The Enhancement of Credit Card Fraud Detection Systems

Dr. Soheila Ehramikar
Dr. TCM, M.Sc., soheila18@yahoo.com

Abstract

Along with the rise of credit card use, fraud is on the rise. In Canada, credit card fraud occurrences have been rising sharply, from $147 million in 1998 to $ 366 million in 2010, causing 150 % increase in losses in 12 years. To address fraud, financial institutions (FIs) are employing preventive measures and fraud detection systems one of them called FDS.

Although FDS has shown good results in reducing fraud occurrences, the majority of cases (approximately 90%) being flagged by this system are False Positives resulting in substantial investigation costs and cardholder inconvenience. The possibilities of enhancing the current operation by post processing the FDS output constitute the objective of this work. The data used for the analysis was provided by one of the major Canadian banks.

Based on several variations and combinations of features and training class distributions, different models (more than fifty) were developed to explore the influence of these parameters on the performance of the desired system. The results indicate that the employed approach and the prototype developed have a very good potential to improve on the existing system leading to significant savings for the FIs.

JEL Classification Numbers: 5 & 7

Keywords:

Financial Institutions, Banking industry, Credit Card Fraud losses, Credit card Fraud Detection Systems, Artificial Intelligence

I. Introduction

Credit cards are one of the most popular methods of payment worldwide and particularly in North America, due to the existence of a widespread point of sale (POS) network. Millions of people around the world use credit cards to purchase goods and services by having access to credit for a period of several weeks. Any convenient system could be abused and credit cards are no exception to this rule. Along with the rise of credit card use, fraud is on the rise.

Financial Institutions (FIs) suffer sophisticated fraudulent activities and bear millions of dollars losses every year. Based on statistics [1], fraud represents more than $1 billion annually for Visa and MasterCard worldwide. Credit card issuers and their member banks try to find new ways to prevent fraud. One of the preventive measures is extensive use of a variety of technologies, mainly Neural Networks (NNs), to track and identify suspicious transactions

and flag them for further investigation. Despite the best efforts of the FIs, law enforcement agencies, and the government, card fraud continues to rise.

At the time of this work, a neural network based software called FDS (for banking secrecy reasons) had been adapted by a large number of FIs for fraud detection. FDS scores the

transactions for the likelihood of fraud in real time. When these scores hit a threshold set by the FIs, a case is created and those accounts are passed to the fraud analysts for further follow up. Fraud analysts are security officers trained to examine the cardholder's historical behavior and by considering different factors determine the potential risk associated with the flagged accounts. Fraud investigation is a difficult task and FIs are reluctant to block an account without making sure that the transaction is, indeed, fraudulent. Very often an 'unusual' transaction is legitimate and issuers are anxious not to inadvertently offend a cardholder by acting too hastily and blocking her/his account, especially in cases where the fraud officer is unable to find the cardholders and verify the transactions with them.

FDS has shown good results in detecting fraudulent transactions, however, the majority of transactions (approximately 90%) being flagged by this system as potentially fraudulent, are in fact, *legitimate*. It should be noted that although fraud analysts, based on their experience and evaluation of the customer's history, might come to the conclusion that the activity of the flagged account is legitimate, bank policy requires them to call every individual cardholder for the verification of transactions [2]. The process of calling cardholders results in three major problems:

1. Not all the suspicious transactions are necessarily fraudulent. This type of error is referred to as *false positive* (FP) which means that the case was not fraud although it was flagged as being potentially fraudulent. The process of confirming every transaction that deviated from the cardholder's usual behavior may result in customer dissatisfaction.
2. The costs associated with investigating a large number of *false positives* are very high.
3. Currently, a substantial amount of time is being spent on investigating a large number of legitimate cases (FPs). If the number of investigation on FPs could be lowered down, fraud analysts can spent more time on real fraud cases, preventing more losses to the industry.

Therefore, any solution that refines the investigation selection process by reducing the number of unnecessary calls is welcomed by the FIs. Collaboration with one of the major Canadian banks was established to examine the potential ways of enhancing the current system. Based on information obtained from this bank, for the set threshold in 2000, FDS flags were close to 50,000 accounts per month all across Canada.

The main objective of this work was to improve the process of personal follow up on a large number of suspicious transactions. The aim was to find a way to preprocess the flagged transactions and identify the most probable legitimate transactions from the stream of legitimate/fraudulent transactions. Thus, the volume of unnecessary investigations was reduced leading to significant savings for the FI. Moreover, the current FDS threshold can also be lowered and a number of fraudulent cases, being missed under this level, can be detected. As a result, the fraud is discovered earlier and the overall losses may be reduced.

II. literature review

In 1951, Franklin National Bank of New York was the first financial institution (FI) to enter the credit card market by issuing its bank card. Within four years about 100 other financial institutions (FIs) introduced their own cards. In 1956, the BankAmericard (now VISA) entered the market followed by Master Charge (now MasterCard). In 1968, four Canadian banks introduced Chargex (now VISA) credit cards to the market and in the first day of their use two cases of fraud were committed [3].

A credit card is a special product with the following characteristics [4]:

- It provides millions of people around the world with the opportunity to purchase goods and services with access to credit for up to 51 days, depending on the posted date of the purchase, at no cost provided that the amount owing is paid back by the statement due date.
- Cardholders do not have to put up collateral against the amount they spend, therefore, it is unsecured.

In North America, credit cards are widely used in purchasing goods and services. The main reasons for this popularity are:

1. The existence of a widespread point of sale (POS) network.
2. Reducing the risk of carrying cash and the advantage of several weeks of free credit plus optional services and benefits such as Air Miles, free insurance plans, and a number of other rewards.
3. Security of funds, that is, in case of card loss or theft, the cardholder's liability at the most is $50 provided that the cardholder reports a lost or stolen card in a timely manner.

The credit card system facilitates commercial transactions and provides profits for the participating parties. The source of income for card issuers (CIs) may come from: (1) merchant user fees, (2) cardholder user fees, and (3) interest charged on unpaid balances.

In purchasing goods and services the buyer pays for a purchase by using a line of credit from the credit card issuer (CI). The CI pays the seller for the purchase, and the buyer then pays the balance on the credit card back to the CI. Since the claim presented in payment is considered a liability of the credit card issuer, this type of transaction transfers much of the risk of insufficient funds in the transaction from the seller to the credit card issuer. In order to make up for these losses, CIs determine annual fees and interest rates based on the unrecoverable amount of money incurred by these losses.

Credit Cards Transaction Process

In purchasing goods or services through credit cards, in on-line processing systems, the authorization is the essential element of the transaction processing system. The authorization process is the first level of protection against fraudulent activities and it also maintains control over the cardholder's credit limit. The authorization process begins when a cardholder uses his/her card for a transaction. The POS machine reads the magnetic stripe embossed on the back of the card which encodes the card holder's name, account number, credit limit, and the expiry date. The authorization is completed when the transaction is approved and the cardholder signs the transaction slip.

The cardholder is required to pay the total or part of the monthly statement balance. If the balance is paid back in full there are no interest charges. If the cardholder's payment is less

than the minimum amount, the credit rating of the cardholder could be affected and the cardholder may be considered *delinquent*. Another type of transaction possible by credit cards, is obtaining cash advances. In this type of transaction the interest is charged from the day when the money is withdrawn even though the balance is paid back in full on the statement due date.

For handling a huge number of daily transactions, FIs and VISA have implemented a real-time, non-stop system of computer hardware and software. This system includes the communications network among the FIs and the VISA network as well as handles the data processing and the record keeping tasks.

Credit Card Fraud

Plastic card based payment systems are booming and being used more extensively by organizations and individuals. Obviously industries with this pace of growth are vulnerable to attacks by fraudsters. In one survey [5] conducted in the United States (U.S.) in 1993, a group of 14 credit card fraudsters admitted to employing over 100 different ways of using credit cards to obtain funds illegally.

Bank card fraud losses to Visa and MasterCard alone have increased from $110 million in 1980 to an estimated amount of $1.63 billion in 1995 worldwide. The U.S. has suffered the bulk of these losses - approximately $875 million for 1995 alone. This is not surprising because 71 percent of all worldwide revolving credit cards in circulation are issued in the US. In 1994, approximately 124 million of the 193 million adults in the U.S. had at least one credit card [6]. While precise figures are not available for the credit card industry as a whole, based on credit card use of $879 billion for 1995, the estimates imply a fraud rate of between 0.1 to 0.2 percent. In the case of bank cards (MasterCard and Visa), a study done by the American Bankers Association in 1996 estimated total gross fraud losses for 1995 at $812 million versus purchases of $451 billion, implying a loss rate of 0.18 percent [7]. In 1997, credit card fraud losses for Visa, MasterCard, American Express and Discover were estimated to be around $2 billion whereas this amount in 1990 was $440 million [8].

Fraud Schemes

Unauthorized use of credit cards for acquiring goods or services is fraud. Visa and MasterCard constitute about 65 percent of all outstanding revolving credit worldwide and the substantial number of fraud occurrences is centered on one or both of these cards [9]. Most credit card fraud schemes fall into the following categories [10]. Figure 1 illustrates statistics on different types of fraud in Canada in 1998.

1. Lost / Stolen
2. Never Received Issued (NRI) (Mail theft)
3. Counterfeit
4. Telemarketing and mail-order
5. Fraudulent applications

Credit Cards in Canada

There are over 600 institutions in Canada that issue VISA or MasterCard. Among these CIs, the number of major institutions that issue VISA or MasterCard are 18; ten banks, one trust company, three credit unions, and four other financial institutions. The other CIs are affiliated issuers, such as the Bay, General Motors, University of Toronto, Petro-Canada, Eaton's, Canadian Tire, and so on. Since more than 64% of Canadians pay their credit card balance in full each month, the interest rate for two-thirds of credit card users is zero. There are more than 70 low interest rate cards on the market and over 40 of those cards have an interest rate of under 12%.

Figure 1 Statistics on different types of fraud in Canada [10]

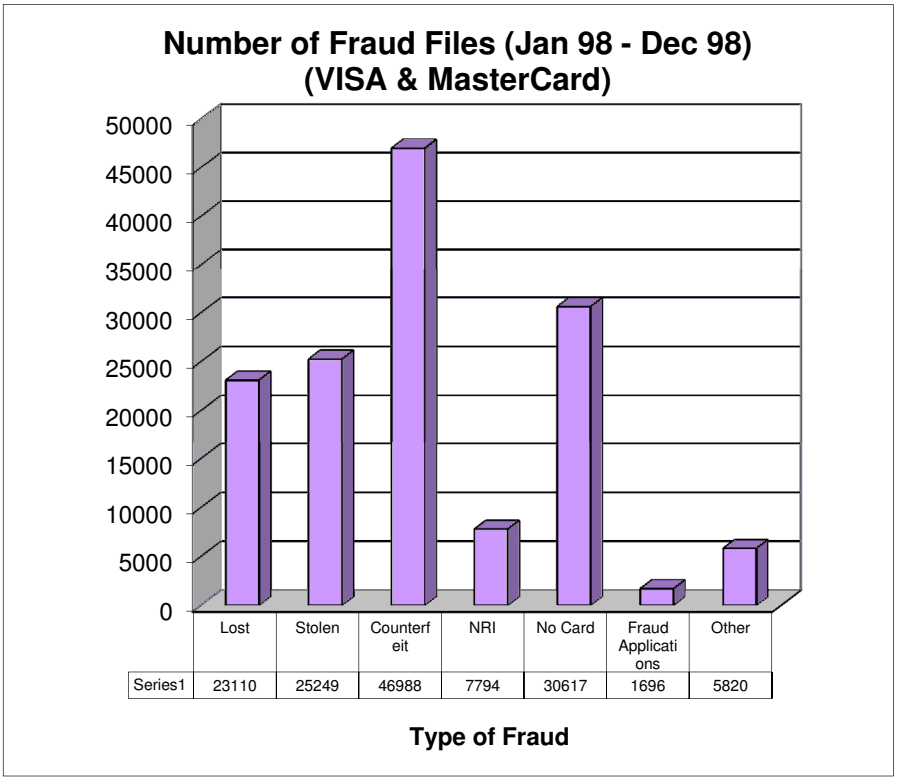

Statistics on Credit Card Fraud

Based on statistics reported by the CBA, credit card fraud occurrences rose sharply in fiscal 1998 (141,274) compared to 1997 (113,264). Based on the information obtained from this report, 34 percent of all credit card fraud occurrences and 50 % of the $147 million written off in 1998 was due to counterfeit card fraud. This report also indicates that approximately 50 % of Canadian credit cards which were used fraudulently, were used outside of Canada.

Current statistic indicates that there are an estimated 71.3 million Visa and MasterCard cards in circulation in Canada[11]. Total outstanding credit-card debt hit $78 billion in September 2009, up from $76 billion in September 2008 [12]. Ninety-day credit card delinquencies jumped 53 % between September 2008 and September 2009, hitting $3.6 billion. Among major cities, Toronto has the nation's highest delinquency rate (2.14 %) in October. The

average national rate is 1.67 % [12]. When fraud occurs, customers have zero liability. In 2010, financial institutions reimbursed more than $365 million to their Canadian credit card customers, representing the losses these customers suffered as a result of criminal activities. The credit card delinquency rate in Canada is half of what it is in the US. Canadians carry an average of two credit cards per household, while US consumers carry an average of six [11].

FDS and Credit Card Fraud

As discussed, FIs make extensive use of NN based software to spot and flag transactions inconsistent with the cardholder's usual behavior. The focus of attention in this work is FDS,

a NN base software being used by 40 of the top 50 [13] large credit card issuers worldwide including our collaborating FI.

FDS is a real-time customized software designed to determine the likelihood of card fraud. By using legitimate and fraudulent transactions, FDS has built an individual behavior profile for each account. Due to the proprietary and business concerns of the software provider, no public documentation is available on the software. Therefore, it is not clear how this profile is established but the conjecture is that the account profile file includes the type of merchant at which the cardholder typically shops, the time of the day that the cardholder normally makes purchases, the geographic locations along with many more characteristics that only software developers are aware of. FDS inspects and evaluates the incoming transactions to see if they fit into the customer's established profile. Any deviation from the usual cardholder's behavior is monitored and scored by this system. Based on the changes that FDS detects in the customer's pattern of behavior, scores between 1 and 1000 is assigned to each transaction. The higher the score is, the higher the likelihood of fraud.

Bank authorities set a threshold value and all transactions scored above this threshold are considered suspicious so that when these scores hit the set threshold, a case is created and is flagged for further investigation. Inherently, FDS makes no assumption about the suspicious transactions and transmits the flagged accounts, in real time, to the FI's fraud department for further follow up and investigation [2].

Fraud Detection Dilemma

Credit card fraud detection is a pattern recognition problem. Every cardholder has a shopping behavior which establishes a profile for her/him. As the result of personal needs or seasonal reasons, patterns of behavior change over time so that s/he may develop new patterns of behavior, which are not known as yet by FDS. Very often an 'unusual' transaction might be legitimate. Currently, FDS identifies many legitimate accounts as fraudulent resulting in a large number of false positives (FPs). As every cardholder has a huge number of possibilities for developing new patterns of behavior, the types of transactions are widely variable. In fact, there are so many variations of behavior for each individual that are exponential in combination and the complexities of enumerating all combinations of cases are enormous. Hence, it is almost impossible to identify consistent and stable patterns for all the transactions. This ever-changing pattern of behavior along with the combination of legitimate and fraudulent cases has left the FIs with a large number of FPs (approximately 90% of flagged accounts) that has to be investigated.

The motivation of this work was to post-process the FDS output and to identify the legitimate transactions (True Negatives, TN) from the stream of flagged transactions. This identification is a classification task, that is, the system we develop has to be able to extract the True Negatives (TNs) from the pool of data while not missing fraudulent transactions. If this goal could be achieved, the bank staff may not need to call the legitimate customers for transaction verification. Pattern recognition for these occurrences is inherently complex and one has to understand the underlying system as much as possible and use this knowledge in the design of the required system. Investigation of some of the AI methodologies and their application revealed that learning is the appropriate approach for addressing this type of classification problems. In fact, learning is very much appropriate for cases where patterns of behavior in real world problems are complex and there is little or no knowledge of the semantics of the

application domain. A further survey on some of the learning methodologies and their application led to learning decision trees methodology for this work.

III. Methodology

Decision trees, a machine learning method, are perhaps the oldest, and one of the most popular ways to represent the outcome of classification learning procedure. It is a method for approximating discrete-valued target functions, in which the learned function is represented by a decision tree [14]. Decision trees are capable of representing the most complex problems given sufficient data, and they are one of the most highly developed techniques for partitioning samples into a set of decision rules. Learned trees can also be represented as sets of if-then rules to improve the human readability. These learning methods are very popular and have been successfully applied to a broad range of tasks from learning to diagnose medical cases to learning to assess credit risk of loan applicants [15].

Although a variety of decision tree learning algorithms have been developed with somewhat different capabilities and requirements, decision tree learning is generally best suited to problems with the following characteristics [15]:

- The target function has discrete output values. For instance, decision tree assigns a ‘yes’ or ‘no’ to each classified example.
- The training data may contain errors. Decision tree learning methods are robust to errors found in the attribute values that describe the input examples.
- The training data may contain missing attribute values. Even though the value of some of the training examples might be unknown, still decision tree learning methods can be employed.

Decision tree learning is easy to implement, to understand and to display and based on the above characteristics is a suitable choice for this work.

IV. Data and Data Transformations

To perform the analysis, the accounts flagged by FDS were used as the input of the learning system. The data was provided by the collaborating FI. The transactions flagged by FDS are taken over 45 days (June, July, and part of August 99) and are related to a limited region of Toronto. Together, ten separate files were provided. The first nine files were related to flagged confirmed *legitimate* accounts, which together consisted of 4919 accounts with 69,182 transactions. Due to the volume of data for the legitimate accounts, they were divided

into nine separate files. The tenth file included 707 *fraudulent* accounts that contained 6,725 transactions. The fraudulent accounts have a combination of fraud/non-fraud transactions consisting of 1,743 legitimate and 4,982 fraudulent transactions. Due to the confidentiality of real account numbers and in order to have all the transactions from each account together, a substitute but unique number was assigned to the transactions of each account by the FI. All ten files had the same fields and each transaction had the information depicted in Table 1.

Although the scores associated to each transaction by the FDS, were of great importance for the analysis by the learning system, due to the proprietary and business concerns of the software provider, the FI was not able to provide this information. It was essential to identify which transactions deviated from the normal behavioral pattern of the legitimate cardholders which caused the system to flag them as potentially fraudulent. The FDS scores was the

indicator of this trend. To make up for this data shortcoming, the case creation date was provided as a proxy to each transaction. Lack of scores not only may have serious impact on the precision of the classifier, but also due to the high volume of data, it caused uncertainty and substantial amount of 'manual' work in selecting the transactions that occurred close to case creation date.

To use the data for training, it was necessary to identify the fraudulent transactions from the legitimate ones. At the time, labeling the fraudulent transactions was done manually and the fraud investigation department kept conventional paper based fraud files on which they mark the transactions that were identified as fraud. Due to this manual process, there was no mechanism to migrate this information back into the transaction tracking system and, therefore, there was no record keeping of them on the system. Fraudulent accounts normally have a mixture of fraudulent and legitimate transactions, therefore, the confirmed fraud transactions in fraud file were labeled by the bank with an asterisk (*).

Preprocessing the Databases

Raw data contains the information that must be extracted but in the meantime it contains too much non-essential information. The raw data provided by the FI, required substantial preprocessing to weed out the irrelevant information and to prepare the data set in a suitable form for the learning system. The original data files were in text format, therefore, Excel was selected as a tool for data manipulation. Manual inspection of data revealed the existence of some inconsistencies in the data. To be able to use the data for the analysis these inconsistencies had to be removed from the datasets. After the preprocessing of databases, the final legitimate database consisted of 13,426 non-fraud transactions and the final fraudulent database consisted of 6,666 transactions (4,969 fraud and 1,698 non-fraud).

Learning Requirements

Learning means behaving better as the result of experience. The task of a learning system is to extract the maximum amount of information from the data samples, and based on this information, to estimate the accuracy of its future classifications and predictions. While conceptually simple, extracting information from a large database requires careful organization and the specification of the goals to be met by the learning system. The simple requirement of the classification methods is that the data be presented in the form of samples composed of patterns of observations with the correct classification. Then the learning procedure will be applied which is an iterative process[16].

Features and Classes

In the problem of predicting whether a flagged transaction is fraud or non-fraud, there are two classes: fraud and non-fraud. The task is to predict which is the correct class based on the observations of a set of transactions. By employing a decision tree learning algorithm, the aim is to learn a definition for the concept, *Transaction* (fraud/non-fraud), where the definition is expressed as a decision tree. In setting this up as a learning problem, the properties or features that are available to describe the examples are presented in Table 1.

Table 1 Credit card observations

Feature	Description
Account No.	Cardholder's account number
Date/Time	Date and time of transaction
Dollar	Dollar amount of transaction
SIC	Merchant category code
Country	Merchant country code
Decision	Authorized (A), Declined (D), Referral (R), Pick up (P)
POS	Card swiped (S) / keyed (K)
Card type	Classic / Gold
Case creation	The day / time case created by FDS
Case action	The day / time fraud analyst investigated on the case

Software Selection

If learning is viewed as a search problem, then it is natural that learning algorithms will examine different strategies for searching the hypothesis space. The algorithms that are capable of efficiently searching very large hypothesis spaces to find the hypotheses that best fit the training data are of great interest. In the field of ML, a variety of programs have been developed. After investigation, See5 [17]which is able to produce decision tress and to transform the generated trees into a set of rules was selected for this analysis.

Distribution Design for Training and Testing Sets

For setting up the experiments with different class distributions, the first step was to identify the class distribution of the selected databases. By taking into consideration that the fraud database had the combination of fraud and non-fraud transactions, a simple calculation revealed that the class distribution of the training and testing databases were 25:75 and 21.9:78.1, respectively. In order to train on data that has an artificial higher fraud rate and observe the result of the constructed classifiers, the following distribution and partitions, shown in Table 2, were formed for the training phases.

Table 2 Design distribution for training and testing sets

	Training cases		Testing cases	
Training Class Distribution	Non-fraud Transactions	Pure fraud Transactions	Non-fraud Transactions	Pure fraud Transactions
25:75	10,054	3,351	5,049	1,416
33:67	6,708	3,351	5,049	1,416
50:50	3,368	3,351	5,049	1,416

Experiments

As Table 2 illustrates, sets of training examples with different class distributions were designed and used as the input of See5 to generate the classifiers and observe the effect of class distribution on training, testing and prediction of new cases. The choice of features can also affect the performance of the learning systems. Therefore different combinations were considered in the classifier construction procedure to discover the most effective combinations. The feature combinations examined are: (1) all features, (2) all features except for the card type, and (3) all features except for the POS and card type. Altogether fifty four experiments were performed to study the effects of class distributions on training and variations of different features on the evaluation of the classifiers constructed. The experiments were conducted using See5 construction options of: (1) decision trees, (2) rulesets, (3) boosting, and (4) tenfold cross validation (CV).

For each option, the program was run to explore the effects of various class distributions and features on training and testing sets. Although the cross validation technique is typically used for intermediate sample sizes (of order 2000)[18], however, to have extra evaluation on the training and testing sets, in the experiments conducted, this option was examined too.

V. Results

The experiments produced 54 sets of results. Thirty six of these results are the classifiers and the rest present the evaluation of tenfold CV trials on training and testing sets. For banking secrecy reasons, only a summary of the results obtained is presented in this paper

Performance Analysis

To do the analysis the classifier attained the lowest error rate among the 36 classifiers was selected. As the result shows, the *boosted decision trees* (BDT) classifier trained on 25:75 class distribution attained the lowest error rate of 11.4%. This classifier was selected as the first choice. To compare the performance of this classifier against another one, the *decision trees* (DT) classifier trained on 25:75 class distribution by attaining the error rate of 13.5% was considered as the second choice.

As the fraud rate increases in the training sets, the error rate also increases leading to the conjecture that there is no need for further analysis on the class distribution with higher fraud rates and the above discussed classifiers are the most effective classifiers for further analysis. However, based on the work of other researchers [6,19,20] who have employed various training class distribution in their analysis for fraud applications, it was decided to study the discussed

BDT and DT classifiers not only for 25:75 class distribution but also for the class distributions of 33:67 and 50:50.

As discussed, in situations where different types of errors have different costs such as in credit card fraud detection, the elements of the confusion matrix such as TN, FP, FN, and TP are the essential metrics for the system performance. Therefore, these metrics were considered to be the true indicators for the performance evaluation of the selected classifiers.

Table 3 Two-class classification performance

Cases	Prediction Negative *(Legitimate)*	Prediction Positive *(Fraudulent)*
Class Negative *(Legitimate)*	True Negatives (TN) *(Normal)*	False Positives (FP) *(False alarm)*
Class Positive *(Fraudulent)*	False Negatives (FN) *(Miss)*	True Positives (TP) *(Hit)*

To compare these metrics for the selected BDT and DT classifiers, a new set of Tables were formed. Tables 6-1 to 6-6 and 6-7 to 6-12 are related to the evaluation of the selected BDT and DT on the training and testing sets. For banking secrecy reasons, only a summary of the results obtained are presented below. Each Table depicts error rate, TN, FN, TP, FP and their associated rates for each class distribution and the sets of features. TN, TP, FN, and FP rates were calculated based on the information available from the confusion matrix of each classifier. Figures 2 and 3 depict the performance of these classifiers. Figure 2 illustrates the performance of the BDT and DT classifiers on the training data. Figure 3 illustrates the performance of the BDT and DT classifiers on the testing data. In these charts, the x-axis represents the percentage of fraud rate in the training set whereas the y-axis represents the FN rate of the classifiers.

Figure 2 Evaluation of BDT and DT classifiers on the training data

Figure 3 Evaluation of BDT and DT classifiers on the testing data

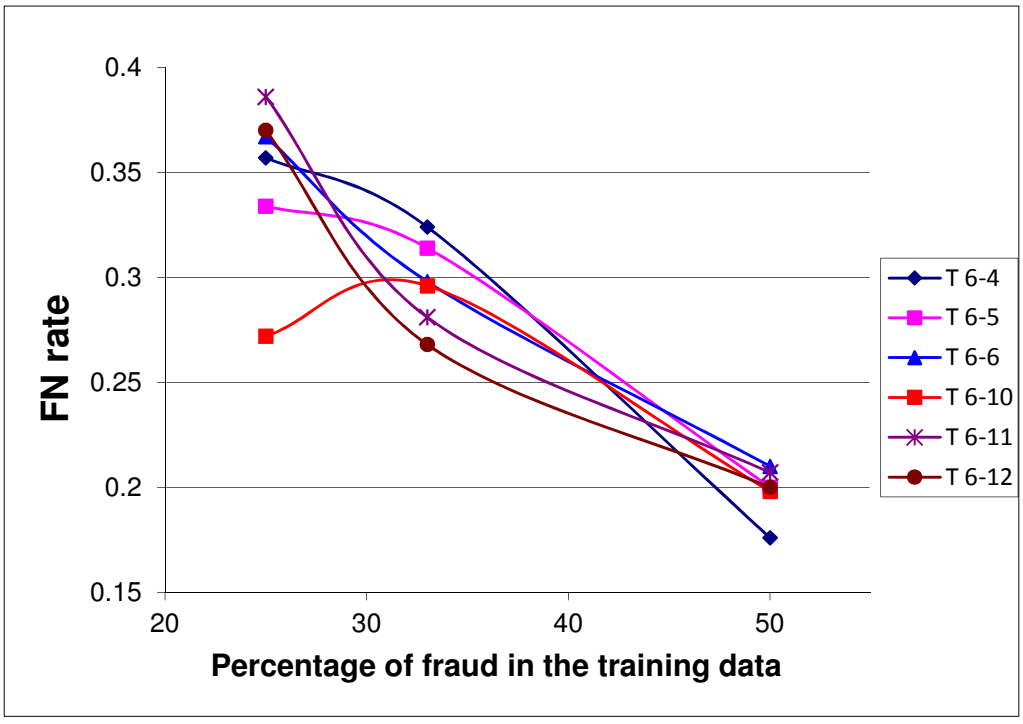

VI. Conclusion

Fifty four experiments were conducted to determine the most predictive classifier. These experiments were performed based on several variations and combinations of features and training class distributions. The evaluation of classifiers, constructed from different sets of experiments, was different on training and testing data confirming that significant attention has to be paid in the class distribution design of the training sets. The performance metrics considered for this analysis were True Negative (TN) and False Negative (FN) rates. The BDT classifier trained on 50:50 class distribution attained a TN rate of 87.4% and 75.3% and FN rate of 24.5% and 17.6% on training and testing data, respectively. Based on this performance, this classifier considered to be the most predictive classifier for this work by having the lowest possible FN rate among all the other classifiers constructed in this analysis.

This analysis reaffirms the importance of training class distribution in the design of the effective classifiers. This study shows that increasing the number of minority instances in the training data will produce classifiers with improved performance. It also confirms that increasing the number of majority instances in the training data will produce classifiers that are adept at classifying the majority of transactions as legitimate and as a result, these classifiers classify a large number of fraudulent cases as legitimate leading to very high FN rate.

The performance of the BDT classifiers on the prediction of new cases was also examined. This analysis showed that the classifier trained on 25:75 distribution of fraud/legitimate transactions attained the TN rate of 98.8% in the prediction of legitimate cases. However, the

performance of this classifier degraded on the identification of fraudulent cases so that the classifier identified half of the fraudulent transactions as legitimate, attaining a FN rate of

49.7%. The degradation in performance makes the system unusable because missed fraud cases are very costly. The classifier trained on 50:50 distribution had lower TN rate (92% against 98.8%) on the prediction of legitimate transactions, however, its FN rate on the prediction of fraud cases was very much lower (26.8% against 49.7%) than the comparative BDT classifier. This analysis reaffirms that classifier trained on 50:50 class distribution is more predictive for the evaluation of new cases.

The other important factor which may have a serious impact on the performance of the classifiers was the limitations of the data sets. The most important limitations were rather small fraud database and the lack of FDS scores associated with the flagged transactions. These scores are an indication of some patterns of behavior in the datasets and contain valuable information. The result of the experiments conducted on the variations of features revealed that the classifiers trained on all features performed much better than the ones trained while disregarding some features such as POS and card type. Based on these empirical results, one would expect that if the FDS scores were provided, they would contribute important information thus leading to better performance result.

There was no information available on the cost of investigation associated with every case created by the FDS, therefore, the savings from the use of the system trained could not be estimated. Based on the observed results on the prediction of new cases, one could expect

that this approach may reduce the volume of personal investigations leading to potentially significant savings for the FI.

Presently, due to the high volume of false positives flagged by the FDS, the FI has set a rather high threshold for this system. Therefore, there are cases that are fraudulent but are being missed (FN) by the FDS. By instituting a post-processor system such as the ML, the FI has the option of lowering the threshold and allowing FDS to flag more cases for investigation. Another important point is the prevention of unnecessary disturbance of the customers which may lead to customer dissatisfaction.

In summary, pattern recognition for legitimate/fraud occurrences is inherently complex and since legitimate cardholders'/ fraudsters' patterns of behavior evolve over time, this work is a basis for further study. Overall this study demonstrates that the approach employed in this work, has a very good potential of identifying the legitimate transactions from the fraudulent ones but there is a need for the enhancement of its predictive accuracy by obtaining the mentioned missing information.

References

[1]: Isabelle Sender; *Detecting and combating fraud*; Chain Store Age; New York; Vol. 74; Issue 7; Page 162; July 1998.

[2]: Elford Dean, Raj Thomas, Lorry; Visa security center; Personal meetings; January and February 1999.

[3]: Donald V. Macdougall, Richard G. Mosley, Garioch J. I. Saunders; *Credit card crime in Canada: Investigation – Prosecution*; The Canadian Association of Crwon Counsel; page 1-56; January 1985.

[4]: Canadian Bankers Association; *Fast Facts – Credit Cards*; June 1999. http://www.cba.ca/eng/Statistics/FastFacts/visamc.htm

[5]: Russell Smith; *Cards Games: Plastic fraud and misuse*; Australian Accountant; Melbourne; Vol. 67; Issue 11; Page 56-58; December 1997.

[6]: S. Stolfo, W.Fan, W.Lee, A. Prodromidis, and P. Chan; *Credit card fraud detection using meta-learning: Issues and initial results*; Work notes AAAI- 97 workshop on AI approaches to Fraud Detection and Risk Management; 1997.

[7]: William Roberds; *The impact of fraud on new methods of retail payment*; Economic Review; Federal Reserve Bank of Atlanta; Atlanta; Vol. 83; Issue1; Page 42-52; First Quarter 1998.

[8]: Peter Hadfield; *Stripe makes credit card fraud tougher;* USA Today; Arlington; July 14, 1998.

[9]: Keith Slotter; *Plastic Payments: Trends in credit card Fraud*; FBI Law Enforcement Bulletin; Washington; Vol. 66; Issue 6; Page 1-7; June 1997.
[10]: Canadian Bankers Association; *Fast Facts – Credit Card Fraud*; June 1999. www.cba.ca/eng/Statistics/FastFacts/credit_card_fraud.htm
[11]: Canadian Bankers Association; August 19, 2011 ww.cba.ca/en/media-rroom/50-backgrounders-on-banking-issues/123-credit-cards
[12]: Equifax Canada; Source: Toronto Star, December 2009.
[13]: Trudy Ring; *Fraud Detection and More;* Credit Card Management; New York; Vol. 10; Issue 6; Page 128; September 1997.
[14]: Sally Jo Cunningham, Matt Humphrey, Ian H. Witten; *Understanding what machine learning produces, Part I: Representations and their comprehensibility*; Department of computer Science, University of Waikato.
[15]: Tom M. Mitchell; Machine Learning; WCB McGraw-Hill; 1997.
[16]: Sholom M. Weiss, Casimir A. Kulikowski; Computer Systems that Learn; Morgan Kaufmann; 1991.
[17]: J. R. Quinlan; C4.5: Programs for Machine Learning; Kaufmann, San Mateo; 1993
[18]: J. R. Quinlan; See5: *An informal tutorial*; August 1999. www.rulequest.com/win.html
[19]: P. Chan, S. Stolfo; *Toward scalable learning with non-uniform class and cost distributions: a case study in credit card fraud*; Proceedings of fourth International conference on knowledge discovery and data mining; p164- 168; 1998.
[20]: T. Fawcett, F. Provost; *Adaptive fraud detection*; Data Mining and Knowledge Discovery 1; 291-316; 1997.