SVM) was implemented with a Gaussian radial basis function. Two parameters were optimized: the width of the Gaussian kernel σ and the cost parameter C . To find the best pair (σ, C) a grid-search was performed using the recipe in Hsu et al. (2007), in which these parameters take values from exponentially growing sequences. All models were implemented using the framework provided by the TMVA package (Hoecker et al., 2007).

3.3 Results

The performance of credit scoring models is measured in terms of the capability of distinguishing the good credit population from the bad credit population in the test sample. As mentioned in Section 2, the BDT algorithm assigns to credit applicants a score according to Equation 3. Good credits will typically have large positive scores while bad credits will have large negative scores. Credit applicants with score above a certain threshold value are granted while the remaining are rejected. For a given cut-off value there are two types of incorrect predictions: the model grants credit to an applicant that will default on the financial obligation (Type I error) and the model rejects credit to an applicant that is creditworthy (Type II error or False Alarm Ratio). The cut-off value represents a compromise between a large efficiency for granting credit and a large rejection of bad credits. An excessively large efficiency for granting credit may result in severe economic losses due to delinquent costumers, while a credit policy that is too strict may result in opportunity costs that surpass the costs of default. The selected cut-off value will ultimately depend on the relative ratio of the misclassification costs associated to Type I and Type II errors.²

Since the cut-off value depends on the credit policy of the financial institution, it is convenient to express the performance of the models in terms of the receiver operating characteristics (ROC) curve. The ROC curve is a plot of the true positive rate (proportion of bad credit that is correctly classified) as a function of the false positive rate (Type II error) for the full range of possible cut-off values. Figure 2 and Figure 3 show the ROC curves for the German and Australian credit datasets obtained by merging the 10 cross-validation test sets. If a model could separate completely the two populations, it would always give correct predictions and never give incorrect predictions. In this case, the ROC curve would pass through the point (0,1) and the area under the ROC curve would be equal to 1. On the other hand, a random guess classifier would result in as many correct predictions as incorrect predictions being made. In this case, for any cut-off value, the

¹A network with a single hidden layer is sufficient to model a complex system to any desired degree of accuracy, provided sufficient hidden nodes are available (Hornik et al., 1989).

²In general, the costs associated with misclassifying bad applicants are financially more damaging than those associated with misclassifying good applicants.




Figure 2: Receiver Operating Characteristics (ROC) curve for the multilayer perceptron (MLP), support vector machine (SVM) and boosted decision trees (BDT), for the German credit dataset.

Hit Rate would be on average equal to the False Alarm Ratio and the ROC curve would be a 45 degree straight line intersecting (0,0) and (1,1). A model that performs better than random guessing gives a concave ROC curve above this straight line. The higher is the model accuracy, the steeper will the ROC curve be. Therefore, the area under the ROC curve (AUC) is a measure of the generalization accuracy which is independent of the cut-off value.

Model	German data	Australian data
MLP	78.32%	92.34%
SVM	79.87%	92.87%
BDT	81.08%	94.03%

Table 1: Comparison of the area under the ROC curve for the multilayer perceptron (MLP), support vector machine (SVM) and boosted decision trees (BDT).

Table 1 gives the AUC predicted by the three models which is obtained by trapezoidal integration. For the German dataset the SVM outperforms the MLP, while BDT outperforms both the MLP and the SVM. Also of note is that the performance of BDT and SVM is roughly equal for false positive rates above 0.3. For the Australian dataset a similar ordering of the predictive performance of the three models is observed. Again, while the global performance of BDT is better than that of SVM, for false positive rates greater than 0.4, the performance of these techniques is comparable.

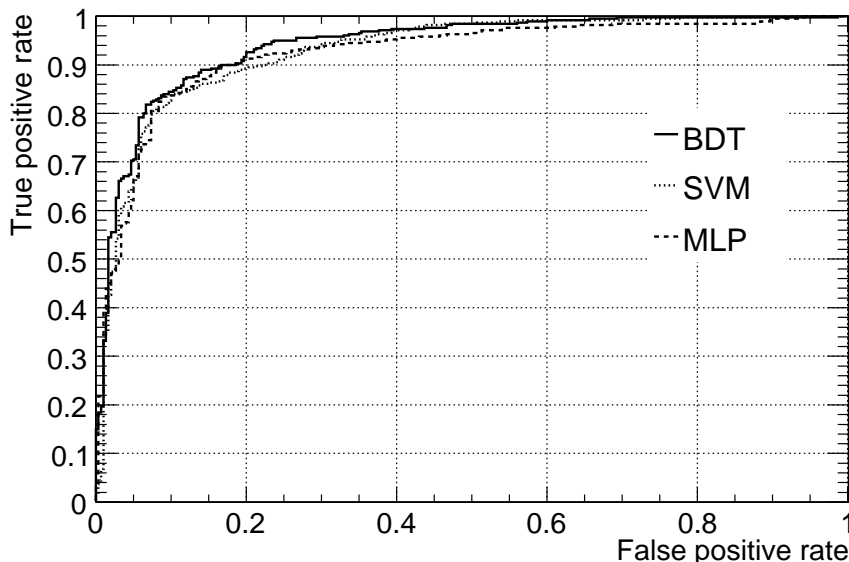


Figure 3: Receiver Operating Characteristics (ROC) curve for the multilayer perceptron (MLP), support vector machine (SVM) and boosted decision trees (BDT), for the Australian credit dataset.

3.4 Comparison of the AUC estimates

In order to test the statistical significance of the differences between the areas under the ROC curves predicted by the models under consideration, the nonparametric approach introduced by DeLong et al. (1988) is followed. The AUC can be interpreted as the probability that the score of a randomly selected good credit applicant is higher than that of a randomly selected bad credit applicant. Therefore, denoting by $X_i^{(g)}$, $i = 1, \dots, n_g$ the estimated scores for the good credit set and by $X_j^{(b)}$, $j = 1, \dots, n_b$ the estimated scores for the bad credit set, an unbiased estimator of the AUC is given by the Wilcoxon-Mann-Whitney statistic

$$\hat{\theta} = \frac{1}{n_b n_g} \sum_{j=1}^{n_b} \sum_{i=1}^{n_g} \mathbf{1}_{X_i^{(g)} > X_j^{(b)}}, \quad (4)$$

where the indicator function $\mathbf{1}_{X_i^{(g)} > X_j^{(b)}}$ is 1 if $X_i^{(g)} > X_j^{(b)}$ and 0 otherwise. In order to obtain an estimate of the variance of $\hat{\theta}$, the structural components of the i th good credit and j th bad credit must be calculated

$$v(X_i^{(g)}) = \frac{1}{n_b} \sum_{j=1}^{n_b} \mathbf{1}_{X_i^{(g)} > X_j^{(b)}}, \quad v(X_j^{(b)}) = \frac{1}{n_g} \sum_{i=1}^{n_g} \mathbf{1}_{X_i^{(g)} > X_j^{(b)}}. \quad (5)$$

Then, an estimator for the variance of $\hat{\theta}$ can be obtained from

$$\hat{\text{Var}}(\hat{\theta}) = \frac{1}{n_g(n_g - 1)} \sum_{i=1}^{n_g} [v(X_i^{(g)}) - \hat{\theta}]^2 + \frac{1}{n_b(n_b - 1)} \sum_{j=1}^{n_b} [v(X_j^{(b)}) - \hat{\theta}]^2. \quad (6)$$

In order to compare the AUC of two alternative models, A and B , the covariance of the corresponding AUC estimators must also be obtained

$$\begin{aligned} \text{C}\hat{\text{ov}}(\hat{\theta}_A, \hat{\theta}_B) &= \frac{1}{n_g(n_g - 1)} \sum_{i=1}^{n_g} \left[v_A(X_i^{(g)}) - \hat{\theta}_A \right] \left[v_B(X_i^{(g)}) - \hat{\theta}_B \right] \\ &+ \frac{1}{n_b(n_b - 1)} \sum_{j=1}^{n_b} \left[v_A(X_j^{(b)}) - \hat{\theta}_A \right] \left[v_B(X_{Bj}^{(b)}) - \hat{\theta}_B \right]. \end{aligned} \quad (7)$$

To test the null hypothesis $H_0 : \hat{\theta}_A = \hat{\theta}_B$ versus the alternative hypothesis $H_1 : \hat{\theta}_A \neq \hat{\theta}_B$ the following test statistic is computed

$$T = \frac{\left(\hat{\theta}_A - \hat{\theta}_B \right)^2}{\text{V}\hat{\text{ar}}(\hat{\theta}_A - \hat{\theta}_B)}, \quad (8)$$

where

$$\text{V}\hat{\text{ar}}(\hat{\theta}_A - \hat{\theta}_B) = \text{V}\hat{\text{ar}}(\hat{\theta}_A) + \text{V}\hat{\text{ar}}(\hat{\theta}_B) - 2\text{C}\hat{\text{ov}}(\hat{\theta}_A, \hat{\theta}_B). \quad (9)$$

The test statistic T is asymptotically χ^2 -distributed with one degree of freedom

Test	German data		Australian data	
	T	p-value	T	p-value
MLP – SVM	2.781	9.54%	1.132	28.75%
MLP – BDT	4.916	2.66%	6.778	0.92%
SVM – BDT	1.774	18.29%	3.737	5.32%

Table 2: Statistical test for comparing the area under the ROC curves estimated by the different models.

Table 2 shows the results of applying this test to the estimated ROC curves. For both datasets one can reject the hypothesis $\hat{\theta}_{BDT} = \hat{\theta}_{MLP}$ with a 95% significance level and, therefore, there is a strong evidence that the performance of BDT is better than that of the MLP. For the Australian dataset there is also strong evidence that BDT outperforms SVM. However, for the German dataset the difference between these methods is not highly significant.

3.5 Relative importance of the attributes

Boosted decision trees provide a straightforward and intuitive measure of the relative contribution of the attributes to separate instances according to the target classification. Using this approach a ranking of the most useful attributes can be established. This ranking is derived by counting the number of times an attribute is employed in the node splitting procedure and by weighting each split by the separation gain-squared it has accomplished and by the number of instances in the node (Breiman et al., 1984).

Figure 4 shows the relative importance of the attributes for the German credit dataset. The first and 4th attributes are the most important. These attributes correspond to the

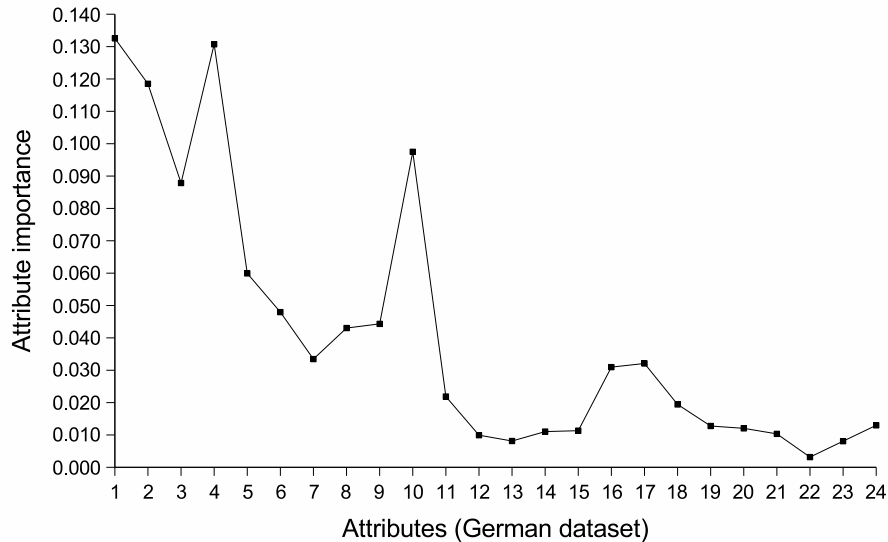


Figure 4: Relative importance of attributes predicted by boosted decision trees for the German dataset.

status of the existing checking accounts and the credit amount, respectively. They are followed by the 2nd attribute (duration of the loan) and the 10th attribute (age of the applicant). Also important is the 3rd attribute, that represents the credit history of the applicant (e.g., if previous credits were paid punctually or there were delays in paying off). Attributes 5th to 9th have moderate importance. They correspond to the status of savings accounts, the employment condition, the marital status and sex, the amount of years living in the present residence and the property that the applicant owns, respectively. Figure 5 shows the relative importance of the attributes for the Australian credit dataset. The nature of the attributes in this dataset is unknown. In this dataset, the 8th attribute is clearly the most important. Also of note is that the contributions of attributes 1, 11 and 12 are almost negligible.

4 Conclusions

This paper introduced a credit scoring model of consumer loans using boosted decision trees: a learning technique that allows to combine several decision trees to form a classifier which is obtained from a weighted majority vote of the classifications given by individual trees. The generalization accuracy of boosted decision trees was compared with that of a multilayer perceptron and support vector machines. Boosted decision trees outperformed the multilayer perceptron and the support vector machines on two real world credit card application datasets. On the basis of these results, it can be concluded that boosted decision trees may be a competitive alternative to these techniques in credit scoring applications. It was also shown that boosted decision trees provide an elegant way to rank the attributes that most significantly indicate the likelihood of default.

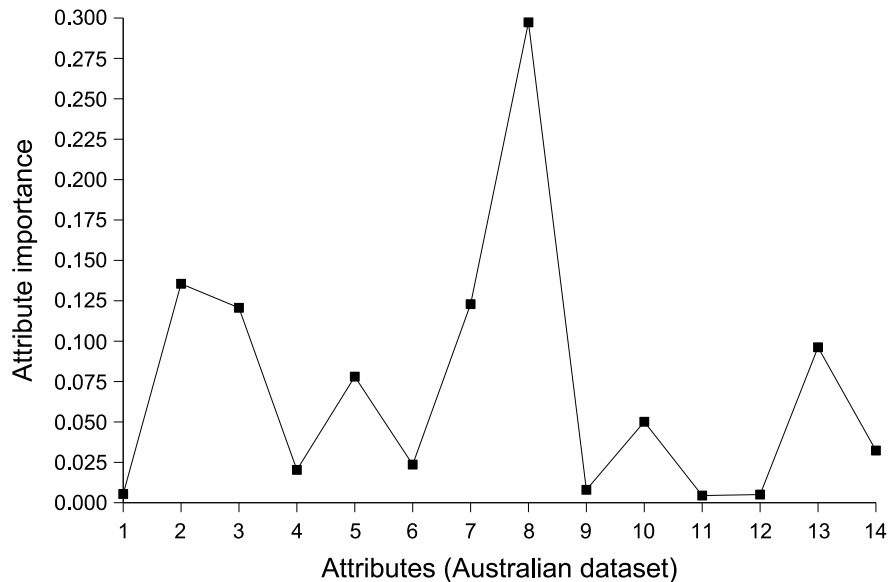


Figure 5: Relative importance of attributes predicted by boosted decision trees for the Australian dataset.

Acknowledgments

This work was supported by grant SFRH/BPD/20616/2004 of Fundação para a Ciência e Tecnologia.

References

- A. Asuncion and D.J. Newman. UCI machine learning repository, 2007. URL <http://www.ics.uci.edu/~mllearn/MLRepository.html>.
- B. Baesens, T. Van Gestel, S. Viaene, M. Stepanova, J. Suykens, and J. Vanthienen. Benchmarking state-of-the-art classification algorithms for credit scoring. *Journal of the Operational Research Society*, 54:1028–1088, 2003.
- L. Breiman, J.H. Friedman, R.A. Olshen, and C.J. Stone. *Classification and regression trees*. Wadworth International Group, Belmont, California, 1984.
- J.N. Crook, D.B. Edelman, and L.C. Thomas. Recent developments in consumer credit risk assessment. *European Journal of Operational Research*, 183:1447–1465, 2007.
- R.H. Davis, D.B. Edelman, and A.J. Gammerman. Machine learning algorithms for credit-card applications. *IMA Journal of Management Mathematics*, 4:43–51, 1992.
- E. DeLong, D. DeLong, and D. Clarke-Pearson. Comparing the area under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, 44:837–845, 1988.

- Y. Freund and R.E. Schapire. A short introduction to boosting. *J. Jpn. Soc. Artif. Intell.*, 14(5):771, 1991.
- Y. Freund and R.E. Schapire. Experiments with a new boosting algorithm. *in: Proceedings of the 13th International Conference on Machine Learning*, pages 148–156, 1996.
- J. Friedman. Recent advances in predictive (machine) learning. *Proceedings of Phystat, Stanford University*, 2003.
- H.E. Frydman, E.I. Altman, and D-L. Kao. Introducing recursive partitioning for financial classification: the case of financial distress. *Journal of Finance*, 269–91:40(1), 1985.
- W.E. Henley and D.J. Hand. A k -nearest neighbor classifier for assessing consumer risk. *Statistician*, 44(1):77–95, 1996.
- A. Hoecker, P. Speckmayer, J. Stelzer, F. Tegenfeldt, H. Voss, and K. Voss. TMVA – toolkit for multivariate data analysis. *arXiv:physics/0703039*, 2007.
- K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- N.-C. Hsieh. Hybrid mining approach in the design of credit scoring models. *Expert Systems with Applications*, 28(4):655–665, 2005.
- C.-W. Hsu, C.-C. Chang, and C.-J. Lin. A practical guide to support vector classification, 2007. URL <http://www.csie.ntu.edu.tw/~cjlin>.
- H.L. Jensen. Using neural networks for credit scoring. *Managerial Finance*, 18(6):15–26, 1992.
- T.-S. Lee and I.-F. Chen. A two-stage hybrid credit scoring model using artificial neural networks and multivariate adaptive regression splines. *Expert Systems with Applications*, 28(4):743–752, 2002.
- T.-S. Lee, C.-C. Chiu, C.-J. Lu, and I.-F. Chen. Credit scoring using the hybrid neural discriminant technique. *Expert Systems with Applications*, 23(3):245–254, 2002.
- C.-S. Ong, J.-J. Huang, and G.-H. Tzeng. Building credit scoring models using genetic programming. *Expert Systems with Applications*, 29(1):41–47, 2005.
- A.K. Reichert, C.C. Cho, and G.M. Wagner. An examination of the conceptual issues involved in developing credit-scoring models. *Journal of Business and Economic Statistics*, 1(2):101–114, 1983.
- R.E. Schapire. The boosting approach to machine learning: an overview. *in: Proceedings of the 2002 MSRI Workshop on Nonlinear Estimation and Classification*, Springer Verlag, pages 149–173, 2002.
- D. West. Neural network credit scoring models. *Computers and Operations Research*, 27: 1131–1152, 2000.

- D. West, S. Dellana, and J. Qian. Neural network ensemble strategies for financial decision applications. *Computers and Operations Research*, 32:2543–2559, 2005.
- J.C. Wiginton. A note on the comparison of logit and discriminant models of consumer credit behavior. *Journal of Financial and Quantitative Analysis*, 15:757–770, 1980.