

Improving Academic Performance Prediction by Dealing with Class Imbalance

Nguyen Thai-Nghe, Andre Busche, and Lars Schmidt-Thieme
Information Systems and Machine Learning Lab
University of Hildesheim
Samelsonplatz 1, 31141 Hildesheim, Germany
{nguyen, busche, schmidt-thieme}@ismll.de

Abstract—This paper introduces and compares some techniques used to predict the student performance at the university. Recently, researchers have focused on applying machine learning in higher education to support both the students and the instructors getting better in their performances. Some previous papers have introduced this problem but the prediction results were unsatisfactory because of the class imbalance problem, which causes the degradation of the classifiers. The purpose of this paper is to tackle the class imbalance for improving the prediction/classification results by over-sampling techniques as well as using cost-sensitive learning (CSL). The paper shows that the results have been improved when comparing with only using baseline classifiers such as Decision Tree (DT), Bayesian Networks (BN), and Support Vector Machines (SVM) to the original data sets.

Keywords-Academic performance; Prediction; Class imbalance; Cost-sensitive.

I. INTRODUCTION

Machine learning has been applied in plenty of areas. For examples, it can be used to predict/classify customer behaviors in marketing or sales, to detect fraudulent or default in banking, to diagnose diseases in medicine, and recently, to a new area which is education. Moreover, the universities desire to improve their educational quality, hence, how to apply machine learning techniques in higher education to help the universities, instructors, and the students getting better in their performances become more and more attractive to both university managers and researchers.

In the first conference of educational data mining 2008, [6] compared different data mining methods and techniques to classify students based on their Moodle usage data and the final marks obtained in their respective courses; [15] proposed a model with different types of education-related questions and the data-mining techniques appropriate for them. For examples, predicting student performance, clustering similar students, and associating types of students with appropriate courses; [18], [21], and [3] used BN, DT, and other common techniques to predict the student results; another study [16] was done by us recently to predict the student performance at two real case studies: Can Tho University, Vietnam (CTU)¹, and Asian Institute of Technology, Thailand (AIT)².

¹<http://www.ctu.edu.vn>

²<http://www.ait.ac.th>

Researchers have spent a lot of time in the task of predicting student performances while achieving only fair results. One major problem is that the number of "pass students" is much higher than the number of "fail students". This skew distribution is the main reason causing the degradation of classifiers.

The purposes of this study are both to improve the results of classification techniques in predicting the student performance of two real world case-studies (CTU, AIT) by dealing the class imbalance problem, and to compare the results of some common classification techniques before and after approaching with class imbalance.

The rest of the paper is organized as follows. Section 2 introduces some common techniques used to tackle with the class imbalance problem; section 3 represents our methodologies; section 4 explores the experimental datasets; section 5 shows the results and evaluations; and finally, section 6 and 7 represent discussion and conclusion respectively.

II. DEALING WITH CLASS IMBALANCE PROBLEM

To deal with imbalanced datasets, researchers used to focus on data level and classifier (algorithm) level. At data level, the common task is the modification of class distribution using over-sampling or under-sampling techniques. At classifier level, some common techniques were introduced such as manipulating classifiers internally, one-class learning, ensemble methods, and CSL [20]. We will briefly review some techniques used in this study.

A. Modifying class distribution

To modify the class distribution, over-sampling and under-sampling are usually used. Under-sampling discards a lot of useful information and usually appropriates for large datasets. The most common over-sampling method is random over-sampling, which randomly duplicates the examples in the datasets. Another over-sampling method was introduced by [14] and called SMOTE (Synthetic Minority Over-sampling Technique). SMOTE generates new artificial minority examples by interpolating between the existing minority examples rather than simply duplicating the original examples. This method, at first, finds k nearest neighbors of each minority example (according to authors, $k=5$); then it selects a random nearest neighbor; finally, the new synthetic

examples are generated along the line segment joining a minority class sample and its nearest neighbor.

B. Cost-sensitive learning

Most of classifiers assume that the misclassification costs (false negative and false positive cost) are the same. In some real-world applications, this assumption may not be true. For examples, the cost of mailing to non-buyers is less expensive than the cost of non-mailing to the buyers [4]; or the cost of predicting non-terrorist to terrorist is much cheaper than the cost of misclassifying an actual terrorist who carries a bomb to a flight. Cost is not necessarily monetary, for examples, it can be a waste of time or even the severity of an illness [17]. In our studies, the cost of misclassifying the actual "fail students" to "pass students" (so we can not help them, consequently, they will be expelled from the university) is much costlier than the false alarm.

When learning with two-class problems, we assume that the positive class (+) represents for the minority examples, and the negative class (-) represents for the majority examples. Let $C(i, j)$ be the cost of predicting an example belonging to class i when in fact it belongs to class j , then the confusion matrix and the cost matrix are described as in tables I and II respectively.

Table I
CONFUSION MATRIX

		Predict classes	
		Positive	Negative
Actual classes	Positive	TP	FN
	Negative	FP	TN

Table II
COST MATRIX

		Predict classes	
		Positive	Negative
Actual classes	Positive	$C(+, +)$	$C(-, +)$
	Negative	$C(+, -)$	$C(-, -)$

Given this cost matrix, an example x can be classified into class i with the minimum expected cost (*conditional risk*) by using Bayes risk in equation (1)

$$\mathcal{H}(i|x) = \arg \min_i \left(\sum_j^N P(j|x)C(i, j) \right) \quad (1)$$

where N is number of classes and $P(j|x)$ is the posterior probability of classifying an example x into class j .

Another term is the cost ratio, which is the proportion of false negative and false positive determined by $C(-, +)/C(+, -)$. The purpose of cost-sensitive learning is to build the model with minimum misclassification cost

as described in equation (2)

$$Totalcost = C(-, +) \times FN + C(+, -) \times FP \quad (2)$$

C. Evaluation metrics

In the case of imbalanced datasets, accuracy metric becomes useless. For example, suppose the dataset has 990 negative examples and only 10 positive examples (this minority is usually very important). Most of the classifiers designed to maximize the accuracy, so in this case, they will classify all examples belong to the majority class to get the maximum of 99% accuracy. This result has no meaning because all the positive examples are classified incorrectly. To evaluate the model in the case of class imbalance, researchers usually use F-measure and the AUC, which are related to some other metrics described in the following.

False negative rate (FNR) is the proportion of positive examples misclassified as belonging to the negative class, $FNR = FN/(TP + FN)$

True negative rate (TNR) is the proportion of negative examples correctly classified as belonging to the negative class, $TNR = TN/(FP + TN)$

True Positive Rate (TPR) (or Recall R) is the proportion of positive examples correctly classified as belonging to the positive class, determined by equation (3)

$$TPR = R = \frac{TP}{TP + FN} \quad (3)$$

False Positive Rate (FPR) is the proportion of negative examples misclassified as belonging to the positive class, $FPR = FP/(TN + FP)$

Precision (P) is the positive predictive value determined by equation (4)

$$P = \frac{TP}{TP + FP} \quad (4)$$

F-Measure is an evaluation metric which considers both precision and recall of the testing results (β is used to set equal to 1)

$$F - Measure = \frac{(1 + \beta^2) \times P \times R}{(\beta^2 \times P + R)} \quad (5)$$

Another important metric is the area under the ROC curve (AUC) [2], in which ROC (Receiver Operating Characteristic) curve is a graphical approach for displaying the tradeoff between TPR (y-axis) and FPR (x-axis) of a classifier. We will use these metrics to evaluate our models in section 5.

III. METHODOLOGIES

The first step in this study is to collect the real datasets from the relational databases. Based on experts and feature selections, we have selected the most important attributes for the prediction tasks. The second step is to apply techniques to tackle the class imbalance based on three methods:

- Method 1: Modifying the class distribution by using SMOTE.

- Method 2: Applying CSL to minimize the total misclassification cost.
- Method 3: Combining SMOTE and CSL

The third step is to evaluate and compare the results of these methods by experimenting with 3 classifiers on both the original and the re-balanced datasets. We use an evaluation protocol visualized as in the figure 1. This evaluation works as 10-folds cross-validation. At first, the dataset was separated; next, the enrichment data was generated on the train set; then the classifiers were built; finally, the evaluations were applied to the original test data. The final results were collected from average of ten run-times.




Figure 1. Evaluation protocol

The final step in this research is to build the intelligent system based on the best model we have evaluated. This system is deployed on the local web to help not only the students in predicting their result to have clever study plans but also the instructors who give them advices or more tutorials.

IV. EXPERIMENTAL DATASETS

The first dataset was got from Can Tho University, Vietnam (CTU), which has 20492 examples and 14 attributes. The second dataset was collected from the Asian Institute of Technology, Thailand (AIT), which has 936 examples and 14 attributes. We use these two datasets for predicting the student results (“pass”, “fail”).

Furthermore, to observe the results on the less imbalance data, two other datasets were also collected. The third dataset was taken from UCI repository³ named Teaching Assistant Evaluation (TAE). This data consists of evaluation of teaching performance for 151 teaching assistant assignments at University of Wisconsin-Madison. The scores were divided into 2 classes (“low”, “med-high”). The fourth dataset was collected from the Journal Statistics of Education Data Archive⁴ named “U.S. News College” (USNC) having 1204 colleges, 33 attributes. This data is used to classify the graduation rate. We artificially assigned binary target labels for USNC by dividing the students at 85% graduation rate to get the imbalanced ratio between two real datasets AIT/CTU and TAE, as in table III

³<http://archive.ics.uci.edu/ml>

⁴<http://www.amstat.org/publications/jse/index.html>

Table III
EXPERIMENTAL DATASETS

Dataset	#Examples	#Attributes	#minor	#major/ #minor
CTU	20492	14	1565	13.09
AIT	936	14	75	12.48
TAE	151	6	49	3.08
USNC	1024	33	132	7.76

V. EXPERIMENTAL RESULTS AND EVALUATIONS

The experiments and the application system in this study were developed based on the WEKA library.⁵

A. Hyperparameters search

We have applied hyperparameter searches to look for the optimal parameters on all classifiers. To observe how the classifiers are influenced by the each dataset, we have also searched for the optimal parameter (percentage) of over-sampling techniques. Each dataset has its own structure, so this percentage is also different. There are some notations such as CTU-SM150 represents for the dataset with SMOTE at 150%, and CTU-O means that the Original CTU dataset was used.

Table IV shows optimal hyper-parameters for DT. The “unpruned” indicates whether pruning is performed; “use-Laplace” indicates whether counts at leaves are smoothed based on Laplace [9].

Table IV
OPTIMAL HYPER-PARAMETERS FOR DT

Dataset	Un pruned	use Laplace	Binary split	#Instances per leaf
CTU-O	True	True	True	5
CTU-SM150	False	True	True	10
AIT-O	True	True	True	2
AIT-SM200	True	True	False	2
TAE-O	True	False	True	2
TAE-SM50	True	False	True	2
USNC-O	True	True	True	7
USNC-SM200	True	True	True	3

Table V describes the hyperparameters for BN. Alpha is used for estimating the probability tables and can be interpreted as the initial count on each value [12]; MaxNoOf-Parents (M) sets the maximum number of parents a node in the BN can have. (M = 1 for Naive Bayes classifier; M = 2 for Tree Augmented Bayes Network [11]; M > 2 for Bayes Net Augmented Bayes Network [10]; The framework K2 developed by [7] for induction of BN from data is used.

In a dataset \mathcal{D} consisting of n examples $(x_i \in \mathcal{X}, y_i)$ with a kernel function $k(\mathbf{x}, \mathbf{x}')$, SVM will predict a new example

⁵<http://www.cs.waikato.ac.nz/ml/weka/>

Table V
OPTIMAL HYPER-PARAMETERS FOR BN

Dataset	MaxNoOfParents	Search Algorithm	Alpha
CTU-O	3	Local.K2	0.3
CTU-SM150	3	Local.K2	0.3
AIT-O	2	Local.K2	0.4
AIT-SM200	2	Local.K2	0.1
TAE-O	2	Local.K2	0.2
TAE-SM50	2	Local.K2	0.1
USNC-O	1	Local.K2	0.1
USNC-SM200	1	Local.K2	0.7

x by using the function

$$f(x) = \text{sign} \left(\sum_{i=1}^n \alpha_i y_i k(\mathbf{x}, \mathbf{x}') + b \right) \quad (6)$$

where b is the threshold and α_i is determined by solving the Lagrangian optimization problem, $L_p =$

$$\frac{1}{2} \|\mathbf{w}\|^2 + C \sum_i \xi_i - \sum_i \alpha_i \{y_i(x_i \cdot w + b) - 1 + \xi_i\} - \sum_i \mu_i \xi_i \quad (7)$$

where ξ_i is a positive slack variable, μ_i is Lagrange multiplier used to enforce ξ_i being positive, and C is user-specified parameter representing the penalty of misclassifying the training instances and it can be chosen based on calibration. In non-linear problems, the kernel trick k is used to maximum margin hyperplanes. Two commonly used kernel functions are Polynomial Function

$$k(\mathbf{x}, \mathbf{x}') = (\mathbf{x} \cdot \mathbf{x}')^p \quad (8)$$

and Radial Basis Function

$$k(\mathbf{x}, \mathbf{x}') = e^{-\gamma \|\mathbf{x} - \mathbf{x}'\|^2} \quad (9)$$

We have searched for the optimal hyperparameters of C , exponent p , and γ in equations (7), (8), and (9) respectively. Table VI represents these results, in which bLogistic indicates whether to fit logistic models to the outputs.

Table VI
OPTIMAL HYPER-PARAMETERS FOR SVM

Dataset	Kernel	p/γ	C	bLogistic
CTU-O	Poly	1	1.5	Yes
CTU-SM150	RBF	0.01	5.0	Yes
AIT-O	Poly	1	1.6	Yes
AIT-SM200	Poly	1	2.0	Yes
TAE-O	Poly	3	3.0	Yes
TAE-SM80	Poly	1	2.0	Yes
USNC-O	Poly	1	1.0	Yes
USNC-SM200	Poly	2	2.0	Yes

When learning with cost-sensitive, this study also searches for the cost ratio between false negative and false positive. We will investigate more clearly in next section.

B. Evaluation and comparison

We have experimented with three classifiers on four datasets. The AUC and F-Measure results are shown in table VII. In each dataset, this study shows that the classification results after dealing the class imbalance problem are normally better than learning on original data.

Another important method to deal with class imbalance is CSL, in which the classifiers will not consider the FN and FP equally. In our application, FN is the misclassification cost from the actual "fail student" to "pass student" and FP is the misclassification from the "pass" to the "fail" one. What is the cost in this case? Suppose that a second year student fails in this year, if he will fail again in next year then he will be expelled from the university, so the cost of both student and the university can be estimated as FN-Cost = (tuition fees + living fees + insurance fees + other fees) * 2-years-have-passed + (tuition fees will be lost from the university) * 2-years-in-future + (invaluable moral-cost). Conversely, the cost of FP can be approximated as FP-Cost = (the fee of some tutorials).

In fact, the cost of FN is much higher than FP. To examine how the total cost change when we increase the cost ratio, we have experimented on the cost ratio from 2 to 20. Figures 2 (only uses CSL) and 3 (combining SMOTE with CSL) display the total cost produced by DT, BN, SVM, and DT without cost-sensitive (DT-O). The results indicate that DT has the lowest total cost when comparing with the others on original data, and SVM has the lowest total cost on re-balanced data. Total cost is also reduced when learning on over-sampling datasets.




Figure 2. Total cost of AIT-O

To observe clearly how CSL effects on our problems, we have also experimented with three other ensemble methods, which can be used for handling class imbalance problem, AdaBoost [22], Bagging, and MetaCost [17] on the baseline DT. (Since the limitation of space, some other results are not presented here).

Figures 4 and 5 compare the total cost on ensemble methods. AdaBoost, Bagging and original-DT increase the total cost as a linear function when the cost ratio is increased. MetaCost is better than CSL in the case of small datasets

Table VII
AUC AND F-MEASURE

Algorithm	Dataset	TPR	FPR	Precision	AUC	F-Measure
Decision Tree	AIT-O	0.160	0.063	0.182	0.692 (0.090)	0.170 (0.120)
	AIT-SM200	0.280	0.038	0.389	0.763 (0.060)	0.326 (0.090)
	CTU-O	0.347	0.029	0.507	0.879 (0.010)	0.412 (0.030)
	CTU-SM150	0.466	0.056	0.550	0.882 (0.010)	0.495 (0.020)
	TAE-O	0.375	0.147	0.545	0.660 (0.130)	0.444 (0.080)
	TAE-SM50	0.563	0.382	0.409	0.733 (0.100)	0.474 (0.090)
	USNC-O	0.341	0.022	0.615	0.811 (0.070)	0.448 (0.010)
	USNC-SM200	0.545	0.062	0.647	0.850 (0.030)	0.533 (0.020)
Bayesian Networks	AIT-O	0.080	0.021	0.250	0.655 (0.080)	0.118 (0.050)
	AIT-SM200	0.120	0.031	0.273	0.731 (0.050)	0.167 (0.100)
	CTU-O	0.502	0.073	0.375	0.845 (0.010)	0.430 (0.030)
	CTU-SM150	0.521	0.071	0.391	0.877 (0.010)	0.456 (0.020)
	TAE-O	0.438	0.206	0.438	0.680 (0.130)	0.296 (0.080)
	TAE-SM90	0.625	0.294	0.500	0.688 (0.100)	0.556 (0.090)
	USNC-O	0.514	0.123	0.380	0.848 (0.060)	0.470 (0.020)
	USNC-SM200	0.591	0.115	0.388	0.864 (0.030)	0.531 (0.020)
Support Vector Machines	AIT-O	0.120	0.021	0.250	0.690 (0.016)	0.167 (0.008)
	AIT-SM100	0.200	0.035	0.330	0.739 (0.053)	0.250 (0.092)
	CTU-O	0.385	0.026	0.500	0.810 (0.012)	0.450 (0.020)
	CTU-SM150	0.454	0.038	0.510	0.880 (0.010)	0.480 (0.015)
	TAE-O	0.500	0.176	0.571	0.715 (0.040)	0.467 (0.039)
	TAE-SM80	0.625	0.088	0.769	0.757 (0.046)	0.615 (0.027)
	USNC-O	0.409	0.011	0.418	0.809 (0.064)	0.416 (0.014)
	USNC-SM200	0.545	0.070	0.490	0.828 (0.043)	0.545 (0.009)

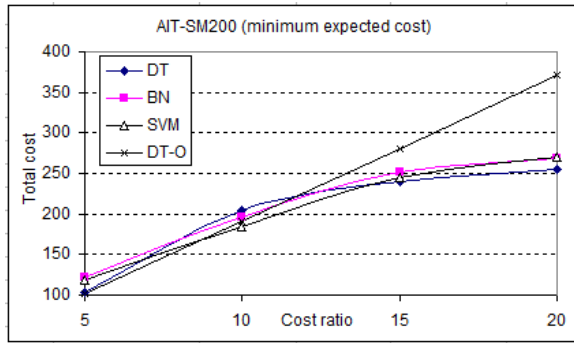


Figure 3. Total cost of AIT-SM200

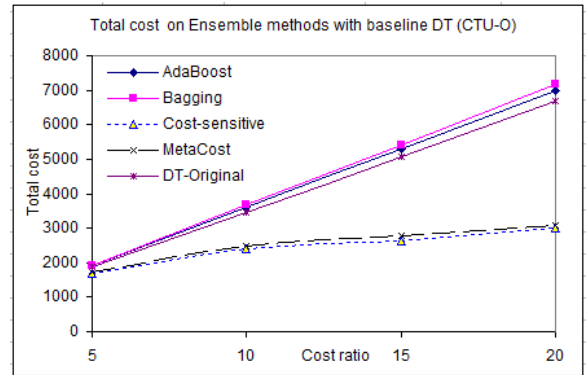


Figure 4. Total cost of ensembles (CTU-O)

(e.g. AIT-O), but when learning on larger datasets (e.g. CTU-O) then CSL outperforms.

VI. DISCUSSIONS

In this study, the results have been improved when comparing with our previous works [16], which only worked on original datasets. Although some papers had researched in this area such as [18], [21], and [3], they evaluated their models in terms of accuracy, which has less meaning in the case of class imbalance as we analyzed in section 2.3. We have used AUC, F-Measure, and total cost to evaluate the models on three methods proposed in section 3 and recognized that the results are reasonable. When the datasets are large enough, CSL is better than MetaCost in the case of increasing cost ratio.

In the ethical issues, it should be aware that there is no segregation from using this system to treat with "low performance" or "high performance" students. As we said in the introduction, the purpose of this study is to improve the prediction results to help the them get better in studying.

VII. CONCLUSIONS

In this study, we have applied machine learning techniques to improve the prediction results of academic performances for two the real case studies. Three methods have been used to deal with the class imbalance problem and all of them show satisfactory results. We first re-balanced the datasets and then used both cost-insensitive and sensitive learning with SVM for the small datasets and with Decision Tree for

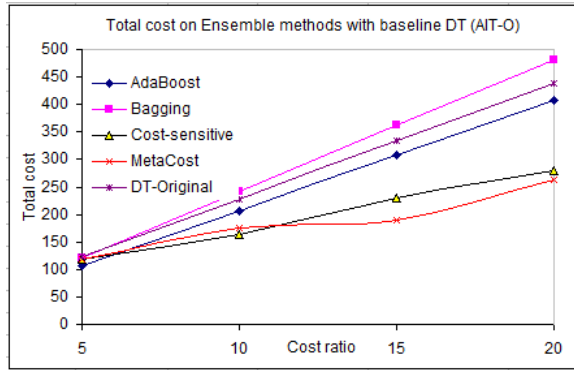


Figure 5. Total cost of ensembles (AIT-O)

the larger datasets. Our models are initially deployed on the local web. In future works, we will cross the results from using this system with the expected results in educational field.

In addition, we can integrate more features such as combining with E-learning system or predicting results in the E-learning. Moreover, most universities today are using the Credit System, in which each semester the students get confused with dozen of courses. So, how to help them choose appropriate courses based on their previous study results becomes the real problem. We can solve this by using association rules or recommender system approach [1] [13].

ACKNOWLEDGMENT

This research was funded by the "Teaching and Research Innovation Grant" Project of Can Tho University-Vietnam, and partly supported by the Information Systems and Machine Learning Lab, University of Hildesheim, Germany.

REFERENCES

- [1] A. Felfernig and G. Friedrich L. Schmidt-Thieme, *Intelligent Systems Special Issue on Recommender Systems*, IEEE Computer Society, Guest Edition, 2007.
- [2] A. P. Bradley, *The use of the area under the ROC curve in the evaluation of machine learning algorithms*, Pattern Recognition, volume 7, number 30, pages 1145-1159, 10997.
- [3] B. Minaei-Bidgoli and D. A. Kashy and G. Kortemeyer and W. F. Punch, *Predicting student performance: an application of data mining methods with an educational web-based system*, 33rd Annual Conference on Frontiers in Education (FIE 2003), volume 1, pages 13-18, 2003.
- [4] C. Elkan, *The foundations of cost-sensitive learning*, 17th International Joint Conference on Artificial Intelligence, Seattle - WA, pages 973-978, 2001.
- [5] C. Romero and S. Ventura, *Educational Data Mining: a Survey from 1995 to 2005*, Expert Systems with Applications, volume 33, number 1, pages 135-146, 2007.
- [6] C. Romero and S. Ventura and P. G. Espejo and C. Hervás, *Data Mining Algorithms to Classify Students*, 1st International Conference on Educational Data Mining, Montral, Quebec, Canada, 2008.

- [7] G. Cooper and E. Herskovits, *A Bayesian method for the induction of probabilistic networks from data*, Machine Learning, number 9, volumn 4, pages 309-347, 1992.
- [8] G. E. A. P. A. Batista and R. C. Prati and M. C. Monard, *A Study of the Behavior of Several Methods for Balancing Machine Learning Training Data*, Sigkdd Explorations, volume 6, number 1, pages 20-26, 2004.
- [9] G. M. Weiss and F. Provost, *Learning when training data are costly: the effect of class distribution on tree induction*, Journal of AI Research, pages 315-354, 2003.
- [10] J. Cheng and R. Greiner, *Learning Bayesian Belief Network Classifiers: Algorithms and System*, in proceeding of 14th Canadian Conference on Artificial Intelligence, 2001.
- [11] N. Friedman and M. Goldszmidt, *Building Classifiers Using Bayesian Networks*, in proceeding of AAAI-96, Vol. 2, pages 1277-1284, 1996.
- [12] Ian H Witten and Eibe Frank, *Data Mining: Practical machine learning tools and techniques*, 2nd Edition, Morgan Kaufmann, San Francisco, 2005.
- [13] M. P. O'Mahony and B. Smyth, *A Recommender System for On-line Course Enrolment: An Initial Study*, ACM Conference on Recommender System, Minnesota USA, pages 973-978, October, 2007.
- [14] N. Chawla and K. Bowyer and L. Hall and W. P. Kegelmeyer, *SMOTE: synthetic minority over-sampling technique*, Journal of AI Research, volume 16, pages 321-357, 2002.
- [15] N. Delavari and M. R. Beikzadeh and M. R. A. Shirazi, *A New Model for Using Data Mining in Higher Educational System*, 5th International Conference on Information Technology Based Higher Education and Training, Turkey, 2004.
- [16] N. Thai-Nghe and P. Janecek and P. Haddawy, *A Comparative Analysis of Techniques for Predicting Academic Performance*, 37th IEEE Frontiers in Education Conference, Milwaukee USA, pages T2G7-T2G12, October 2007.
- [17] P. Domingos, *MetaCost: A general method for making classifiers cost-sensitive*, 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego CA USA, pages 155-164, 1999.
- [18] R. Bekele and W. Menzel, *A Bayesian approach to predict performance of a student (BAPPS): A case with Ethiopian students*, International Conference on AI and Applications, Vienna, Austria, 2005.
- [19] R. Barandela and J. S. Sanchez and V. Garcia and E. Rangel, *Strategies for learning in class imbalance problems*, Pattern Recognition, volume 36, pages 849-851, 2003.
- [20] S. B. Kotsiantis and D. Kanellopoulos and P. E. Pintelas, *Handling imbalanced datasets: A review*, International Transactions on Computer Science and Engineering, volume 30, pages 25-36, 2006.
- [21] S. B. Kotsiantis and P. E. Pintelas, *Predicting Students Marks in Hellenic Open University*, IEEE Conference on Advanced Learning Technologies, volume 30, pages 664-668, 2005.
- [22] Y. Freund and R. E. Schapire, *Experiments with a New Boosting Algorithm*, International Conference on Machine Learning, pages 148-156, 1996.
- [23] Y. Sun and M. S. Kamela and A. K. C. Wong and Y. Wang, *Cost-sensitive boosting for classification of imbalanced data*, Pattern Recognition, volume 40, pages 3358-3378, 2007.