Tuesday , June 14, 2005

Subscription Fraud Prevention in Telecommunications using Fuzzy Rules and Neural Networks

Pablo A. Estévez *, Claudio M. Held and Claudio A. Perez

Department of Electrical Engineering, University of Chile, Casilla 412-3, Santiago, Chile

Abstract

A system to prevent subscription fraud in fixed telecommunications with high impact on long-distance carriers is proposed. The system consists of a classification module and a prediction module. The classification module classifies subscribers according to their previous historical behavior into four different categories: subscription fraudulent, otherwise fraudulent, insolvent and normal. The prediction module allows us to identify potential fraudulent customers at the time of subscription. The classification module was implemented using fuzzy rules. It was applied to a database containing information of over 10,000 real subscribers of a major telecom company in Chile. In this database a subscription fraud prevalence of 2.2% was found. The prediction module was implemented as a multilayer perceptron neural network. It was able to identify 56.2% of the true fraudsters, screening only 3.5% of all the subscribers in the test set. This study shows the feasibility of significantly preventing subscription fraud in telecommunications by analyzing the application information and the customer antecedents at the time of application.

Keywords: fraud prevention, fraud detection, subscription fraud, neural networks, fuzzy rules.

* Corresponding author. Tel. +56-2-6784207; fax: +56-2-6720162. E-mail address: pestevez@cec.uchile.cl

Review Copy

Tuesday , June 14, 2005

1. Introduction

The biggest revenue leakage area in the telecom industry is fraud (Wieland, 2004). Global telecommunications fraud losses are estimated in the tens of billions of dollars every year (FML, 2003; Hoath, 1998). The history of telecommunications crime, including several types of fraudulent activities, was reviewed by Collins (1999a, 1999b, 2000). Some authors have emphasized the importance of distinguishing between fraud prevention and fraud detection (Bolton & Hand, 2002). Fraud prevention describes measures to avoid fraud to occur in the first place. In contrast, fraud detection involves identifying fraud as quickly as possible once it has been committed.

Shawe-Taylor et al. (2000) distinguished six different fraud scenarios: subscription fraud, the manipulation of Private Branch Exchange (PBX) facilities or dial through fraud, freephone fraud, premium rate service fraud, handset theft and roaming fraud. Subscription fraud, which is defined as the use of telephone services with no intention of paying, is probably the most significant and prevalent worldwide telecom fraud (FML, 2003; Hoath, 1998). Subscription fraud can be subdivided into two categories: a) for profit, i.e. mainly for selling long distance calls and b) for personal usage. Subscription fraud can be committed upon fixed and mobile telephones, and it is usually difficult to distinguish from bad debt, particularly if the fraud is for personal usage. Both subscription fraud and bad debt are major problems to telecoms in developing and third world countries (Hoath, 1999). Two strategies have been proposed for detecting subscription fraud: examining account applications and tracking customer behavior (Fawcett & Provost, 2002b). Other efforts have focused on formalizing and predicting the deceiving intention of fraudsters (Barghava et al., 2003).

Review Copy

Tuesday , June 14, 2005

The detection of fraud in mobile telecommunications was investigated in the European project ASPeCT (Advance Security for Personal Communications Technologies) (Shawe-Taylor, Howker & Burge, 1999; Shawe-Taylor et al., 2000; Burge & Shawe-Taylor, 2001). The ASPeCT fraud detection tool is based on investigating sequences of call detail records (CDRs), which contain the details of each mobile phone call attempt for billing purposes. The information produced for billing also contains usage behavior information valuable for fraud detection. A differential analysis is performed to identify a fraudster through profiling the behavior of a user. The analysis of user profiles are based on comparison of recent and longer-term behavior histories derived from the toll ticket data. Alarms are activated when the usage pattern of a mobile phone changes significantly over a short period of time. The ASPeCT fraud detection tool utilizes a rule-based system for identifying certain frauds, and neural networks (NNs) to deal with novel or abnormal instances or scenarios. Rosset et al. (1999) used customer data, in addition to CDRs, to discover rules for identifying subscription fraud.

According to Cahill et al. (2002), a fraud detection algorithm has two components: a) a summary of the activity on an account that can be kept current and b) rules that are applied to account summaries to identify accounts with fraudulent activity. A popular approach is to reduce the CDRs for an account to several statistics that are computed for each period, e.g. average call duration, and compare them to thresholds. Fawcett and Provost (1997a, 1997b, 2002a) developed a method for choosing account-specific thresholds rather than universal thresholds. Their procedure takes daily traffic summaries for a set of accounts that experienced at least 30 days of fraud-free traffic activity followed by a period of fraud. This method was applied to cellular cloning, in which fraudulent usage is superimposed upon the

Review Copy

Elsevier

legitimate usage of an account. For each account a set of rules that distinguish fraud from non-fraud was developed. The superset of the rules for all accounts was then pruned by keeping only those that cover many accounts, with possibly different thresholds for different accounts. Cahill et al. (2002) defined account signatures to track legitimate call behaviors in real time. An account signature describes which call variables (e.g. call duration) are likely and which are unlikely for the account. Signatures evolve with each new call that is not considered fraudulent, so each established customer eventually has its own signature. Likewise fraud signatures are defined for each kind of fraud using the same structure as an account signature. A call is scored by comparing its probability to belong to the account signature and to a fraud signature. For new accounts the first calls are used to assign signature components, associating them with calling patterns of a given segment of customers with similar initial information.

Cortes et al. (2001, 2003) applied large dynamic graphs, represented as the union of small sub-graphs called communities of interest, to the area of telecommunications fraud detection. The nodes in the graphs are network IDs and the edges represent communications between pairs of network IDs. In one application the "guilt by association" argument was used to detect new cases of fraud in the network, one week after the new accounts were activated. It was found that the probability of an account to be fraudulent is an increasing function of the number of fraudulent nodes in its community of interest. A second example used a distance metric between communities of interest to suggest when an individual whose account had recently been disconnected for fraud had assumed a new network identity. This assumed that the calling patterns of the new account had not changed very much from the previous account.

Review Copy

In the last decade modern intelligent systems have been applied to fraud detection. Bolton and Hand (2002) reviewed the statistical and machine learning technologies for fraud detection, including their application to detect activities in money laundering, e-commerce, credit card fraud, telecommunication fraud and computer intrusion. Weatherford (2002) presented several real-world applications of intelligent fraud detection technologies. Kou et al. (2004) made a survey of fraud detection techniques used in telecommunication, as well as in credit card fraud and computer intrusion. Phua et al. (2005) made a comprehensive survey of data mining techniques applied to fraud detection. Hong and Weiss (2001) presented several predictive models for data mining applied to fraud detection and insurance risk assessment. Some authors have provided comprehensive surveys of NNs (Wong, Bodnovich, & Selvi, 1997; Vellido, Lisboa, & Vaughan, 1999) and Expert Systems (ES) (Liao, 2005; Wong & Monaco, 1995) applications in business. Vellido et al. (1999) found that published applications of NNs in real-world scale are scant. One difficulty for publishing results is the need for confidentiality of private companies operating in a tough competitive environment. The main advantages of NNs are: a) their suitability to handle incomplete, missing or noisy data; b) being a non-parametric method, it does not require any a-priori assumptions about the distribution and/or mapping of the data; and c) their demonstrated capability to approximate any continuous function. The lack of explanatory capabilities is considered as the main shortcoming of the application of NNs. Hence, several attempts have been made to integrate NNs and ES; a synergistic effect between them is expected, as ES are characterized by their capability of explaining their own reasoning process.

Review Copy

Tuesday , June 14, 2005

Other authors have used data mining techniques to develop a decision support system for predicting customer insolvency in telecommunications (Daskalaki et al., 2003). In their approach, it is assumed that insolvent customers behave differently on the average from the rest of the customers, especially during a critical period preceding the due-date for payment. The prediction of customer insolvency for a telecommunications company as a problem was found to be similar to the fraud detection problems in mobile and conventional telecommunications as well as in credit or calling card operations. Among the common characteristics found are the following: significant loss of revenue, unpredictability of human behavior, information retrieval involves processing huge amounts of data from several different sources; fraudulent cases are rare compared to legitimate ones. Ezawa and Norton (1996) constructed Bayesian networks to predict uncollectible telecommunications accounts.

The related problem of subscriber churning in mobile telecommunications, i.e. the movement of subscribers from one provider to another, has been investigated using NNs (Mozer et al., 2000) and data mining (Wei & Chiu, 2002). Mozer et al. (2000) used techniques from statistical machine learning to evaluate the benefits of predicting churn. Experiments were carried out using a database of 47,000 subscribers that included information about their usage (CDRs, quality of service), billing, credit, application for service (contract details, rate plan, and credit report), and complaint history. The outcome was expressed using a lift curve which plots the fraction of all churners having churn probability above a threshold versus the fraction of all subscribers having churn probability above the threshold. Wei and Chiu (2002) built a model that predicts churning from subscriber contractual information and call patterns changes extracted from CDRs. The

Review Copy

Elsevier

Tuesday , June 14, 2005

proposed churn-prediction technique used a decision tree induction algorithm for learning. A randomly selected data set included 1.5-2% churners and 98-98.5% non-churners. The proposed technique was capable of identifying 10% of the subscribers that contained 54% of the true churners.

According to Chan et al. (1999) the fraud detection task is characterized by a) skewed distribution of data, i.e. many more transactions are legitimate than fraudulent, and b) non-uniform cost per error, e.g. the cost of failing to detect a fraud varies with each transaction. The authors addressed skewness by partitioning the data set into subsets with a desired distribution, applying mining techniques to the subsets, and combining the mined classifiers. The issue of non-uniform cost was addressed by developing the appropriate cost model and biasing the methods towards reducing cost. Stolfo et al. (1997) argued that for the fraud detection domain, the fraud catching rate (true positive rate) and false alarm rate (false positive rate) are better metrics than overall accuracy due to unequal error costs and uneven class distributions. Given a skewed distribution on the original data, artificially balanced training data with a 50%/50% fraud/non-fraud distribution lead to classifiers with the highest true positive rate and low false positive rate.

Provost and Fawcett (1997) presented a method called *ROC convex hull* which combined techniques from Receiver Operating Characteristic (ROC) analysis and a decision analysis method for analyzing and visualizing classifier performance. ROC graphs depict tradeoffs between the hit (true positive) rate and the false alarm (false positive) rate. However, ROC graphs illustrate the behavior of a classifier disregarding class distribution or error cost (Fawcett, 2003). An iso-performance line is defined in ROC space, where all classifiers

Review Copy

corresponding to points in the line have the same expected cost. Each set of classes and distributions define a family of iso-performance lines. The optimal classifier will be the point on the convex hull that intersects the iso-performance line with largest true positive rate intercept. The method assumes that there are only two classes and that costs do not vary within a given type of error.

Chile has one of the most liberal telecommunications regimes in the world (Stehmann, 1995). Chile has pioneered the privatization and deregulation of both the long distance and the local telephony market (Paredes, 2005). By June 2002, there were 20 carriers operating in the long-distance market; 13 operators in the local telephony market and 6 operators in the mobile communications market (SUBTEL, 2002). Although the ratio of mobile phones to fixed lines reached 1.6 in 2002, the traffic generated in the latter network was 6 times higher than the traffic generated in the mobile network.

The aim of this research is to develop a system for scoring the risk of subscription fraud at the time of application for fixed telephone lines. In particular, our study focused on the identification of subscribers who would order new fixed lines to make use of long distance services without paying the corresponding telephone bills. This corresponds mainly to the fraud-for-profit category. For this problem the objective was to detect as many subscription fraudsters as possible while minimizing false alarms. Our study was conducted with real data provided by a major telecom operating in Chile. For confidentiality reasons the telecom name and some particular information, such as the full list of variables and rules, are not published.

Review Copy

Tuesday , June 14, 2005

2. Methods

Creating a Data Set

The first step consisted in collecting two lists of phone numbers. The first list consisted of 1,523 fixed phone lines blocked by the telecom from July 2001 to July 2002, as a result of operating a traffic-based commercial fraud detection system for monitoring long distance calls online. The data corresponded mostly to residential subscribers. The second list corresponded to a sample of 12,805 new fixed lines installed by the telecom from October through December 2001. This sample excluded business subscribers with more than 30 lines and phone lines without long distance calls during that period.

The second step consisted in searching for the data sources which were distributed over eight different databases maintained by the telecom. These databases are mainly related to customer, account, application and billing information. The CDRs were not directly used in this research, since our goal was to evaluate the risk of fraud at the time of the application, before the installation of a new line. All the information contained in the eight different sources from July 2001 through December 2002, related to the two lists of phone numbers mentioned above and their subscribers, was extracted and integrated into a single repository built for this purpose. The repository contained 234 fields associated with each subscriber. The total collected raw data was over 10 GB. This historical repository was implemented as a relational database in PostgresSQL[1]. The list of new phone lines was monitored regularly in intervals of two weeks for at least a year after the installation date. The list of blocked phone lines was followed up for at least 6 months after the blocking date.

Elsevier

Tuesday , June 14, 2005

Categories of Subscribers

Fraud cases would generally be detected online triggered by traffic measures by the commercial fraud detection system, and confirmed later on as such during the billing process. In order to generate a database of known fraudulent/legitimate cases, it was necessary to formalize the definition of subscribers' categories. Consequently the following four categories of subscribers were defined:

- **Subscription fraudulent**. Most of the users in this category do not pay their bills at all, but if they do, the debt/payment ratio is very high. The line is typically blocked due to suspicious behavior in long distance calls within 6 months after the installation date.
- **Otherwise fraudulent**. Subscribers for more than a year who present a sudden change in their calling behavior, generating an abnormal rise in their newer billing accounts.
- **Insolvent**. Subscribers with a total debt of less than 10 times their monthly payments, having two or more unpaid bills. This category includes new customers that have never paid their bills but whose monthly expenditures are similar to average residential lines.
- **Normal**. Customers with their bills up to date or at most a single unpaid bill for less than 30 days after the due date.

System Architecture

The proposed system for preventing subscription fraud consists of two modules: a classification module and a prediction module. The classification module separates subscribers according to their historical behavior into one out of four of the defined categories: subscription fraudulent, otherwise fraudulent, insolvent or normal. This module uses as inputs the information available about bills, payments, phone line blockings, etc.

[1] http://www.postgresql.org/

Review Copy

Elsevier

Tuesday , June 14, 2005

The main purpose of the classification module is to generate a database of known fraudulent/legitimate cases. The prediction module allows the identification of potential fraudulent subscribers at the time of application. This module uses as inputs the information available about new subscribers, such as demographics and commercial antecedents, as well as application and account information.

Classification Module

First 672 cases were drawn from the repository and classified manually into the four categories described above. The manual classification procedure was assisted by an expert with many years of experience in telecommunication fraud management. This was a time-consuming procedure since for each case, all the information available in the repository had to be examined on the computer screen. Table 1 shows the number of cases manually classified for each of the four categories of subscribers.

The classification module was designed with a hierarchical tree structure, including three layers and five nodes, as shown in Figure 1. The first layer consists of the root node, which discriminates between fraudulent and normal subscribers, but assigns the insolvent subscribers to any of the two groups. The second layer has two nodes. Node N/I discriminates between normal and insolvent cases. Node F/I discriminates between fraudulent and insolvent cases. The third layer has two nodes that discriminate among subscription fraudulent, otherwise fraudulent and insolvent cases. Node I/O distinguishes between insolvent and otherwise fraudulent. Node S/O discriminates between subscription fraudulent and otherwise fraudulent.

Review Copy

Elsevier

Tuesday , June 14, 2005

The data set of 672 cases was used to select the variables of the classification module, as well as to design fuzzy rules to discriminate among the categories. Some of the variables used here were: customer longevity (**clo**), time elapsed between installation date and blocking date (**tbi**), debt/payment ratio (**dpr**), phone blocked flag (**pbf**), account balance (**aba**), number of days with unpaid bills **(ndu),** maximum debt with (long distance) carriers **(mdc)**, and call forwarding traffic **(cft)**. As an example, Figure 2 shows the values taken by the variable **dpr** for the 672 manually classified cases. In the x-axis the cases are ordered by category: S (subscription fraudulent), O (otherwise fraudulent), I (insolvent) and N (normal). In the y-axis the variable was saturated at a maximum value of 100. As can be seen, the variable **dpr** allows to discriminate between the S and N categories, but the I category showed cases with low and high **dpr** values.

For continuous variables, three Gaussian-like fuzzy membership functions were defined to measure low-risk (LR), medium-risk (MR) and high-risk (HR) of subscription fraud. For binary variables, the value 1 was associated to HR and 0 to LR. A total of 54 fuzzy rules were defined for the classification module, using 17 variables. Here we present some examples. At the root node of the tree-classifier shown in Fig. 1, the first three rules generated were:

Rule 1. IF (clo is LR) AND (tbi is LR) AND (dpr is LR) AND (pbf is LR) AND (aba is LR) THEN (Output_RootNode is Node N/I).

Rule 2. IF (clo is HR) AND (tbi is HR) AND (dpr is HR) AND (pbf is HR) AND (aba is HR) THEN (Output_RootNode is Node F/I).

Rule 3. IF (aba is LR) AND (ndu is LR) THEN (Output_RootNode is Node N/I).

Review Copy

Tuesday , June 14, 2005

The first three rules generated at the F/I node were:

Rule 4. IF (mdc is HR) AND (tbi is HR) AND (dpr is HR) AND (pbf is HR) AND (aba is HR) THEN (Output_F/I_Node is Node S/O).

Rule 5. IF (mdc is BR) AND (dpr is BR) AND (aba is BR) THEN (Output_F/I_Node is Node I/O).

Rule 6. IF (cft is HR) THEN (Output_F/I_Node is Node S/O).

The proposed classification module was applied to the cases available in the repository, in order to automatically label them into the different categories of subscribers. This procedure generated a database of known fraud/legitimate cases.

Prediction module

The aim of the prediction module is to discriminate whether an application for a new phone line corresponds to a subscription fraud case or not. In contrast with the classification task which is done *a posteriori*, there is scant information for predicting subscription fraud before the installation of a new phone line. For new subscribers there is no billing or account information available. One source of information comes from the application process, where some customer demographics are obtained: name, address, national ID number, contact phone number, date of birth (not always available), etc. These variables were crossed with the information available in the account database, to get for example information about other lines installed at the same address. Another source of information was a commercial records database, that is publicly available in Chile through a paid service called DICOM. The DICOM database registers insolvency situations mostly related to banks and department stores. Some of the DICOM variables considered were: total

Review Copy

Elsevier

unpaid debt amount, number of months of insolvency, and time elapsed since the due date of the last unpaid document.

The variables above mentioned were investigated in their predictive power using the labeled data in the repository. For example, the contact phone number which is required by the telecom in the application process, allowed the detection of sequences of cases of subscription fraud. Fig. 3 shows a sequence of seven fraud cases committed within a period of three months which are related to each other through the contact phone number used for ordering new lines. Typically a fraudster ordered a new phone line and committed fraud a few days after the installation. In the meantime, previous to the blocking, the fraudster ordered another line using the first line installed as the contact phone number.

The prediction module consists of a multilayer feedforward neural network with 22 inputs, 10 hidden units and 2 output units. The output units indicate whether the case is fraudulent or legitimate. The labeled data set was split into a training set, a validation set and a testing set. The few samples corresponding to the category "otherwise fraudulent" were eliminated from the database since their prevalence was too small (see Table 1). Moreover, subscription fraud is the only fraud type duly identifiable at the time of the application. In addition, the “normal” and “insolvent” categories were merged into one category of non-fraudulent cases. The training set was used to adjust the parameters of the NN. It consisted of 1680 examples: 1440 legitimate cases and 240 fraudulent cases. To avoid biasing the NN towards the category with more samples, the fraudulent cases were repeated three times. The validation set was used for selecting the best predictor and avoid overlearning on the

Review Copy

Tuesday , June 14, 2005

training set. It consisted of 240 cases. The testing set contained 4216 cases, which were utilized to evaluate the performance of the prediction module with independent data.

Clearly a naive classifier can get a high accuracy on a skewed data set, where the prevalence of fraud is very low, but at the cost of misclassifying all the true fraudulent cases. For this reason, the ROC curve was chosen as the performance criterion instead of the classifier accuracy. The true positive (TP) rate is defined as the per unit rate of fraudulent cases that are correctly identified over the total of fraudulent cases. The false positive (FP) rate is defined as the per unit rate of legitimate (normal and insolvent) cases that are incorrectly identified as fraudulent over the total of legitimate cases. The ROC curve plots the TP rate versus the FP rate. To take into account the class distribution and the error costs, an iso-performance line is defined in ROC space, where all classifiers corresponding to points in the line have the same expected cost (Provost & Fawcett, 1997). Two points (TP_1, FP_1) and (TP_2, FP_2) have the same performance if

$$\frac{TP_2 - TP_1}{FP_2 - FP_1} = \frac{p(N)c(FP)}{p(P)c(FN)}, \qquad (1)$$

where p(N) and p(P) are the prior probabilities of obtaining a negative and a positive example, respectively. The c(FP) and c(FN) represent the costs of a FP and a FN error, respectively. This equation defines the slope of the iso-performance line. Given an ROC curve and an iso-performance line, the intersection defines the operation point.

3. Results

The data set of 672 manually labeled cases was split into a training set of 381 cases and a test set of 291 cases. The training set was used to adjust the functions and parameters of the

Elsevier

tree-classifier. The test set was used to evaluate the performance of the proposed classification module. In both sets a 100% accuracy was achieved. The resulting classification module was applied to automatically label 10,610 cases available in the repository. Table 2 shows the statistics per subscriber category obtained for the list of new phone lines using the automatic classification module. The list of blocked lines also available in the repository was not considered in Table 2, in order to estimate the statistics of a random sample of new lines. The prevalence of subscription fraud obtained was 2.2%, i.e. a ratio of legitimate-to-fraud cases of approximately 50:1.

For the predictive study, 6,106 cases having complete information about the customers' commercial antecedents and the application data at the time of subscription were considered. This data set included cases from the two lists of phones. It was split into three sets for training, validation and testing of the prediction module. The multilayer perceptron NN was trained on the training set. To avoid overlearning the validation set was used to select the best predictor. Table 3 shows the composition of the test set, that had a fraud prevalence of 2.9%, which is similar to the ratio found with the classification module for the set of new subscribers.

Figure 4 shows the ROC curve obtained with the prediction module. The top left point (0,1) represents the perfect predictor where all fraudulent cases are caught without generating false alarms. Therefore, the closer the ROC curve is to the point (0,1) the better is the performance. To plot the iso-performance line, the prevalence of fraud was set at 2%, thus p(P) = 0.02 and p(N) = 0.98. We chose a scenario where the cost of a FN error double the cost of a FP error, yielding a rate of $c(FN)/c(FP) = 2,$ i.e., missing a case of fraud is worse

Review Copy

Tuesday , June 14, 2005

than a false alarm. Since the real costs are difficult to specify, a range of values could be considered, or alternatively a limit could be set to the fraction of all susbscribers to be examined. At the operation point the TP and FP rates were 56.2% and 2.4%, respectively. In other words, the prediction module was able to identify 56.2% of the true fraudsters, screening only 3.5% of all subscribers.

4. Discussion and Conclusions

In contrast to fraud detection systems that operate once the fraud has been committed, the proposed system is predictive and operates at the application time. Demographics and commercial antecedents, as well as other characteristics associated to the application for a new phone line, were used as predictors. The predictive module was able to identify 3.5% of the subscribers containing 56.2% of the true fraudsters. A manual analysis of errors showed that most of the FP cases corresponded to the insolvent category. One third of these corresponded to customers that never paid the bills but had a typical residential average expenditure. This pattern corresponds to the category of fraud for personal usage, and could be considered as a kind of subscription fraud.

In the future, the proposed system could be enhanced by adding information about the subscriber's call patterns. In this way, cases marked as potentially risky by the system at the time of application, could be followed up after the installation date for closer examination.

It is well-known that the patterns as well as the levels and costs of fraud change very quickly in time. Because of this complexity, any fraud system could become rapidly obsolete. In our system architecture, the classification module should operate continuously

Review Copy

Elsevier

Tuesday , June 14, 2005

to monitor the prevalence of fraud, and to provide new fraud cases for adjusting the prediction module.

This study was carried out on fixed telecommunications but the techniques proposed here could be extended to subscription fraud in mobile communications, as well as other markets.

Acknowledgments

This work was supported in part by grants Fondef D99I1050 and Fondecyt 1050751 from Conicyt-Chile. We thank Francisco Rivera for his collaboration as an expert in several stages of our research project. We also thank research assistants Leonardo Causa, Cristián Díaz, Alejandro Ehrenfeld, Cristián Figueroa, Rodrigo Flores, Jaime Heiss, José Riquelme, Paul Vallejos, and Alvaro Santibañez for their help in building the repository and making the software implementations. Alfredo Schnell, an experienced IT entrepreneur, joined the research team to start a business to develop and expand the proposed fraud prevention system.

Review Copy

Elsevier

Tuesday , June 14, 2005

References

Barghava, B., Zhong, Y., & Lu, Y. (2003). Fraud formalization and detection. In Y. Kambayashi, M. Mohania, W. Wöb (Eds.), *Data Warehousing and Knowledge Discovery, 5th Int. Conference, DaWak 2003, Prague, Czech Republic* (pp. 330-339). LCNS Vol. 2737, Springer-Verlag.

Bolton, R.J., & Hand, D.J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235-249.

Burge, P., & Shawe-Taylor, J. (2001). An unsupervised neural network approach to profiling the behavior of mobile phone users for use in fraud detection. *Journal of Parallel and Distributed Computing*, 61, 915-925.

Cahill, M. H., Lambert, D., Pinheiro, J.C., & Sun, D. X. (2002). Detecting fraud in the real world. In J. Abello, P. M. Pardalos, M. G. C. Resende (Eds.), *Handbook of Massive Data Sets*, (pp. 913-930). Dordrecht: Kluwer Academic Publishers.

Chan, P. K., Fan, W., Prodromidis, A. L., & Stolfo, S. J. (1999). Distributed data mining in credit card fraud detection. *IEEE Intelligent Systems*, 14(6), 67-74 .

Collins, M. (1999a). Telecommunications Crime – Part 1. *Computer & Security*, 18, 577-586.

Collins, M. (1999b). Telecommunications Crime – Part 2. *Computer & Security*, 18, 683-692.

Collins, M. (2000). Telecommunications Crime – Part 3. *Computer & Security*, 19, 141-148.

Cortes, C., Pregibon, D., & Volinsky, C. (2001). Communities of interest. In F. Hoffman, D. J. Hand, N. Adams, G. Guimaraes (Eds.), *Advanced in Intelligent Data Analysis, 4th*

Review Copy

Elsevier

Tuesday , June 14, 2005

Int. Conference IDA 2001, Cascais, Portugal (pp. 105-114). LCNS Vol. 2189, Springer-Verlag.

Cortes, C., Pregibon, D., & Volinsky, C. (2003). Computational methods for dynamics graphs. *Journal of Computational and Graphical Statistics*, 12(4), 950-970.

Daskalaki, S., Kopanas, I., Goudara, M., & Avouris, N. (2003). Data mining for decision support on customer insolvency in telecommunications business. *European Journal of Operational Research*, 145, 239-255.

Ezawa, K.J., & Norton, S.W. (1996). Constructing Bayesian networks to predict uncollectible telecommunications accounts. *IEEE Expert*, 11(5), 45-51.

Fawcett, T., & Provost, F. (1997a). Combining data mining and machine learning for effective fraud detection. In AI Approaches to Fraud Detection & Risk Management, Workshop Technical Report WS-97-07 (pp. 14-19). AAAI Press, Menlo Park, California.

Fawcett, T., & Provost, F. (1997b). Adaptive fraud detection. *Data Min. Knowl. Disc*, 1(3), 291-316.

Fawcett, T. (2002a). Case study: Adaptive fraud detection. In W. Klösgen, & J. Zytkow (eds.), *Handbook of Data Mining and Knowledge Discovery* (section H1.2.1). Oxford U. Press, UK.

Fawcett, T., & Provost, F. (2002b). Fraud detection. In W. Klösgen, & J. Zytkow (eds.), *Handbook of Data Mining and Knowledge Discovery* (section F2). Oxford U. Press, UK.

Fawcett, T. (2003). ROC graphs: Notes and practical considerations for data mining researchers. HP Labs. Technical Report HPL-2003-4.

FML (2003). FML Revenue Assurance and Fraud Management Yearbook.

Review Copy

Elsevier

Tuesday , June 14, 2005

Hoath, P. (1999). What's new in telecoms fraud?. *Computer Fraud & Security*, 1999(2),13-19.

Hoath, P. (1998). Telecoms fraud, the gory details. *Computer Fraud & Security*, 1998(1),10-14.

Hong, S. J., & Weiss, S. M. (2001). Advances in predictive models for data mining. *Pattern Recognition Letters*, 22, 55-61.

Kou, Y., Lu, C.T., Sirwongwattana S., & Huang Y.P. (2004). Survey of fraud detection techniques. *Proceedings of the IEEE International Conference on Networking, Sensing and Control* (pp. 749-754). Taipei, Taiwan.

Liao, S. H. (2005). Expert systems methodologies and applications - a decade review from 1995 to 2004. *Expert Systems with Applications*, 28, 93-103.

Mozer, M. C., Wolniewicz, R., Grimes, D. B., Johnson, E., & Kaushansky, H. (2000) Predicting subscriber dissatisfaction and improving retention in the wireless telecommunications industry. *IEEE Trans. Neural Networks*, 11(3), 690-696.

Paredes, R.D. (2005). Lessons from the deregulation transition in Chile's local telephony market. *Telecommunications Policy*, 29, 333-350.

Phua, C., Lee, V., Smith, K., & Gayler R. (2005). A comprehensive survey of data mining - based fraud detection research. Available at http://www.bsys.monash.edu.au/ people/cphua/

Provost, F., & Fawcett, T. (1997). Analysis and visualization of classifier performance with nonuniform class and cost distributions. In AI Approaches to Fraud Detection & Risk Management, Workshop Technical Report WS-97-07 (pp. 57-63). AAAI Press, Menlo Park, California.

Review Copy

Elsevier

Tuesday , June 14, 2005

Rosset, S., Murad, U., Neumann, E., Idan, Y., & Pinkas, G. (1999). Discovery of fraud rules for telecommunications - challenges and solutions. In *Proceedings of the fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 409-413.

Shawe-Taylor, J., Howker, K., & Burge, P. (1999). Detection of fraud in mobile telecommunications. *Information Security Technical Report*, 4(1), 16-28.

Shawe-Taylor, J., Howker, K., Gosset, P., Hyland, M., Verrelst, H., Moreau, Y., Stoermann, C., & Burge, P. (2000). Novel techniques for profiling and fraud in mobile telecommunications. In P. J. G. Lisboa, B. Edisbury, & A. Vellido (eds.), *Business Applications of Neural Networks. The State-of-the-Art of Real World Applications* (pp. 113-139). World Scientific, Singapore.

Stehmann, O. (1995). Network liberalization and developing countries: The case of Chile. *Telecommunications Policy*, 19(9), 667-684.

Stolfo, S. J., Fan, D. W., Lee W. L., Prodromidis, A. L. & Chan, P. K. (1997) Credit card fraud detection using meta-learning: issues and initial results. In AI Approaches to Fraud Detection & Risk Management, Workshop Technical Report WS-97-07 (pp. 83-90). AAAI Press, Menlo Park, California.

SUBTEL (2002). Estadísticas del Sector de las Telecomunicaciones en Chile: 2000-Primer semestre 2002. Informe Estadístico 6. Subsecretaría de Telecomunicaciones, Gobierno de Chile (available at http://www.subtel.cl).Vellido, A., Lisboa, P. J. G, & Vaughan, J. (1999). Neural networks in business: a survey of applications (1992-1998). *Expert Systems with Applications*, 17, 51-70.

Weatherford, M. (2002). Mining for fraud. *IEEE Intelligent Systems*, 17, 4-6.

Wei, C. P., & Chiu, I. T. (2002) Turning telecommunications call details to churn

Review Copy

Elsevier

Tuesday , June 14, 2005

prediction: a data mining approach. *Expert Systems with Applications*, 23, 103-112 .

Wieland, K., (2004) The last taboo? Revenue leakage continues to hamper the telecom industry. *Telecommunications (International Edition)*, 38, 10-11.

Wong, B. K., Bodnovich, T. A., & Selvi, Y. (1997). Neural network applications in business: a review and analysis of the literature (1988-95). *Decision Support Systems*, 19, 301-320.

Wong, B. K., & Monaco, J. A. (1995). Expert system applications in business: a review and analysis of the literature (1977-1993). *Information & Management*, 29, 141-152.

TABLE CAPTIONS

Table 1. Number of cases manually classified by an expert for each of the four categories of subscribers.

Table 2. Results of the automatic labeling of the list of new phone lines using the classification module.

Table 3. Test set for the prediction module for subscription fraud.

Review Copy

Elsevier

Tuesday , June 14, 2005

FIGURE CAPTIONS

Figure 1. Classification Module: It consists of a tree structure including 3 layers and 5 nodes. The classes correspond to N (normal), I (insolvent), O (otherwise fraudulent) and S (subscription fraud).

Figure 2. Debt-payment ratio (dpr) for the 672 cases manually classified. The cases are ordered according to their class: S (subscription fraud), O (otherwise fraudulent), I (insolvent), and N (normal).

Figure 3. Sequence of seven fraud cases linked by the contact phone number (CP) given at the time of application. All lines were blocked (BD: blocking date) within a few days after their installation date (ID). The telephone number (TN) at the top left box was used as contact phone number for ordering two new lines before being blocked. Then, the TN at the top center box was used as CP to order two new lines, which in turn were all blocked within a few days.

Figure 4. ROC curve obtained with the prediction module. The operation point is the intersection point between the ROC curve and the iso-performance line, which takes into account the class distribution and the cost of errors. In this example, the operation point is a TP rate of 56.2% with a FP rate of 2.4%.

Review Copy

Elsevier

Tuesday , June 14, 2005

Table 1

Subscriber category	**Number of cases**
Normal	251
Insolvent	206
Otherwise fraudulent	10
Subscription fraudulent	205
TOTAL	672

Review Copy

Elsevier

Table 2

Subscriber Category	N° Cases	[%]
Normal or Insolvent	**10363**	**97.6**
Otherwise fraudulent	**18**	**0.2**
Subscription fraudulent	**229**	**2.2**
TOTAL	**10610**	**100.0**

Review Copy

Tuesday , June 14, 2005

Table 3

Subscriber Category	N° Cases	[%]
Non-fraudulent	**4095**	**97.1**
Subscription Fraudulent	**121**	**2.9**
TOTAL	**4216**	**100.0**

Review Copy

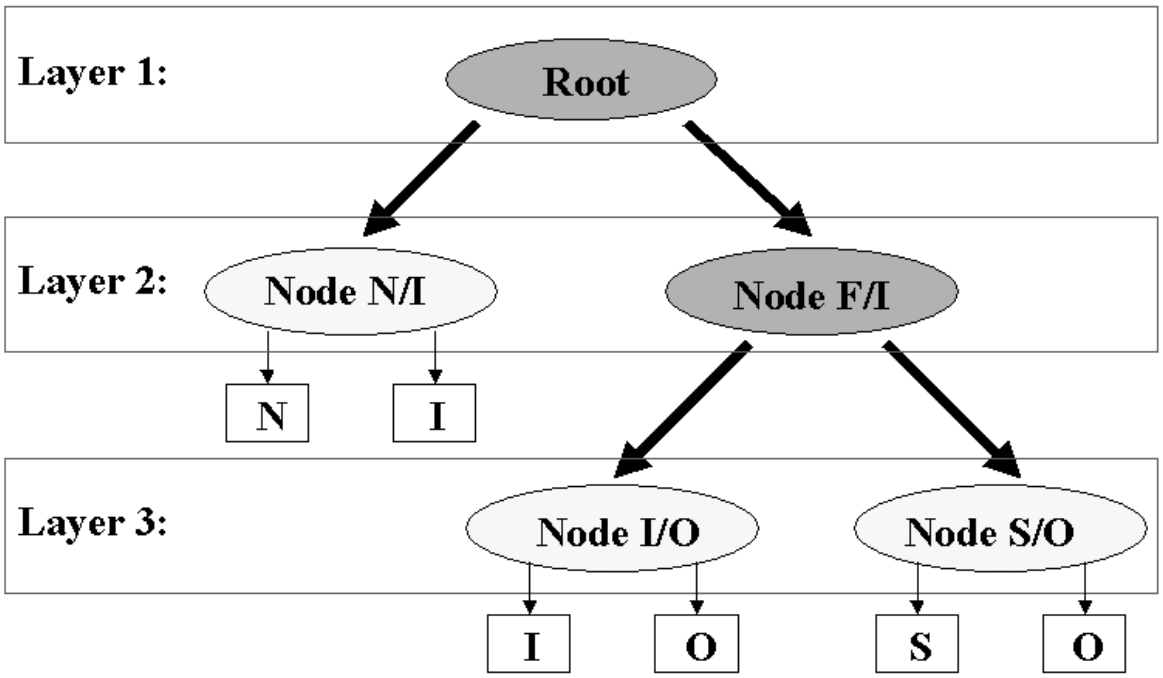

Figure 1

Review C

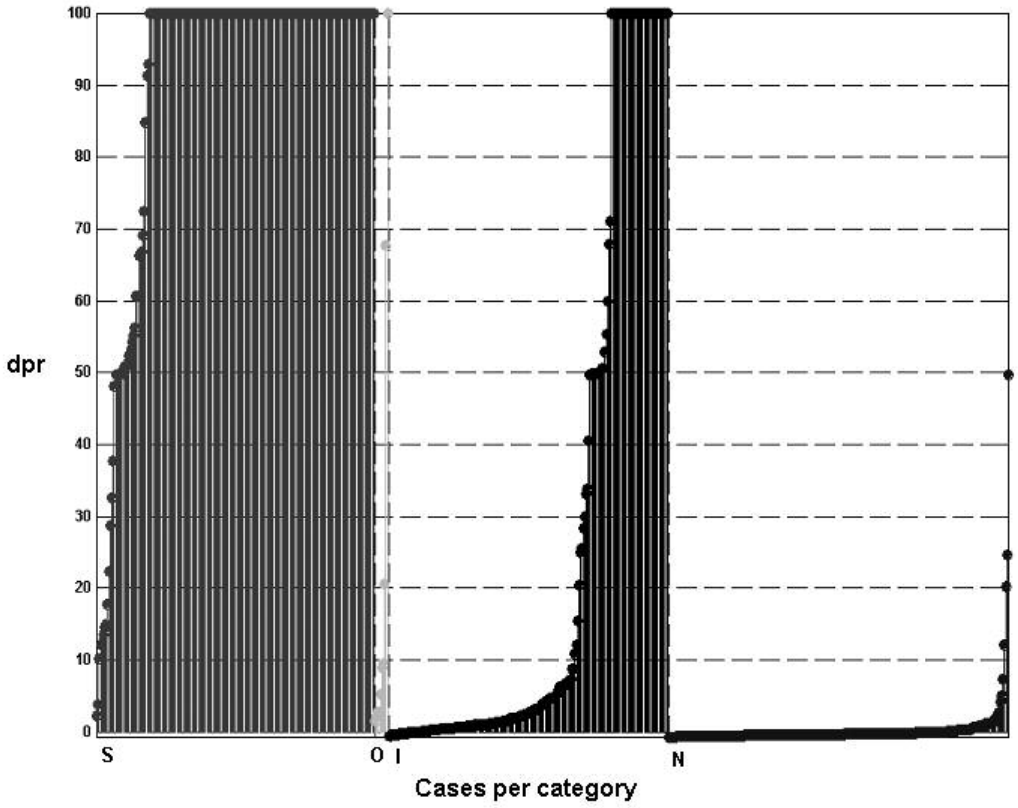

Figure 2

Review

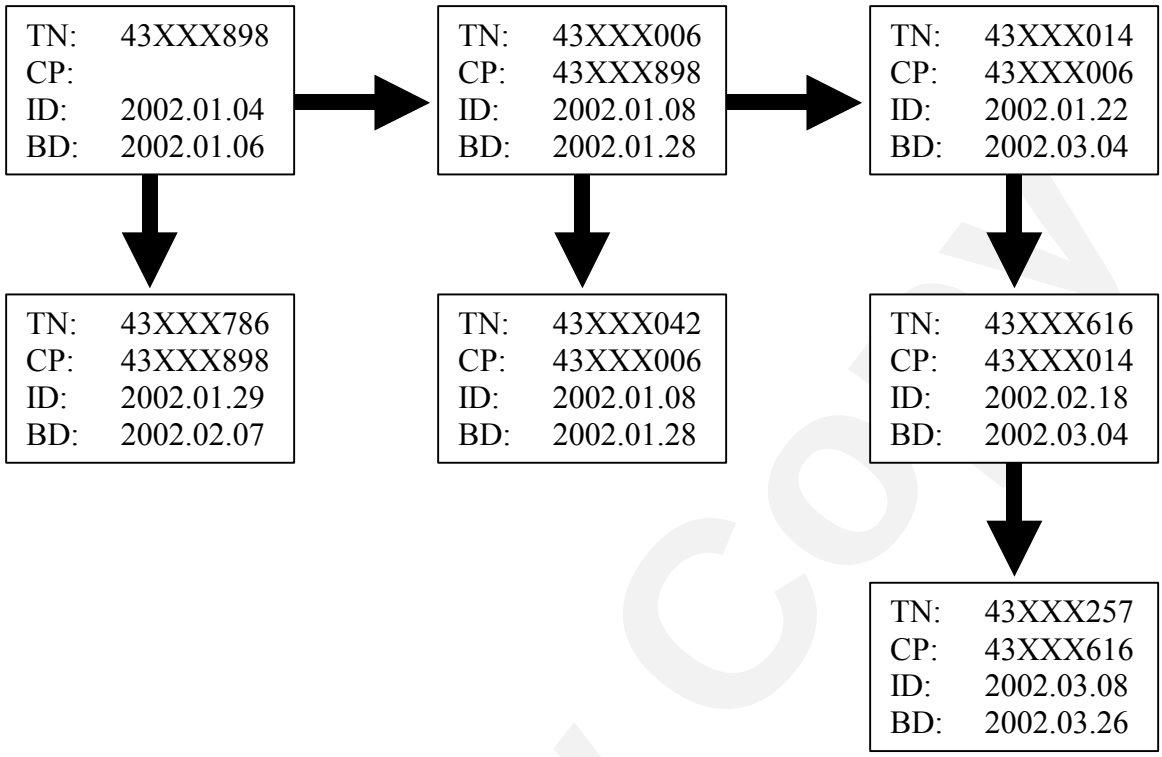

Figure 3

Review Copy

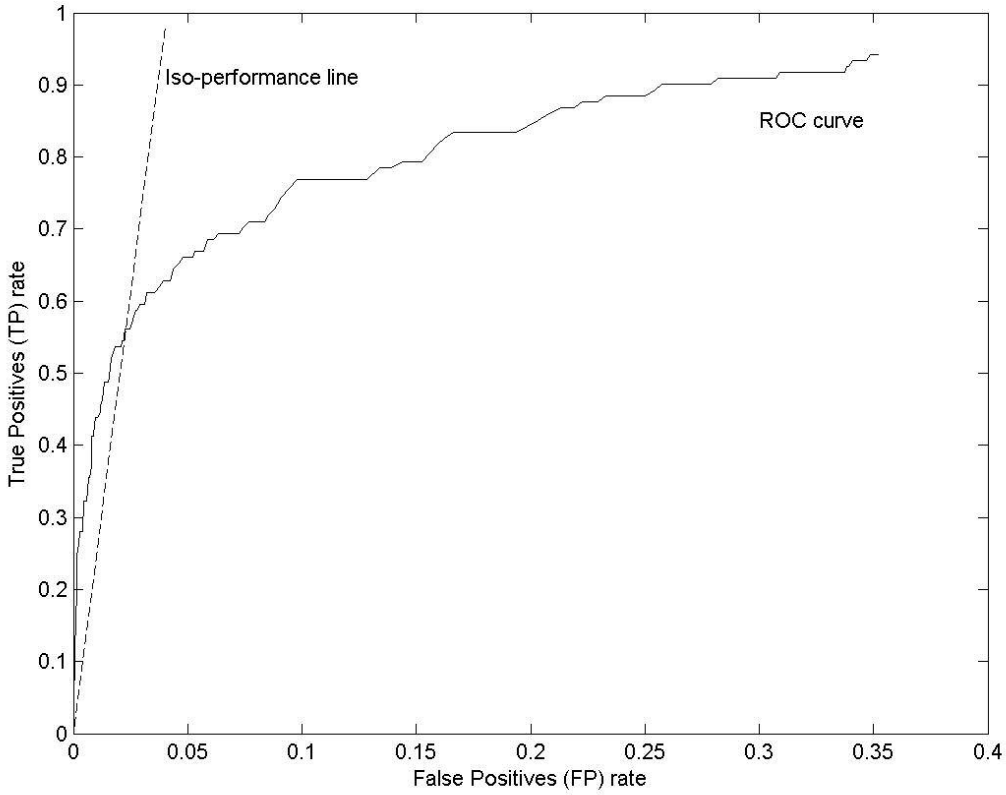
Iso-performance line
ROC curve
True Positives (TP) rate
False Positives (FP) rate
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0.05
0.15
0.25
0.35

Figure 4