ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

A hybrid system for fraud detection in mobile communications

Yves Moreau[1], Ellen Lerouge[1], Herman Verrelst[1], Joos Vandewalle[1], Christof Störmann[2], Peter Burge[3]

[1]Katholieke Universiteit Leuven, Belgium
[2]Siemens Research, München, Germany
[3]Royal Holloway University of London, UK

Abstract.

During the course of the European project "Advanced Security for Personal Communication Technologies" (ASPeCT), we have developed some rule-based and neural network architectures as a number of different fraud detection tools for GSM networks. We have now integrated these different techniques into a hybrid detection tool. We optimized the performance of the hybrid system in terms of the number of subscribers raising alarms. More precisely, we optimized performance curves showing the trade-off between the percentage of correctly identified fraudsters versus the percentage of new subscribers raising alarms. We report here on a common suite of experiments we performed on these different systems.

1. Introduction

Mobile telecommunications are an attractive target for fraudsters. In the U.S. alone, the industry estimates its loss of revenue from fraud to over 650 million dollars a year. We must thus develop systems to detect these fraudulent users.

We have been taking part to a European Commission funded ACTS project, the Advanced Security for Personal Communications Technologies (ASPeCT) project. One of the objectives of the project was the development of fraud detection and management tools for GSM networks [1]. The project chose to use three separate AI techniques [2] in four stand-alone tools, which we then combined into a powerful fraud detection engine. The Siemens research group has developed a rule-based tool and the research groups at the Katholieke Universiteit Leuven and at Royal Holloway University of London have implemented neural network tools, respectively supervised and unsupervised.

We first review the common principles behind the different fraud detection tools. We then describe the different tools briefly. After this, we describe the integration of these tools and evaluate the respective performances of the different systems.

ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

2. Fraud Scenarios and Fraud Indicators

A typical example of fraud would be subscription fraud, where a fraudster acquires a subscription to the mobile network under a false identity; and starts reselling the use of his phone to unscrupulous customers (typically for international calls to distant foreign countries) at a rate lower than the regular tariff. The fraudster accumulates a large number of expensive calls, but disappears before the bill can be collected.

After identifying possible fraud scenarios, we have identified the possible indicators that could be extracted from the information available on the network to detect fraud. An example could be an excessive number of international calls.

The information about the activity on the network is encoded in the toll tickets of all the calls placed on the network. A toll ticket is a bill issued by the network after each call, which contains all relevant information about the call. The information we use is (1) the International Mobile Subscriber Identity (the IMSI, which identifies a user uniquely), (2) the starting date of the call, (3) the starting time of the call, (4) the duration of the call, (5) number that was called, (6) the type of call (national or international).

3. Fraud detection tools

The rule-based tool is a white box approach and hence the end-user can be given a reason why the tool has flagged an alarm for a particular user. A supervised neural network is also implemented to avoid the necessity for an expert design of the rules. Two unsupervised neural networks are used to look at how a user's behavior changes over time. These systems need no prior knowledge of fraud unlike the previous two tools. The unsupervised systems are (1) an A-number analysis (which detects changes in the user behavior) and (2) an international B-number analysis (which looks at specific changes in behavior of a user making international calls).

When combined, these four tools form a powerful fraud detection tool. This integrated tool, together with its data handling module and its web-based graphical user interface (GUI) front-end, is known as BRUTUS (an acronym for the names of its components). The GUI provides access to records of suspicious users and allows the operator to look at calls of individual users to judge whether they are fraudulent and to take further action.

3.1. User profiling

A profile for each user, derived from relevant toll ticket information, is used as a basis for identifying abnormal activity which may be indicative of fraudulent use.

ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

3.1.1. Absolute or differential analysis

Existing fraud detection systems tend to interrogate sequences of toll tickets comparing a function of the various fields with fixed criteria known as triggers. Such fixed trigger systems perform what is known as an absolute analysis. But certain behavioral patterns may be indicative of fraud for one type of user (e.g., pre-paid cards) while they are considered acceptable for another (e.g., business users). Another approach to the problem is to perform a differential analysis. Here, behavioral patterns of the mobile phone are monitored, by comparing its most recent activities with a history of its usage. Triggers are activated when usage patterns of the mobile phone change significantly over a short period of time.

A differential usage system requires information concerning the history of behavior of the user plus a more recent sample of the mobile phone activity. This information is stored in a long term history, the User Profile History (UPH), and a short term behavior pattern, the Current User Profile (CUP).

3.1.2. Current User Profile (CUP) in the rule-based system and supervised neural network

At any time the CUP's values reflect aggregated information for the last day. The CUP realized in the rule-based and neural network FDT uses the hopping window technique with 24 hour intervals. This results in the following vector of features, which is the input to the classifier: (1) the number of days since first activity, (2) the short-term mean/standard deviation of the duration of national/international calls, (3) the short-term mean/standard deviation of the call interval between national/international calls, (4) the long-term mean/standard deviation of the duration of national/international calls, and (5) the long-term mean/standard deviation of the call interval between national/international calls. On each incoming toll ticket, the fields of the CUP are updated.

3.1.3. User Profile History (UPH)

The UPH will be updated whenever a CUP has completed its life span. This is done by decaying the current UPH-values and adding "fresh" data taken from the existing CUPs. Exponential fading is used to compute the components of the UPH:

$$\mathrm{UPH}_{\mathrm{new}} = (1-\alpha)\mathrm{UPH}_{\mathrm{old}} + \alpha\mathrm{CUP}$$

with fading-factors $\alpha, 0 \leq \alpha < 1$. An important quality of the UPH is that it slowly adapts to the user behavior in the following way; smooth and long-lasting changes in behavior—usual for normal behavior–are adopted by the UPH while significant short-term changes—indicative of fraud—are still detected by a contrast between the CUP and UPH.

ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

4. Supervised neural network tool

The classifier we use in the fraud detection engine is a multilayer perceptron with a single hidden layer of logistic-sigmoid neurons. We have a two-class problem for which we use a weighted squared error as cost function. We perform the minimization using the Levenberg-Marquardt algorithm. To maximize the performance on previously unseen data we use weight decay as regularization procedure. We determine the optimal weights using the error minimization procedure, but we use a multi-start procedure to avoid local minima. Furthermore, we have to repeat this procedures for different architectures of the neural network (here, the number of hidden neurons) to determine the optimal one. Once we have found the optimal neural network, we simply have to use it on top of the front-end and it will produce an alarm value between 0 and 1 each time a toll ticket is presented to the fraud detection tool. A more detailed description of the neural network is available in [9].

5. Rule-based tool

Most of today's fraud detection tools are either rule-based or at least comprise a rule-based detection component. A rule-based approach allows detecting the definite frauds with a low rate of false alarms. Moreover, the rule-based tool can easily provide reasons for an alarm being raised. The rule-based tool uses the profiling strategy described above and features similar to those of the supervised neural network. The rules for the trigerring of an alarm are designed manually by an expert. A more concrete description of the tool can be found in [2].

6. Unsupervised neural network tools

As in the first two tools, we generate two profile records for each user by considering two different time spans over the toll tickets. The unsupervised neural network is based on a prototyping technique [7] reminiscent of vector quantization or clustering. Prototyping is a method of forming an optimal discrete representation of a naturally continuous variable. For lack of space, we refer to [6] for a detailed description of the unsupervised neural network tool.

7. B - Number analysis tool

The B-number analysis tool monitors the destination countries of calls on a per subscriber basis. The destinations of calls (the B-Number) are weighted differently so that well known destinations for fraudulent calls can be given special attention. The profile is maintained as a probability distribution of

ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

the call destination for the CUP and UPH. The fraud engine takes the B-number profile record consisting of the CUP and UPH as input and calculates a modified distance over all the entries of the profile record. More details are available in [4].

8. Discussion

The rule based tool can give explanations on how and why the alarm was raised. On the other hand, this tool is not very flexible. A thorough understanding of the problem of fraud detection is needed to construct a set of satisfying rules. This also means, that with each new fraud scenario that shows up, new rules have to be made. The main advantage of the neural network approach is the flexibility and adaptability, which should make it easy to cope with new fraud scenarios. The unsupervised neural network tool has the potential to detect new types of fraud as and when they occur. The system is able to do this because it is not being trained to recognize specific fraud scenarios, but rather altered or unusual usage.

9. BRUTUS: a hybrid detection tool

We have then integrated the existing fraud detection techniques into a hybrid detection tool, which we call BRUTUS. Indeed, our ultimate purpose is to build a common framework where all the tools will be integrated so we can combine the respective strengths of the different tools [3].

9.1. Combination of the different tools

The integrated tools process toll tickets in a sequential manner. Toll tickets flow through the architecture accumulating information pertaining to the analysis as it takes place. Subsequent modules have the ability to use this information in support of their own decisions. However, the working in this development prototype is still parallel. The alarms generated by the fraud detection tool are then handled by an intelligent monitoring tool, which serves as a (graphical) interface to the human operator.

As an intermediate step towards full information exchange between the different modules, the approach followed by ASPeCT in the technical trials only integrated the respective alarm levels A_1, A_2, A_3, A_4. The ultimate performance is based on a combined alarm level $A_{\text{com}} = f(w_1, w_2, w_3, w_4, A_1, A_2, A_3, A_4)$. We opted for a well-known approach from statistical theory: logistic regression modeling (which we can seen as a neural network with a single hidden neuron). The combination function has then the form $f = 1/(1+\exp(-(\sum_i w_i A_i))$. The advantages of this combination function are (1) the number of parameters to estimate is low and (2) the resulting parameters w_i are also statistically meaningful in that contributions with large parameter values contribute exponentially more to the probability of fraud than contributions with low parameter

ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

values. We estimated the parameters of the integrated model by optimizing the ROC performance measure that we describe in the next section. This optimization was performed using simulated annealing because this performance measure is not differentiable.

10. Results

The available examples of frauds were collected from the TACS network of a mobile phone operator. This data contains a total of 317 fraudulent users who generated 131.594 toll tickets. The example of live network data consists of a three-month download (from 16-02-1998 to 16-05-1998) from 20.212 users for a total of about three million toll tickets. The number of fraudulent users in the live data is expected to be low. From the live data, we selected 562 random users for which we retained the toll tickets from the first 40 days. We limited the time range to 40 days so that the average period of activity of the normal users is the same as the average lifetime of a fraudulent example. We selected 500 users from the live data purely at random and added another 62 random users from those users who had placed international call. The evaluation data is split into a training set and a test set. We use the test set to evaluate the different fraud detection tools and also to evaluate the integrated tool.

10.1. Evaluation of technical effectiveness

The striking feature of fraud detection is the importance of the trade-off between detection of fraudulent users and the production of false alarms. Indeed, telecom network operators are, from a commercial point of view, extremely cautious about unduly bothering good subscribers. Moreover, even levels of false alarms that would be excellent for many applications (let us say, one percent of misclassification), would be completely unacceptable in our case because of the high number of users. Therefore, the problem of the fraud detection tool will be to find the right balance between false alarms and correct detection.

The Receiver-Operating Characteristic [8] plots the percentage of correct detection of fraudulent users (sensitivity), versus the percentage of false alarms for non-fraudulent users (specificity) for varying values of the threshold. The index of performance that we need to maximize is the surface under the curve. Such a trade-off curve will give the user of the fraud detection tools control over the fraud detection rate and false alarm rate.

10.2. Performance of the individual tools and integration

In this section we report on the individual performance of the tools, see [3, 5] for a more in-depth analysis. For the B-number tool, the area under the ROC-curve is 54,54% (training set) and 57,97% (test set). For the unsupervised neural network tool, mainly low alarm values are excited. The area under the ROC-curve is 54,56% (training set) and 57,75% (test set). For the supervised

ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

neural network tool, the area under the ROC-curve is 87,17% (training set) and 85,63% (test set). The results of the supervised neural network are reported in Figure 10.2. For the rule based tool, the area under the ROC-curve is 81,24% (training set) and 85,63% (test set). For the integration of the four systems, the area under the ROC-curve is 88,45% (training set) and 90.08% (test set).

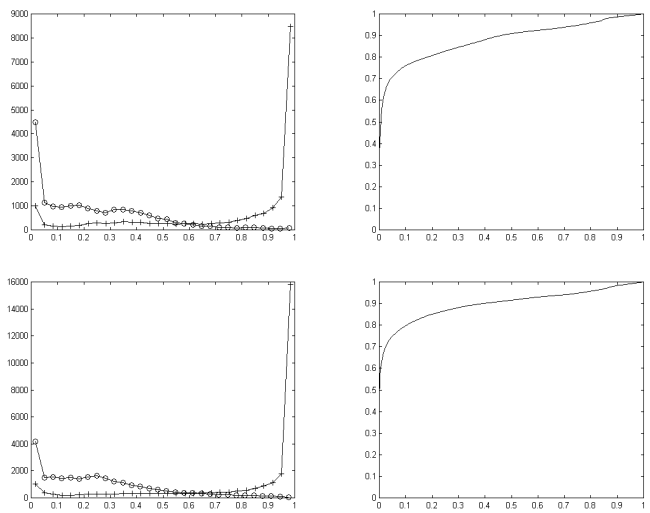

Figure 1: Results of the supervised neural network tool: histogram of the alarm on fraudulent and normal toll tickets (left) and ROC curve (right), for training (top) and test sets (bottom). The area under the ROC-curve is 87,17% (training set) and 85,63% (test set).

10.3. Integration of tools

The integration of the tools allows a two-percent increase of the ROC performance against the best sub-module (supervised neural network), which is statistically significant ($p < 0.0001$) [8]. But the main improvement is that the behavior of the integrated tool in the region of low false positive rates has bettered significantly. At 0.02% false alarms, the supervised neural network detects 40% of the fraudulent users (both on the training and test sets) while the integrated tool detects 45% of the fraudulent users on the training set and 50% on the test set.

The user acceptability of the fraud tool has been evaluated in a number of ways. The integrated tool generated a list of 27 suspicious users over the first forty-day period of data—of which 3 were fraudulent. The network operator considered this proportion of false alarms useful enough for daily operations.

ESANN'1999 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 21-23 April 1999, D-Facto public., ISBN 2-600049-9-X, pp. 447-454

11. Conclusions

For the ASPeCT project, we developed three approaches to fraud detection, namely rule-based, supervised learning and unsupervised learning. Each of these tools used profiling techniques and based its decision on a differential analysis. We have reported on the final stage of the project where we integrated all four tools into an overall fraud detection system. This tool was developed under the name BRUTUS and was demonstrated on the available data. The users who raised the most significant alarms were then given to the operators for them to investigate an investigation and ascertain whether any genuine fraud had been detected. The results obtained in the technical trials were at the level of a pre-competitive product.

Acknowledgments

Joos Vandewalle is a full Professor at the K.U. Leuven. Yves Moreau is a post-doctoral researcher of the K.U.Leuven. Herman Verrelst is a research assistant of the IWT. This work is supported by several institutions: (1) The European ACTS ASPeCT project, (2) The Flemish Government: Concerted Research Action GOA-MIPS, (3) The Belgian State, Prime Minister's Office: Interuniversity Poles of Attraction IUAP P4-02: Modeling, Identification, Simulation and Control of Complex Systems.

References

[1] ACTS AC095, ASPeCT, "Definition of fraud detection concepts", 1996.

[2] ACTS AC095, ASPeCT, "Fraud management tools: first prototype", 1997.

[3] ACTS AC095, ASPeCT, "Fraud management tools: evaluation", 1997.

[4] ACTS AC095, ASPeCT, "Fraud detection concepts: final report", 1998.

[5] ACTS AC095, ASPeCT, "Report on final trial and demonstration", 1998.

[6] P. Burge, J. Shawe-Taylor, "Detecting cellular fraud using adaptive prototypes", *Proceedings of the AAAI-97 Workshop on AI Approaches to Fraud Detection and Risk Management*, Providence, RI, USA, 1997.

[7] "Self-organization of neurons described by the second maximum entropy principle", *Proceedings of the 1st IEE International Conference on Artificial Neural Networks*, London, 1989.

[8] J.A. Hanley, B.J. McNeil, "A method of comparing the areas under receiver-operating characteristic curves derived from the same cases", *Radiology*, Vol. 148, pp. 839–843, 1983.

[9] Y. Moreau, H. Verrelst, J. Vandewalle, "Detection of mobile phone fraud using supervised neural networks: a first prototype", *International Conference on Artificial Neural Networks 97*, pp. 1065-1070, 1997.