A Frequency-domain-based Pattern Mining for Credit Card Fraud Detection

Roberto Saia and Salvatore Carta
Dipartimento di Matematica e Informatica, Università di Cagliari, Italy

Keywords: Business Intelligence, Fraud Detection, Pattern Mining, Fourier, Metrics.

Abstract: Nowadays, the prevention of credit card fraud represents a crucial task, since almost all the operators in the E-commerce environment accept payments made through credit cards, aware of that some of them could be fraudulent. The development of approaches able to face effectively this problem represents a hard challenge due to several problems. The most important among them are the heterogeneity and the imbalanced class distribution of data, problems that lead toward a reduction of the effectiveness of the most used techniques, making it difficult to define effective models able to evaluate the new transactions. This paper proposes a new strategy able to face the aforementioned problems based on a model defined by using the Discrete Fourier Transform conversion in order to exploit frequency patterns, instead of the canonical ones, in the evaluation process. Such approach presents some advantages, since it allows us to face the imbalanced class distribution and the cold-start issues by involving only the past legitimate transactions, reducing the data heterogeneity problem thanks to the frequency-domain-based data representation, which results less influenced by the data variation. A practical implementation of the proposed approach is given by presenting an algorithm able to classify a new transaction as reliable or unreliable on the basis of the aforementioned strategy.

1 INTRODUCTION

Studies conducted by *American Association of Fraud Examiners*[1] show that the financial frauds represent the 10-15% of the entire fraud cases, by involving the 75-80% of the entire financial value, with an estimated average loss per fraud case of 2 million of dollars, in the USA alone. Fraud represents one of the major issues related to the use of credit cards, an important aspect considering the exponential growth of the E-commerce transactions. For these reasons, the research of effective approaches able to detect the frauds has become a crucial task, because it allows the involved operators to eliminate, or at least reduce, the economic losses.

Since the fraudulent transactions are typically less than the legitimate ones, the data distribution is highly unbalanced and this reduces the effectiveness of many machine learning strategies Japkowicz and Stephen (2002). Such problem is worsened by the scarcity of information that characterizes a typical financial transaction, a scenario that leads toward an overlapping of the classes of expense of a user Holte et al. (1989).

There are many state-of-the-art techniques designed to perform the fraud detection task, for instance, those that exploit the Data Mining Lek et al. (2001), the Artificial Intelligence Hoffman and Tessendorf (2005), the Fuzzy Logic Lenard and Alam (2005), the Machine Learning Whiting et al. (2012), and the Genetic Programming Assis et al. (2010) techniques.

Almost all the aforementioned techniques mainly rely on the detection of outliers in the transactions under analysis, a basic approach that could lead toward many wrong classifications (i.e., reliable transactions classified as unreliable). Most of these wrong classifications happen due to the absence of extensive criteria during the evaluation process, since many techniques are not able to manage some non-numeric transaction features during the evaluation process, e.g., one of the most performing approaches, such as Random Forests, is not able to manage types of data that involve a large number of categories.

The idea behind this paper is a new representation of the data obtained by using the *Fourier transformation* Duhamel and Vetterli (1990) in order to move a time series (the sequence of discrete-time data given

[1] http://www.acfe.com

Saia, R. and Carta, S.
A Frequency-domain-based Pattern Mining for Credit Card Fraud Detection.
DOI: 10.5220/0006361403860391
In *Proceedings of the 2nd International Conference on Internet of Things, Big Data and Security (IoTBDS 2017)*, pages 386-391
ISBN: 978-989-758-245-5

by the feature values of a transaction) in the frequency domain, allowing us to analyze the data from a new point of view.

It should be observed that the proposed evaluation process involves only the past legitimate transactions, presenting some advantages: first, it operates in a proactive way, by facing the imbalanced class distribution and the *cold-start* (i.e., scarcity or total absence of fraudulent transaction cases) problems; second, it reduces the problems related to the data heterogeneity, since the data representation in the frequency domain is more stable than the canonical one, in terms of capability of recognizing a peculiar pattern, regardless of the value assumed by the transaction features.

The contributions of this paper are as follows:

(i) definition of the *time series* to use in the Fourier process, on the basis of the past legitimate transactions;

(i) formalization of the comparison process between the *time series* of an unevaluated transaction and those of the past legitimate transactions, in terms of difference between their frequency magnitude;

(i) formulation of an algorithm, based on the previous comparison process, able to classify a new transaction as *reliable* or *unreliable*.

The remainder of the paper is organized as follows: Section 2 introduces the background and related work; Section 3 provides a formal notation, makes some assumptions, and defines the faced problem; Section 4 describes the steps necessary to define the proposed approach; Section 5 gives some concluding remarks.

2 BACKGROUND AND RELATED WORK

Many studies consider the frauds as the biggest problem in the E-commerce environment. The challenge faced by the fraud detection techniques is the classification of a financial transaction as reliable or unreliable, on the basis of the analysis of its features (e.g., description, date, total amount, etc.).

The study presented in Assis et al. (2010) indicates how in the fraud detection field there is a lack of public real-world datasets, configuring a relevant issue for those who deal with the research and development of new and more effective fraud detection techniques. This scenario mainly depends on the restrictive policies adopted by the financial operators, which for competitive or legal reasons do not provide information about their business activities. These policies are also adopted because the financial data are composed by real information about their customers, which even anonymized may reveal potential vulnerabilities related to the E-commerce infrastructure.

Supervised and Unsupervised Approaches. In Phua et al. (2010) it is underlined how the *unsupervised* fraud detection strategies are still a very big challenge in the field of E-commerce. In spite of the fact that every *supervised* strategy in fraud detection needs a reliable training set, the work proposed in Bolton and Hand (2002) takes in consideration the possibility to adopt an *unsupervised* approach during the fraud detection process, when no dataset of reference containing an adequate number of transactions (legitimate and non-legitimate) is available.

Data Heterogeneity. Pattern recognition can be considered an important branch of the machine learning field. Its main task is the detection of patterns and regularities in a data stream, in order to define an evaluation model to exploit in a large number of real-world applications Garibotto et al. (2013). One of the most critical problems related to the pattern recognition tasks is the data heterogeneity. Literature describes the data heterogeneity issue as the incompatibility among similar features resulting in the same data being represented differently in different datasets Chatterjee and Segev (1991).

Data Unbalance. One of the most important problems that makes the definition of effective models for the fraud detection difficult is the imbalanced class distribution of data Japkowicz and Stephen (2002); He and Garcia (2009). This issue is given by the fact that the data used in order to train the models are characterized by a small number of default cases and a big number of non-default ones, a distribution of data that limits the performance of the classification techniques Japkowicz and Stephen (2002); Brown and Mues (2012).

Cold Start. The *cold start* problem Donmez et al. (2007) arises when there is not enough information to train a reliable model about a domain. In the context of the fraud detection, such scenario appears when the data used to train the model are not representative of all classes of data Attenberg and Provost (2010) (i.e., default and non-default cases).

Detection Models. The *static approach* Pozzolo et al. (2014) represents a canonical way to operate in order to detect fraudulent events in a stream of transactions. This approach divides the data stream into blocks of the same size, and the user model is trained by using a certain number of initial and contiguous blocks of the sequence, which are used to infer the future blocks. The *updating approach* Wang et al. (2003), instead, when a new block appears, trains the user model by using a certain number of latest and contiguous blocks of the sequence, then the model can

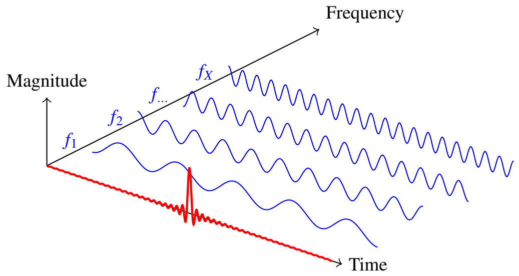

Figure 1: Time and Frequency Domains.

be used to infer the future blocks, or it can be aggregated into a big model composed by several models. In another strategy, based on the so-called *forgetting approach* Gao et al. (2007), a user model is defined at each new block, by using a small number of non fraudulent transactions, extracted from the last two blocks, but by keeping all previous fraudulent ones. Also in this case, the model can be used to infer the future blocks, or it can be aggregated into a big model composed by several models.

The main disadvantages related to these approaches of user modeling are: the incapacity to track the changes in the users behavior, in the case of the *static approach*; the ineffectiveness to operate in the context of small classes, in the case of the *updating approach*; the computational complexity in the case of the *forgetting approach*. However, regardless of the used approach, the problem of the heterogeneity and unbalance of the data still remains unaltered.

Discrete Fourier Transform. The basic idea behind the approach proposed in this paper is to move the process of evaluation of the new transactions (*time series*) from their canonical time domain to the frequency one, in order to obtain a representative pattern composed by their frequency components, as shown in Figure 1. This operation is performed by recurring to the *Discrete Fourier Transform* (*DFT*), whose formalization is reported in Equation 1, where i is the imaginary unit.

$$F_n \overset{\text{def}}{=} \sum_{k=0}^{N-1} f_k \cdot e^{-2\pi ink/N}, \quad n \in \mathbb{Z} \tag{1}$$

As result we obtain a set of sinusoidal functions, each corresponding to a particular frequency component. It is possible to return to the original time domain by using the *inverse Fourier transform* shown in Equation 2.

$$f_k = \frac{1}{N} \sum_{n=0}^{N-1} F_n \cdot e^{2\pi ikn/N}, \quad n \in \mathbb{Z} \tag{2}$$

3 PRELIMINARIES

Formal notation, assumptions, and problem definition are stated in the following:

3.1 Formal Notation

Given a set of classified transactions $T = \{t_1, t_2, \dots, t_N\}$, and a set of features $V = \{v_1, v_2, \dots, v_M\}$ that compose each $t \in T$, we denote as $T_+ \subseteq T$ the subset of legitimate transactions, and as $T_- \subseteq T$ the subset of fraudulent ones. We also denote as $\hat{T} = \{\hat{t}_1, \hat{t}_2, \dots, \hat{t}_U\}$ a set of unclassified transactions. It should be observed that a transaction only can belong to one class $c \in C$, where $C = \{reliable, unreliable\}$. Finally, we denote as $F = \{f_1, f_2, \dots, f_X\}$ the frequency components of each transaction obtained through the *DFT* process.

3.2 Assumptions

A periodic wave is characterized by a frequency f and a wavelength λ (i.e., the distance in the medium between the beginning and end of a cycle $\lambda = \frac{w}{f_0}$, where w stands for the wave velocity), which are defined by the repeating pattern, the non-periodic waves that we take into account during the *Discrete Fourier Transform* process do not have a frequency and a wavelength. Their fundamental period T is the period where the wave values were taken and sr denotes their number over this time (i.e., the acquisition frequency).

Assuming that the time interval between the acquisitions is equal, on the basis of the previous definitions applied in the context of this paper, the considered non-periodic wave is given by the sequence of values $v_1, v_2, \dots, v_M$ with $v \in V$, which composes each transaction $t \in T_+$ (i.e., the past legitimate transactions) and $\hat{t} \in \hat{T}$ (i.e., the unevaluated transactions), and that representing the *time series* taken into account. Their fundamental period T starts with v_1 and it ends with v_M, thus we have that $sr = |V|$; the sample interval si is instead given by the fundamental period T divided by the number of acquisition, i.e., $si = \frac{T}{|V|}$.

We compute the *Discrete Fourier Transform* of each *time series* $t \in T_+$ and $\hat{t} \in \hat{T}$, by converting their representation from the time domain to the frequency one. The obtained frequency-domain representation provides information about the signal's magnitude and phase at each frequency. For this reason, the output (denoted as x) of the *DFT* computation is a series of complex numbers composed by a real part x_r and an imaginary part x_i, thus $x = (x_r + ix_i)$. We can obtain the x magnitude by using $|x| = \sqrt{(x_r^2 + x_i^2)}$ and the x phase by using $\varphi(x) = \arctan\left(\frac{x_i}{x_r}\right)$, although in the context of this paper we take into account only the frequency magnitude.

3.3 Problem Definition

On the basis of a process of comparison (denoted as Θ) performed between the frequency patterns of the *time series* related to the set T_+ and to the set $\hat{T}$, the goal of the proposed approach is to classify each transaction $\hat{t} \in \hat{T}$ as *reliable* or *unreliable*.

Given a function $evaluation(\hat{t}, \Theta)$ created to evaluate the correctness of the $\hat{t}$ classification, which returns a boolean value β (*0=misclassification*, *1=correct classification*), we can formalize our objective as the maximization of the results sum, as shown in Equation 3.

$$\max_{0 \leq \beta \leq |\hat{T}|} \beta = \sum_{u=1}^{|\hat{T}|} evaluation(\hat{t}_u, \Theta) \tag{3}$$

4 PROPOSED APPROACH

The implementation of our approach is carried out through the following steps:

- **Data Definition**: definition of the *time series* in terms of sequence of transaction features;
- **Data Evaluation**: comparison of the frequency patterns of two transactions, made by processing the related *time series*;
- **Data Classification**: presentation of the algorithm able to classify a new transaction as *reliable* or *unreliable*, on the basis of the previous comparison process.

In the following, we provide a detailed description of each of these steps.

4.1 Data Definition

In the first step of our approach we define the *time series* to use in the *Discrete Fourier Transform* process.

Formally, a *time series* represents a series of data points stored by following the time order and usually it is a sequence captured at successive equally spaced points in time, thus it can be considered a sequence of discrete-time data.

In the context of the proposed approach, the *time series* taken into account are defined by using the set of features V that compose each transaction in the T_+ and $\hat{T}$ sets, as shown in Equation 4, by following the criterion reported in Equation 5.

$$T_+ = \begin{vmatrix} v_{1,1} & v_{1,2} & \cdots & v_{1,M} \\ v_{2,1} & v_{2,2} & \cdots & v_{2,M} \\ \vdots & \vdots & \ddots & \vdots \\ v_{N,1} & v_{N,2} & \cdots & v_{N,M} \end{vmatrix} \quad \hat{T} = \begin{vmatrix} v_{1,1} & v_{1,2} & \cdots & v_{1,M} \\ v_{2,1} & v_{2,2} & \cdots & v_{2,M} \\ \vdots & \vdots & \ddots & \vdots \\ v_{U,1} & v_{U,2} & \cdots & v_{U,M} \end{vmatrix} \tag{4}$$

$$\begin{matrix} (v_{1,1}, v_{1,2}, \ldots, v_{1,M}), (v_{2,1}, v_{2,2}, \ldots, v_{2,M}), \cdots, (v_{N,1}, v_{N,2}, \ldots, v_{N,M}) \\ (v_{1,1}, v_{1,2}, \ldots, v_{1,M}), (v_{2,1}, v_{2,2}, \ldots, v_{2,M}), \cdots, (v_{U,1}, v_{U,2}, \ldots, v_{U,M}) \end{matrix} \tag{5}$$

The *time series* related to an item $\hat{t} \in \hat{T}$ will be compared to the *time series* related to all the items $t \in T_+$, by following the criteria explained in the next steps.

4.2 Data Evaluation

The frequency domain representation allows us to perform a transaction analysis in terms of the magnitude assumed by each frequency component that characterizes the transaction, allowing us to detect some patterns in the features that are not discoverable otherwise. As preliminary work, we compared the two different representation of a transaction (i.e., these obtained in the time and frequency domains), observing some interesting properties for the context taken into account in this paper, which are described in the following:

- The *phase invariance property* shown in Figure 2 proves that also in case of translation[2] between transactions, a specific pattern still exists in the frequency domain. In other words, by working in the frequency domain we can detect a specific pattern, also when it shifts along the features that compose a transaction.
- The *amplitude correlation property* shown in FIgure 3 evidences that a direct correlation exists between the feature values in the time domain and the magnitudes assumed by the frequency components in the frequency domain. It grants that our approach is able to differentiate the transactions on the basis of the values assumed by the transaction features.

Practically, the process of analysis is performed by moving the *time series* of the transactions to compare from their time domain to the frequency one, by recurring to the *DFT* introduced in Section 2.

The process of comparison between a transaction $\hat{t} \in \hat{T}$ to evaluate and a past legitimate transaction $t \in T+$ is performed by measuring the difference Δ between the magnitude $|f|$ of each component $f \in F$

[2] A translation in time domain corresponds to a change in phase in the frequency domain.

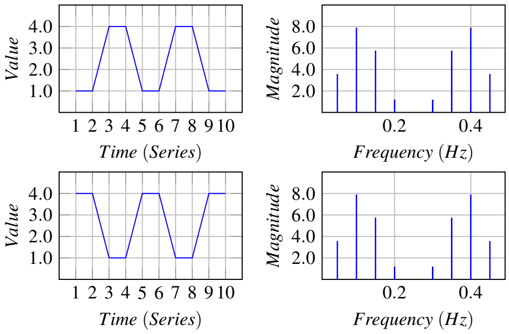

Figure 2: Phase Invariance Property.

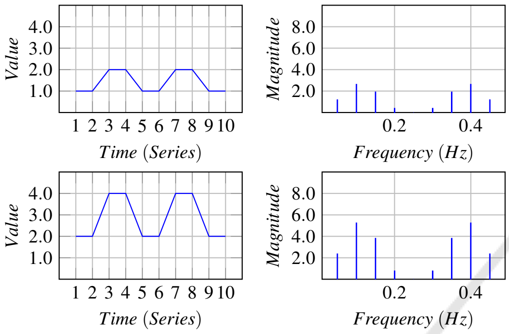

Figure 3: Amplitude Correlation Property.

in the frequency components of the involved transactions.

It is shown in the Equation 6, where f_x^1 and f_x^2 denote, respectively, the same frequency component of an item $t \in T_+$ and an item $\hat{t} \in \hat{T}$.

$$\Delta = \left(|f_x^1| - |f_x^2| \right), \; with \; |f_x^1| \geq |f_x^2| \qquad (6)$$

It should be noted that, as described in Section 4.3, for each transaction $\hat{t} \in \hat{T}$ to evaluate, the aforementioned process is repeated by comparing it to each transaction $t \in T_+$. This allows us to evaluate the variation Δ in the context of all the legitimate past cases.

4.3 Data Classification

The proposed approach is based on the Algorithm 1. It takes as input the set T_+ of past legitimate transactions and a transaction $\hat{t}$ to evaluate. It returns a boolean value that indicates the classification of the transaction $\hat{t}$ (*true*=*reliable* or *false*=*unreliable*).

From *step 1* to *step 21* we process the unevaluated transaction $\hat{t}$, by starting with the extraction of its *time series* (*step 2*), which is processed at *step 3* in order to get the frequency components. From the *step 4* to *step 14* we instead process each non-default transaction $t \in T_+$, by performing the extraction of the *time series* (*step 5*) and by obtaining its frequency components (*step 6*). The *steps* from *7* to *13* verify if the

Algorithm 1: Transaction classification.

Input: T_+=Legitimate past transactions, $\hat{t}$=Unevaluated transaction
Output: β=Classification of the transaction $\hat{t}$

```
1:  procedure TRANSACTIONCLASSIFICATION($T_+, \hat{t}$)
2:      $ts1 \leftarrow getTimeseries(\hat{t})$
3:      $F_1 \leftarrow getDFT(ts1)$
4:      for each $t$ in $T_+$ do
5:          $ts2 \leftarrow getTimeseries(t)$
6:          $F_2 \leftarrow getDFT(ts2)$
7:          for each $f$ in $F$ do
8:              if $(|F_2(f)| - |F_1(f)| \in getVariationRange(T_+, f)$ then
9:                  $reliable \leftarrow reliable + 1$
10:             else
11:                 $unreliable \leftarrow unreliable + 1$
12:             end if
13:         end for
14:     end for
15:     if $reliable > unreliable$ then
16:         $\beta \leftarrow true$
17:     else
18:         $\beta \leftarrow false$
19:     end if
20:     return $\beta$
21: end procedure
```

difference between the magnitude of each frequency components $f \in F$ of the non-default transactions and the correspondent component of the current transaction, is within the interval given by the minimum and maximum variation measured in the set T_+, by comparing all magnitudes of the current frequency component f. On the basis of the result of this operation we increase the *reliable* value (when the difference is within the interval) or the *unreliable* one (otherwise) (steps *9* and *11*). The *reliable* and *unreliable* values will determine the classification of the transaction under evaluation (*steps* from *15* to *19*), and the result is returned by the algorithm at the *step 20*.

5 CONCLUSIONS

Fraud detection techniques cover a crucial role in many financial contexts, since they are able to reduce the losses due to fraud, suffered directly by the traders or indirectly by the credit card issuers.

This paper introduces a novel fraud detection approach aimed to classify the new transactions as *reliable* or *unreliable* by evaluating their characteristics (pattern) in the frequency domain instead of the canonical one. It is performed through the Fourier transformation, defining our model by only using the past legitimate user transactions.

Such approach allows us to avoid the data unbalance problem that affects the canonical classification approaches, because it only uses a class of data during the process of definition of the model, allowing us to operate in a proactive way, by also reducing the *cold-start* problem.

Even the problems related to the data heterogeneity are reduced thanks to the adoption of a more stable model (based on the frequency components) able to recognize peculiar patterns in the transaction features, regardless of the value assumed by them.

Future work would be oriented to the implementation of the proposed approach in a real-world context, by comparing its performance to those of the most widely used state-of-the-art approaches.

ACKNOWLEDGEMENTS

This research is partially funded by Regione Sardegna under project "Next generation Open Mobile Apps Development" (NOMAD), "Pacchetti Integrati di Agevolazione" (PIA) - Industria Artigianato e Servizi - Annualità 2013.

REFERENCES

Assis, C., Pereira, A. M., de Arruda Pereira, M., and Carrano, E. G. (2010). Using genetic programming to detect fraud in electronic transactions. In Prazeres, C. V. S., Sampaio, P. N. M., Santanchè, A., Santos, C. A. S., and Goularte, R., editors, *A Comprehensive Survey of Data Mining-based Fraud Detection Research*, volume abs/1009.6119, pages 337–340.

Attenberg, J. and Provost, F. J. (2010). Inactive learning?: difficulties employing active learning in practice. *SIGKDD Explorations*, 12(2):36–41.

Bolton, R. J. and Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, pages 235–249.

Brown, I. and Mues, C. (2012). An experimental comparison of classification algorithms for imbalanced credit scoring data sets. *Expert Syst. Appl.*, 39(3):3446–3453.

Chatterjee, A. and Segev, A. (1991). Data manipulation in heterogeneous databases. *ACM SIGMOD Record*, 20(4):64–68.

Donmez, P., Carbonell, J. G., and Bennett, P. N. (2007). Dual strategy active learning. In *ECML*, volume 4701 of *Lecture Notes in Computer Science*, pages 116–127. Springer.

Duhamel, P. and Vetterli, M. (1990). Fast fourier transforms: a tutorial review and a state of the art. *Signal processing*, 19(4):259–299.

Gao, J., Fan, W., Han, J., and Yu, P. S. (2007). A general framework for mining concept-drifting data streams with skewed distributions. In *Proceedings of the Seventh SIAM International Conference on Data Mining, April 26-28, 2007, Minneapolis, Minnesota, USA*, pages 3–14. SIAM.

Garibotto, G., Murrieri, P., Capra, A., Muro, S. D., Petillo, U., Flammini, F., Esposito, M., Pragliola, C., Leo, G. D., Lengu, R., Mazzino, N., Paolillo, A., D'Urso, M., Vertucci, R., Narducci, F., Ricciardi, S., Casanova, A., Fenu, G., Mizio, M. D., Savastano, M., Capua, M. D., and Ferone, A. (2013). White paper on industrial applications of computer vision and pattern recognition. In *ICIAP (2)*, volume 8157 of *Lecture Notes in Computer Science*, pages 721–730. Springer.

He, H. and Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.*, 21(9):1263–1284.

Hoffman, A. J. and Tessendorf, R. E. (2005). Artificial intelligence based fraud agent to identify supply chain irregularities. In Hamza, M. H., editor, *IASTED International Conference on Artificial Intelligence and Applications, part of the 23rd Multi-Conference on Applied Informatics, Innsbruck, Austria, February 14-16, 2005*, pages 743–750. IASTED/ACTA Press.

Holte, R. C., Acker, L., and Porter, B. W. (1989). Concept learning and the problem of small disjuncts. In Sridharan, N. S., editor, *Proceedings of the 11th International Joint Conference on Artificial Intelligence. Detroit, MI, USA, August 1989*, pages 813–818. Morgan Kaufmann.

Japkowicz, N. and Stephen, S. (2002). The class imbalance problem: A systematic study. *Intell. Data Anal.*, 6(5):429–449.

Lek, M., Anandarajah, B., Cerpa, N., and Jamieson, R. (2001). Data mining prototype for detecting e-commerce fraud. In Smithson, S., Gricar, J., Podlogar, M., and Avgerinou, S., editors, *Proceedings of the 9th European Conference on Information Systems, Global Co-operation in the New Millennium, ECIS 2001, Bled, Slovenia, June 27-29, 2001*, pages 160–165.

Lenard, M. J. and Alam, P. (2005). Application of fuzzy logic fraud detection. In Khosrow-Pour, M., editor, *Encyclopedia of Information Science and Technology (5 Volumes)*, pages 135–139. Idea Group.

Phua, C., Lee, V. C. S., Smith-Miles, K., and Gayler, R. W. (2010). A comprehensive survey of data mining-based fraud detection research. *CoRR*, abs/1009.6119.

Pozzolo, A. D., Caelen, O., Borgne, Y. L., Waterschoot, S., and Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Syst. Appl.*, 41(10):4915–4928.

Wang, H., Fan, W., Yu, P. S., and Han, J. (2003). Mining concept-drifting data streams using ensemble classifiers. In Getoor, L., Senator, T. E., Domingos, P. M., and Faloutsos, C., editors, *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 24 - 27, 2003*, pages 226–235. ACM.

Whiting, D. G., Hansen, J. V., McDonald, J. B., Albrecht, C. C., and Albrecht, W. S. (2012). Machine learning methods for detecting patterns of management fraud. *Computational Intelligence*, 28(4):505–527.