

Premium Rate Services Fraud Detection

Mariella Rivas*, Richard Roach**, Patrick Hosein*

The University of the West Indies, St. Augustine, Trinidad and Tobago*

Telecommunications Services of Trinidad and Tobago, Port of Spain, Trinidad and Tobago**

mariellarivas93@hotmail.com, rroach@tstt.co.tt, patrick.hosein@sta.uwi.edu

Abstract—Premium Rate Services (PRS) destination fraud occurs when telephone calls are made to high-cost (premium) destinations through fraudulent means. Early detection and termination of such calls can significantly reduce the cost they incur. We investigate various features of such calls to determine which can be used to quickly predict and terminate them. This is achieved through feature selection and clustering analyses followed by fraud detection using a Classification and Regression Tree (CART) model. Additionally, we propose the use of costs/benefits to evaluate the model's performance through the consideration of an adjusted decision tree cost function at the training level. This approach is compared to the traditional decision tree to highlight its advantages.

Index Terms—Fraud Detection, Machine Learning, Network Management, Premium Rate Destination

I. INTRODUCTION

Premium Rate Services (PRS) Destination Fraud arises from the fraudulent occurrences of outgoing calls from a telecommunication network to high-cost destinations. Such fraud can result in significant revenue leakage and hence early detection and prevention is important [1], [2]. We use a case of a telecommunication company to illustrate a proposed approach for early detection and termination of such calls.

Profiling fraud is time-consuming and it is difficult to detect in advance certain PRS destinations, which in turn increases the associated risk of financial loss. In fact, companies are often times made aware of suspicious/fraudulent activities by account managers and customers after the impact has reached their bills. Moreover, they must absorb these losses in an effort to maintain their customer base. We present a model that can detect PRS fraud more proactively, so as to reduce the high costs incurred to the company and its customers.

This paper first outlines the steps taken to achieve the desired goal by using historical fraud cases to determine the most influential call level features that the company may use to better identify PRS in the future. This step, referred to as feature selection, will be particularly useful in ensuring that the most important indicators of PRS are indeed observed when attempting to identify a PRS case. It will also reduce the time taken for a model to flag an activity as possible fraud by protecting the model from being over-fitted with unnecessary features. We then perform a cluster analysis using the selected features, then describe two prediction methods (traditional approach and cost-sensitive approach) and apply and evaluate these on a real dataset.

II. RELATED WORK AND CONTRIBUTIONS

Several papers have been written on telecommunication fraud detection (see [3] for an overview) but the situation varies by country. Popular, successful approaches to this issue have explored the application of several machine learning algorithms to a variety of features such as telephone numbers, call times, call duration, called networks to detect fraudulent call activities (for instance, [4] [5] [6] and [7]). Accounting for the costs/benefits related to fraud detection is also of high value, as made evident in [8], which focused on the design of Key Performance Indicators (KPI) to optimize revenue assurance and manage fraud. The authors in [9] and [10] also highlighted the advantages of considering costs/benefits in classification, by applying cost-sensitive machine learning methods to medical data.

We focus on the case of a Small Island Developing State with a sophisticated telecommunication infrastructure but with limited resources to monitor and manage fraud. Over the last two decades, the issue of telecommunication fraud has sparked international concern [1] and has recently become a more serious issue especially for developing countries [2]. In this paper, we initially conduct feature selection analysis to outline the most influential call features to be extracted and used in fraud detection. This step was intended to allow more efficient model training, saving both time and financial resources, as illustrated in [11] and [12]. We then perform Ward's Hierarchical Divisive Clustering (WHDC) algorithm on a sample of the data, using the selected features to further reveal any underlying patterns in the data that may have assisted in model training. The WHDC method was specifically chosen due to its efficiency regarding large datasets. The use of hierarchical agglomerative clustering in telecommunication fraud detection is also outlined in [7]. Next we trained and tested a CART decision tree and a radial-kernel Support Vector Machine (SVM) to detect fraud, given values for the selected call features. Similar work involving decision trees is presented in [4], [13] and [14], while [5] and [15] explored SVM approaches to the problem. Lastly, we explored a cost-sensitive CART decision tree to account for costs/benefits associated with different outcomes in fraud detection and compared this approach to the traditional approach aforementioned. Cost-based analyses have been less popular in the telecommunication industry but the approach has been successfully applied in other areas as detailed in [16], [17], [9], and [10].

TABLE I
LIST OF CALL LEVEL INPUT VARIABLES

Variable Name	Type	Example
ACCOUNT_NO	Numerical	2.867409e+13
SERV_NO	Numerical	224145067
DIALLING_NO	Numerical	224145067
DIALED_NUMBER	Numerical	6228332367
CALL_DATE	Date	2018-06-03 08:50:30.4
ACTUAL_DURATION	Numerical	0.2
IM_CHARGE_TOTAL	Numerical	0.8
LEAD_DIALED_NUMBER	Numerical	18686718140
LEAD_CallDate	Date	2018-06-05 14:07:07
CallTime_Plus_Duration	Date	2018-06-03 08:50:42.4
groups_consecutive_dialed_no	Numerical	542
NumberOfCalls_PerDialled_Daily	Numerical	5
ServNo_NoOfDailyCalls_ToDest	Numerical	3
ServNo_DailyRev_ToDest	Numerical	13.60
ServNo_DailyMOUs_ToDest	Numerical	11.25
NumberOfCalls_PerGrouped	Numerical	130414
ServNo_NoOfDailyCalls_ToDest_avg	Numerical	8.02
ServNo_DailyRev_ToDest_avg	Numerical	14.08
ServNo_DailyMOUs_ToDest_avg	Numerical	4.97
Consecutive_dialed_no_with_non_matching_duration	Numerical	1024
TOS_CAT	Categorical	N
DayOfWeek	Numerical	6
TimeOfDay	Numerical	1

III. DATA DESCRIPTION

The data examined in this analysis consists of 34 explanatory (input) call level variables for individual instances of PRS fraud over a two-year period. However, we found that 11 variables (also called features or attributes) had several missing values and so the number of variables was reduced to 23 from 34. These 23 variables are outlined in Table I. Together this gave a rich sample dataset of call level information for 199,771 fraudulent and non-fraudulent PRS calls. The occurrence of PRS fraud was taken as a binary response (output) variable with occurrence being denoted by 1 and non-occurrence denoted by 0. For the purposes of the initial feature selection analysis and model development, we used all available cases. The sampled data was split 75%-25% between training and testing respectively. However, a representative random sample of 1,000 cases was employed for the cluster analysis stage, to provide ease of visualization.

IV. FEATURE SELECTION

All mathematical and statistical work conducted for this analysis was done using the R Statistical Software. After removing missing data, the 4 non-numerical variables were easily converted to their numerical equivalents to facilitate




Fig. 1. Information Gain of each Feature

easier, more accurate analysis. Next, the `FSelector` R library was used to fundamentally assess the information gain of the input variables. Information gain measures the reduction in entropy by splitting the dataset based on given values of a random variable, and is commonly used in feature selection and construction of decision trees. A variable's information gain usually lies between 0 and 1 inclusive, with a higher value being preferred, as this minimizes entropy and aids in effective classification. The information gain achieved for each input variable is provided in Figure 1.

After this, the `caret` and `mlbench` R libraries were used to calculate the correlation between each pair of input variables. This allows a better understanding of the relationships between the variables and aids in the identification of redundancies. A variable was considered redundant if its correlation to another variable was higher than 0.80. Feature

TABLE II
HIGHLY CORRELATED FEATURE PAIRS

Feature 1	Feature 2	Corr
SERV_NO (abbreviated SN)	DIALLING_NO	1.00
SN_ToDest	SN_DailyMOUs_ToDest	0.84
SN_DailyRev_ToDest	SN_DailyRev_ToDest_avg	0.85
SN_DailyMOUs_ToDest	SN_DailyMOUs_ToDest_avg	0.86
SN_NoOfDailyCalls_ToDest	NumberOfCalls_PerDialled_Daily	0.98
SN_NoOfDailyCalls_ToDest	SN_NoOfDailyCalls_ToDest_avg	0.89
Consecutive_dialed_no_with _non_matching_duration	NumberOfCalls_PerGrouped	0.81

pairs with high correlation values are provided in Table II.

From Figure 1, it is evident that the date of occurrence, dialling and dialled numbers, and variables relating to frequencies of calls made are the most critical features in determining a fraudulent PRS call. They are considered highly influential, whereas the remaining variables are only fairly influential.

Moreover, a closer examination of the results in Table II shows that a server number's daily minutes of use (MOUs) to a destination is redundant to its daily revenue to a destination, and so on. In fact, examining column 2 of Table II alongside the results in Figure 1 led to 5 of the 23 input variables being considered redundant and therefore withheld from further analyses. This leaves 18 variables to be selected for inclusion in our analyses, ACCOUNT_NO, SERV_NO, DIALLING_NO, DIALLED_NO, CALL_DATE, ACTUAL_DURATION, Call-Time_Plus_Duration, IM_CHARGE_TOTAL, LEAD_DIALED_NO, groups_consecutive_dialed_no, NumberOfCalls_PerDialled_Daily, NumberOfCalls_PerGrouped, ServNo_NoOfDailyCalls_ToDest_avg, ServNo_DailyRev_ToDest_avg, ServNo_DailyMOUs_ToDest_avg, DayOfWeek, TimeOfDay and TOS_CAT.

V. CLUSTER ANALYSIS

Next we provide a clustering analysis on the selected 18 features for both PRS fraud and non-fraud cases. This will allow a more detailed understanding of how each feature interacts with the response variable of fraud occurrence and set a solid foundation for the development of a model that detects PRS fraud more proactively and efficiently. To conduct this analysis, the original dataset was truncated to a representative random sample of 1,000 PRS calls in an effort to better visualize results. In this sample, 787 calls were non-fraud while 213 of them were fraudulent. The binary dependant variable (occurrence of fraud) was withheld from the algorithm as necessary, and numerical equivalents deduced earlier for date and categorical variables were used as replacement values in the data to be examined.

We use Ward's Hierarchical Divisive Clustering method with Euclidean distances under complete linkage. Theoretically, this method builds a top-down tree-like hierarchy of



Fig. 2. Cluster Plot of Sampled Individual PRS Calls

clusters (data groups) by beginning with all the data points in one initial cluster and splitting clusters recursively until each data point has been assigned to a single cluster. It is particularly useful for large datasets.

Different R packages were used to apply this method, each generating different visualizations that broadened perspective and facilitated better interpretation of the results. Firstly, the `cluster` R library was used to calculate the divisive coefficient, which measures the amount of clustering structure found, with values close to 1 being preferable. In our case, it was satisfactorily found to be 0.97. The complete linkage dendrogram on individual data points and the corresponding silhouette plot were also visualized using this library. Additionally, the `factoextra` R library was used to create a mapping of individual data points within clusters, and produce an elbow plot to confirm a suitable number of clusters for the data. The cluster plot is provided for reference in 2. It shows a mapping of the data into 2 distinct clusters, one larger than the other. This coincides with the separation of the clusters as generated by the other visualization methods used, as well as the actual split of the calls into fraud and non-fraud categories. Thus, we can safely deduce that it is efficient to categorize PRS calls into either of 2 groups, which is desirable as the response variable of fraud occurrence is, intuitively, indeed binary.

VI. FRAUD DETECTION METHOD

We investigated two basic supervised learning approaches so as to determine the accuracy of each model. After the conversion of the non-numerical input variables, the dataset of 199,771 cases was partitioned 75% - 25% as a training set and testing set respectively. This resulted in data for 149,828 PRS calls to be used to train the models, and 49,943 calls to be used for testing for the 18 input variables being considered. Two simple supervised learning models were evaluated, a classification and regression decision tree (CART DT) and a radial-kernel Support Vector Machine (SVM). The CART DT was chosen due to its transparent, comprehensive nature and its ability to provide the researcher with ease of use. The SVM approach was also selected due to its proven

TABLE III
CONFUSION MATRIX FOR CART DT

Predicted \ Actual	0	1
0	41098	2511
1	2634	3700

TABLE IV
CONFUSION MATRIX FOR SVM

Predicted \ Actual	0	1
0	35850	4496
1	5842	3755

power in addressing classification and prediction problems. The Decision Tree for the CART model is shown in Figure 3. It is important to note that such visualizations were not available for the SVM due to its more complex nature.

Looking at the decision tree in Figure 3 suggests the base node of the tree is call date and time, implying that this is the most influential call variable in initially separating PRS calls into fraud or non-fraud groups. Its relationship with the other input variables to achieve such categorization is shown at the middle to lower nodes of the tree. The probabilities in darker blue represent the outcome of a PRS call being fraudulent, given the course through the nodes to arrive at the particular outcome. Those in lighter blue represent the outcome of a PRS call being non-fraudulent.

Following training, both models were employed to predict fraud outcomes for the test set in order to assess their predictive abilities. Firstly, the confusion matrix of each model was constructed (shown in Tables III and IV). The confusion matrix for the CART DT indicates that the model made accurate predictions for 44,798 out of the 49,943 test cases, but mis-classified 5,145 cases. This resulted in a predictive accuracy of the model being expected to predict PRS fraud correctly 89.7% of the time. The area under the decision tree's Receiver Operating Characteristic (ROC) Curve, called its AUC, was found to be 0.91. Together with the model's predictive accuracy, this implies that the CART DT performs satisfactorily well, but not without room for improvement.

On the other hand, the SVM method produced a confusion matrix showing correct predictions for 39,605 of the test cases and 10,338 mis-classified cases. This gave a predictive accuracy of 79.3% for the model, which is fair but noticeably lower than that of the CART DT. The SVM's AUC value associated with its ROC curve was also lower than the CART DT's, with a value of 0.84. Generally, when compared, the CART DT seems to be more accurate in predicting fraud outcomes for given PRS call data. However, both models can be subjected to further tuning and improvement of their fit and predictive abilities.

Note that in fraud detection we have two types of errors,

False Positives and False Negatives. In the above analysis the accuracy was determined based on the assumption that the samples were correctly tagged. If a call was tagged as fraudulent but after further investigation was found not to be so then that tag would be removed. Therefore we believe that fraudulent tags would typically be correct. However, if a fraudulent call was missed then its tag would continue to be incorrect. The prediction algorithm would potentially pick this case up as fraudulent but later this would be listed as being mis-classified. With time this problem will be reduced because further investigation would be performed on the flagged calls and hence those that were missed should be captured.

VII. COST SENSITIVE CLASSIFICATION

In the above analysis the objective is to minimize error rate and false positive errors are treated the same as false negative errors. However this is not the case here. If a false positive error is made then a call may be incorrectly terminated resulting in poor service to the customer. On the other hand, if a false negative error is made then this can result in significant financial loss to the company. In this section we assign costs associated with each outcome of True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN) and then make predictions with the objective of minimizing total cost (as opposed to minimizing error).

Let us denote the outcome probabilities by $P_{tp}, P_{tn}, P_{fp}, P_{fn}$ with associated costs of these outcomes being $C_{tp}, C_{tn}, C_{fp}, C_{fn}$. For a given model the expected cost of the outcomes is given by

$$\bar{C} = P_{tp}C_{tp} + P_{tn}C_{tn} + P_{fp}C_{fp} + P_{fn}C_{fn} \quad (1)$$

The objective is to provide a classifier that minimizes $\bar{C}$. Using the fact that $P_{tp} = 1 - P_{fn}$ and $P_{tn} = 1 - P_{fp}$ we can rewrite this as

$$\bar{C} = (C_{fn} - C_{tp}) \left\{ 1 + \left\{ P_{fn} + P_{fp} \frac{C_{fp} - C_{tn}}{C_{fn} - C_{tp}} \right\} \right\} \quad (2)$$

Note that we can ignore constant terms and so this is equivalent to finding the classifier that minimizes $P_{fn} + \eta P_{fp}$ where

$$\eta = \frac{C_{fp} - C_{tn}}{C_{fn} - C_{tp}}$$

is the cost associated with a false positive and the cost of a false negative is 1. Note that, since the cost of errors are likely to be greater than the cost of a correct prediction, then $\eta \geq 0$.

Let us consider the problem at hand. A false positive error occurs when one incorrectly predicts fraud (and potentially aborts a call resulting in a displeased customer) but eventually further investigation would have identified that this was incorrect. On the other hand a false negative error results in a fraudulent call going unnoticed and this could lead to significant losses. We can consider the cost of true predictions to be zero. If, for example, the cost of a false negative is 10 times the cost of a false positive (with correct predictions having zero cost) then $\eta = 0.1$.

To illustrate the approach, a cost-sensitive CART DT model was trained in a similar manner to the traditional CART DT




Fig. 3. Decision Tree for the CART Model

TABLE V
CONFUSION MATRIX FOR THE COST SENSITIVE CART DT

Predicted \ Actual	0	1
0	38681	1125
1	6640	3497

presented above but incorporating costs with $\eta = 0.1$. The resulting tree is shown in Figure 4. Note that branching is performed to provide more false positive errors than false negative ones because of the lower associated cost. Using this model the predicted fraud outcomes for the test data yielded the Confusion Matrix given in Table V. Here we clearly see the trade-off in the types of errors, for which predictive accuracy stands at 84.5%, roughly 5% lower than that of the traditional CART DT shown earlier.

Since the exact costs associated with errors are not known, we evaluated the costs for various values of η . We then compare the cost of the cost-based optimal result with the cost of the accuracy-based optimal result. This ratio is

$$R = \frac{P_{fn}(\text{cost}) + \eta P_{fp}(\text{cost})}{P_{fn}(\text{acc}) + \eta P_{fp}(\text{acc})} \quad (3)$$

Note that when $\eta = 0$ then false positive errors cost nothing and so it is optimal to always choose a positive outcome in which case $R = 0$. If $\eta = 1$ then this corresponds to unit costs for both types of errors and hence the resulting optimal probabilities are the same that would be obtained for accuracy and hence in this case we get $R = 1$. In Figure VII we plot the value of this ratio R for various values of η .

In Figure VII, it is clear that savings are expected to increase as error rates increase. This has the potential to provide the telecommunication company with tangible cost savings when faced with the possibility of prediction errors, and thus

provides a major advantage over the traditional approach to the problem at hand.

VIII. CONCLUSIONS AND FUTURE WORK

Training and testing the CART DT and SVM models produced acceptable performance results. Our analysis created a suitable platform for further development of a PRS fraud detection model by providing useful insight on the matter. At the core, including other call level variables as input, or revisiting/resampling the training and testing data, may also be of value. Moreover, it was illustrated that using a cost-sensitive approach can lead to tangible cost savings, despite sacrificing some predictive accuracy as compared to a traditional approach. This approach is highly recommended for future research, and will be examined further.

IX. CONFLICT OF INTEREST


On behalf of all authors, the corresponding author states that there is no conflict of interest.

REFERENCES

- [1] C. Agubor, G. Chukwudebe, and O. Nosiri, "Security challenges to telecommunication networks: An overview of threats and preventive strategies," in *2015 International Conference on Cyberspace (CYBER-Abuja)*. IEEE, 2015, pp. 124–129.
- [2] Y.-J. Zheng, X.-H. Zhou, W.-G. Sheng, Y. Xue, and S.-Y. Chen, "Generative adversarial network based telecom fraud detection at the receiving bank," *Neural Networks*, vol. 102, pp. 78 – 86, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0893608018300698>
- [3] C. Duffy and R. T. Coupe, "Detecting and combating internet telephony fraud," in *Crime Solvability Factors*. Springer, 2019, pp. 127–148.
- [4] Y. Alraouji and A. Bramantoro, "International call fraud detection systems and techniques," in *Proceedings of the 6th International Conference on Management of Emergent Digital EcoSystems*, ser. MEDES '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 159–166. [Online]. Available: <https://doi.org/10.1145/2668260.2668272>



ICACT2022 February 13 ~ 16, 2022



385

- ICACT2022 February 13 ~ 16, 2022



Mariella Rivas was born in Trinidad and Tobago, in May 1993. She received the BSc degree in actuarial science and MSc degree in statistics from The University of the West Indies, St Augustine, Trinidad and Tobago, in 2016 and 2019 respectively. In 2016, she joined the Finance Department of the Trinidad and Tobago Unit Trust Corporation (TTUTC), Port-of-Spain, Trinidad and Tobago as an Associate Professional. Ms. Rivas later became a part of the Strategic Analytics and Performance Department, Telecommunications Services of Trinidad and Tobago (TSTT), Port-of-Spain, Trinidad and Tobago as an Intern Data Scientist in 2019. Since January 2021, she has performed as a Data Scientist at Guardian Group, Port-of-Spain, Trinidad and Tobago. Presently, her main research interests include machine and deep learning, optimization, analytics translation, and real-time big data analysis.

applied data science, operations research and performance, and pricing optimization for cellular networks.



Richard Roach was born in Trinidad and Tobago, in August 1983. He received the BSc degree in economics with minor studies in mathematics and finance from The University of the West Indies, St Augustine, Trinidad and Tobago and the International Master of Business Administration (MBA) from Arthur Lok Jack Graduate School of Business, Mount Hope, Trinidad and Tobago in 2006 and 2011 respectively. He joined Repsol YPF as a Commercial Analyst in November 2007, and has been with the Strategic Analytics and Performance Department, Telecommunications Services of Trinidad and Tobago (TSTT), Port-of-Spain, Trinidad and Tobago as a Strategic Analyst since 2017. His major areas of research interest are financial modelling, relationship management, and dashboard design with analytics.



Dr. Hosein has received five degrees, including a PhD in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology (MIT). He has worked at Bose Corporation, Bell Laboratories, ATT Laboratories, Ericsson and Huawei. Currently, he is the administrative and technical contact for the TT top level domain, CEO of the TTNIC, and a Professor of Computer Science at The University of the West Indies. He has published extensively with over 150 refereed journal and conference publications. He holds 41 granted patents in the areas of telecommunications and wireless technologies. In 2004, he was nominated for the Ericsson Inventor of the Year award and was later awarded as the Huawei US Wireless Research Employee of the year and an Anthony Sabga Caribbean Laureate for Science and Technology in 2007 and 2015 respectively. His present research interests are