A classification approach to identify fraudulent electricity usage in Nelson Mandela Bay Municipality

Gary Sharp*
NMMU, Port Elizabeth, South Africa – gary.sharp@nmmu.ac.za

Sisa Pazi
NMMU, Port Elizabeth, South Africa – s209067822@nmmu.ac.za

Abstract

This study investigates the classification ability of a support vector machine in trying to identify fraudulent behaviour on electricity consumption in Nelson Mandela Bay, South Africa. The ability is assessed on commonly used performance measures and results discussed in relation to the limitations of the data and previous studies. The classifier exhibits reasonable performance with successful identification close to 80%.

Keywords: Support vector machines; fraud; electricity theft.

1. Introduction

Each year electricity utilities record losses due to electricity theft. These losses are financial and impact both the source supplier and the intermediary parties. In South Africa the intermediaries are local municipalities which use electricity sales as a revenue source to supplement other operational expenses.

In South Africa the predominant utilities supplier is a quasi-state operation called Eskom. Electricity theft from Eskom has been estimated to exceed R2 billion annually (PowerNews, 2013). South Africa is divided into nine provincial boundaries, and thereafter partitioned into smaller units called municipalities. The Nelson Mandela Bay Municipality (NMBM) is one of the larger units and has experienced electricity losses estimated at R218 million in the financial year ending June 2014 (Algoa fm news, 2014).

Electrical losses are often partitioned into two categories, technical losses such as transmission losses, and non-technical losses such as those experienced from illegal electricity connections. In this study, we use a classification method to identify non-technical losses based on consumer electricity theft. The data in this assessment is used to develop a support vector machine (SVM) classifier based on load profiles related to non-technical activities of individual households. Customer's data with known fraudulent and clean cases were used to train the classifier which was then used to classify customers with unknown class labels.

2. Literature review

The literature demonstrates that there have been several studies to assess non-technical losses. These studies include Monedero, I., Biscarri, F., Leon, C., Biscarri, J., & Millan, R. (2006) who constructed a prototype for the detection of non-technical losses by means of neural networks. They used unsupervised neural networks to distinguish fraudulent customers from customers with normal consumption patterns. Simultaneously they used an outlier detection approach in an attempt to identify possible fraudsters. Using these two techniques they developed methodologies for the detection of non-technical losses for Endesa, an important electrical company in Spain. Using these methodologies, the company reported a success rate of around 50% when identifying fraudulent customers. Nizar, A.H., Dong, Z.Y., Zhao, J.H., & Zhang, P. (2007) used two methods, decision trees and a naïve Bayesian approach to determine what arrangement of profile data provided the most successful fraudulent detection scenario. The data was rearranged into different time frames that was averaged

over days, weeks, months and years. The decision tree algorithm was found to be more accurate when the data was arranged into daily and monthly averages. The limitation to the success of the decision tree method was the computational time, which was much slower than the naïve Bayesian method used. For both methods, the lowest accuracy was found when the data was arranged on a yearly basis. The highest accuracy for both classification systems was when daily average consumptions were used. Nagi, J., Mohammed, A.M., Yap, K.S., Tiong, S.K,. & Ahmed, S.K. (2008) proposed a new approach towards non-technical loss analysis. They proposed the use of SVM as a classifier. The objective of the study was to develop an SVM classifier that identified transactions of potential fraudulent customers in order to effect an onsite inspection. This study, for Tenaga Nasional Berhad (TNB) in Malaysia, reported that out of the total number of candidates shortlisted by the SVM classifier, 53% were found to be fraudulent cases upon inspection.
These research studies all reported satisfactory results when daily average consumption data was used. Based on these findings, and the success of the SVM classifier, it was decided to follow a similar approach in an attempt to identify fraudulent activity in the NMBM.

3. Methodology and the data

Data available for this study was provided by the NMBM. This was a collaborative study to assist the municipality who had out-sourced the responsibility of tracking and identifying non-technical losses to a company who had the specialist expertise to expose the fraudulent clients. The data was incomplete and required considerable cleaning prior to modelling. The data consisted of 219 940 customers for a period of 21 months from March 2013 to November 2014. After cleaning and coding only 110 740 (50.4%) customer records were found usable for statistical assessment. This reduction in the number of useable observations was a result of many inconsistencies in the source database. There were many zero consumption users in the source database. These were all removed based on valid reasons. Similarly there were many duplicate entries in the database, again, these were removed.
Table 1 shows the structure of the data for a particular month. The “FBE Customer” column indicates whether or not a customer is eligible for free basic electricity. Customers that are eligible for free basic electricity are given a “Yes”, otherwise a customer is given a “No”. The tariff code is assigned to a customer based on where the customer resides. Units billed indicate the customer’s electricity consumption for the whole month in kWh. The customer’s electricity consumptions for all the months on the database were merged into one table using SQL Server 2014 Management Studio queries. Each customer was given a unique identity ranging from 1 to 219 940. The monthly customer consumption is shown in Table 2.

Table 1: Electricity customer’s information for a particular month

Meter Number	FBE Customer	Tariff Code	Units Billed (kWh)	Address
Meter Number 1	Yes	T01	36	Address 1
⋮	⋮	⋮	⋮	⋮
Meter Number 219940	No	A32	256.71	Address 219940

Table 2: Electricity customer’s consumptions in kWh

ID	Mar-13	Apr-13	May-13	Jun-13	...	Nov-14
1	467.7	470.5	432.9	418.0	...	512.7
2	896.4	843.1	802.4	799.1	...	763.1
3	198.6	176.4	171.0	197.8	...	119.0
4	187.7	243.8	253.9	243.9	...	246.4
5	0	0	0	0	...	0
6	433.5	357.6	316.4	501.3	...	344.1
⋮	⋮	⋮	⋮	⋮	...	⋮
219940	595.9	676.6	679.9	710.5	...	788.8

The database was cleaned, the unique IDs recoded and the daily averages determined for each ID, these observations are given in Table 3.

Table 3: Load profiles of 110 740 customers for 21 periods					
ID*	Mar 2013	Apr 2013	May 2013	...	Nov 2014
1	6.95	6.45	6.45	...	7.67
2	21.49	23.25	21.07	...	12.00
3	5.18	4.84	2.82	...	4.68
⋮	⋮	⋮	⋮	...	4.30
110 740	22.84	11.54	15.92	...	16.08

The daily average consumptions were normalised for the SVM classifier. These daily average consumptions were normalised such that all the values fall on the interval [0,1]. The daily average consumptions were normalised by $NL = \frac{L - \min(L)}{\max(L) - \min(L)}$, where L represent the daily average consumption of a customer, $\min(L)$ and $\max(L)$ minimum and maximum values in the period.
As with any supervised learning technique, a SVM is trained first. The trained machine is then used to classify (predict) new data, this approach was followed using the software, Statistica ver12. The choice of kernel function for estimation was a Gaussian kernel with single parameter. For the purpose of this study the correct rate, error rate, sensitivity and specificity were used to evaluate the performance of the classifier. The four performance measures are defined as:

$$\text{correct rate} = \frac{\text{The number of correctly classified samples}}{\text{The total number of classified samples}} \times 100\%,$$

$$\text{error rate} = \frac{\text{The number of incorrectly classified samples}}{\text{The total number of classified samples}} \times 100\%,$$

$$sensitivity = \frac{\text{The number of correctly classified clean customers}}{\text{The total number of clean customers}} \times 100\%, \text{and}$$

$$specificity = \frac{\text{The number of correctly classified fraudulent customers}}{\text{The total number of fraudulent customers}} \times 100\%.$$

These measures are used in research on credit/debit card fraud and provide an indication of the modelling ability of the methodology used.

4. Results

Inspections were performed on randomly selected customers in the Nelson Mandela Bay region. These inspections were done to identify fraudulent customers. Given the time and financial restrictions, it was infeasible to select a large sample of customers for inspection, hence only 2136 customers were inspected. The status of the electricity meters for these customers was tested using electrical engineering devices. All those customers that were found to have committed fraud were tagged in the database accordingly. This data was used to train the SVM classifier. Consequently classification was done on the remaining customers in the database.
The 2136 customer’s data, of which approximately a third were tagged as fraudulent, were randomly divided into training and test data. The partitioned data is shown in Table 4, with 75% of the data used as a training set and 25% used as test set. Clean customers refer to the customers that were not guilty of electricity fraud after inspection.
The SVM classifier was then trained using 1602 observations and the classification summary for the training set shown in Table 5. The correct rate and error rate for the training set were 82.83% and 17.17% respectively. The sensitivity was found to be 82.69% and the specificity was 83.14%.

Table 4: The summary of the data used for SVM classifier			
Category	Total	Training	Test
Clean	1454	1086	368
Fraudulent	682	516	166
Total	2136	1602	534

Table 5: Classification summary for the training data					
Category	Total	Correct	Incorrect	Correct (%)	Incorrect (%)
Clean	1086	898	188	82.69	17.31
Fraudulent	516	429	87	83.14	16.86

Thereafter the trained classifier was used to classify the test data to check the classification ability on an unseen dataset. Table 6 shows the classification summary for the test data.

Table 6: Classification summary for the test data					
Category	Total	Correct	Incorrect	Correct (%)	Incorrect (%)
Clean	368	290	78	78.80	21.20
Fraudulent	166	130	36	78.31	21.69

The correct rate for the test data was 78.65% and the error rate was 21.35%. The test data set had a sensitivity of 78.80% and a specificity of 78.31%. These results are marginally less than the training data, but demonstrate reasonable ability of the classifier.

5. Discussion and Conclusion

The classifier trained provided some support in its capabilities. The test data set yielded a fraud detection rate, defined as specificity of 78.3%. These results compare favourably with results of Monedero et al. (2006) and Nizar et al. (2007). The limitations to this study was the relatively small proportion of customers visited, the choice of the 2136 customers was judgemental and dependent on whether or not customers could be contacted. It was impractical to return to a site in cases where customers were not available.

The SVM classification abilities are affected by the choice of training and test sets. In statistical learning the number outliers or unusual observations such as fraudulent customers is often very small compared to sample size. This imbalance can have negative impact on the classifier. As the true proportion of fraudulent customers is unknown, the relatively high proportion in the sampled data needs further research.

The specificity is the indication of how well the SVM classifier classifies fraudulent customers. The specificity of 78.3% for the unseen data implies that the classifier identified more than ¾ of the fraudulent customers. Therefore the SVM classifier that was developed can be used for classification.

Future research can include improving the specificity. The specificity may be improved by using a period longer than 21 months or by using a different data mining technique. The grid-search method can also be used to find optimal parameter estimates of the Gaussian kernel.

References

Algoa fm news, 16:28 (GMT+2), Friday, 11 July 2014. Former mayor's company to help curb electricity theft in Nelson Mandela Bay. www.algoafm.co.za/article.aspx?id=100018.

Monedero, I., Biscarri, F., Leon, C., Biscarri, J., & Millan, R. (2006). Detection of non-technical losses in electrical consumption using neural networks and statistical techniques, Conference: Computational Science and its Applications, M. Gavrilova et al. (Eds.): ICCSA, pp 725-734.

Nagi, J., Mohammed, A.M., Yap, K.S., Tiong, S.K,. & Ahmed, S.K. (2008). Non-technical loss analysis for detection of electricity theft using support vector machines. In the proceedings of the 2nd IEEE International conference on power and energy (PECon 08), 1 – 3 Dec 2008, Johor Bahru, Malaysia, pp 907 - 912.

Nizar, A.H., Dong, Z.Y., Zhao, J.H., & Zhang, P. (2007). A data mining based non-technical loss analysis method, IEEE Power engineering society general meeting, 24 – 28 June 2007, Tampa, Florida, pp 1-8.

PowerNews, July 2013. Municipalities, utilities meet in bid to protect dwindling revenues. www.powernews.co.za/news/entry/municipalities-utilities-meet-in-bid-to-protect-dwindling-revenues.