Evaluating Credit Card Transactions in the Frequency Domain for a Proactive Fraud Detection Approach

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

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

Abstract: The massive increase in financial transactions made in the e-commerce field has led to an equally massive increase in the risks related to fraudulent activities. It is a problem directly correlated with the use of credit cards, considering that almost all the operators that offer goods or services in the e-commerce space allow their customers to use them for making payments. The main disadvantage of these powerful methods of payment concerns the fact that they can be used not only by the legitimate users (cardholders) but also by fraudsters. Literature reports a considerable number of techniques designed to face this problem, although their effectiveness is jeopardized by a series of common problems, such as the imbalanced distribution and the heterogeneity of the involved data. The approach presented in this paper takes advantage of a novel evaluation criterion based on the analysis, in the frequency domain, of the spectral pattern of the data. Such strategy allows us to obtain a more stable model for representing information, with respect to the canonical ones, reducing both the problems of imbalance and heterogeneity of data. Experiments show that the performance of the proposed approach is comparable to that of its state-of-the-art competitor, although the model definition does not use any fraudulent previous case, adopting a proactive strategy able to contrast the cold-start issue.

1 INTRODUCTION

The credit card frauds (i.e., when someone makes purchases without authorization or counterfeits a credit card) represent one of the major issues that affect the e-commerce environment, especially in this age of exponential growth of it. Authoritative studies performed by *American Association of Fraud Examiners*[1] report that this kind of fraud involves the *10-15%* of all fraud cases, for a total financial value close to *75-80%*. Only in the United States of America, this scenario leads toward an estimated average loss per fraud case of *2* million of dollars.

Such situation has given rise to a great interest by the research community in order to develop increasingly effective techniques able to detect the fraudulent transactions. This is a task that can be performed by exploiting several techniques but there are some common issues that the researchers must face. The most important between them are the imbalanced distribution and the heterogeneity of the involved data. The first issue is given by the fact that the fraudulent transactions are usually less than the legitimate ones, configuring an unbalanced distribution of data that reduces the effectiveness of the machine learning strategies (Japkowicz and Stephen, 2002).

The common criterion used in almost all the state-of-the-art approaches of fraud detection is substantially based on the comparison between the set of previous legitimate transactions of a user and the new transactions under evaluation. This is a rather trivial criterion that in many cases, due to the high heterogeneity of data, leads toward misclassifications. In order to overcome this problem, a fraud detection approach should be able to use as much as possible information about the transactions during the evaluation process, but this is not always possible due to the inability of some approaches to manage some information (e.g., *Random Forests*, one of the most performing approaches, is not able to manage types of data that involve a large number of categories).

The idea around which this paper was born is the introduction of a new model of evaluation based on the spectral pattern of the transactions, a novel representation of the information obtained by exploiting the *Fourier transformation* (Duhamel and Vetterli, 1990). This operation offers us a new point of view on data, providing the following advantages: consid-

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

Saia, R. and Carta, S.
Evaluating Credit Card Transactions in the Frequency Domain for a Proactive Fraud Detection Approach.
DOI: 10.5220/0006425803350342
In *Proceedings of the 14th International Joint Conference on e-Business and Telecommunications (ICETE 2017) - Volume 4: SECRYPT*, pages 335-342
ISBN: 978-989-758-259-2

ering that our process involves only the previous legitimate transactions, it allows us to operate proactively; such proactivity effectively face the *cold-start* issue (i.e., scarcity or absence of fraudulent cases during the model training); the representation in the frequency domain reduces the issues related to the data heterogeneity, since the spectral model is less influenced by the data variations.

The main contributions of this paper are summarized below:

- definition of the *time series* to use in a *Fourier Transform* process, in terms of sequence of values assumed by the transaction features;
- formalization of the comparison process between the *spectral pattern* of an unevaluated transaction and those of the previous legitimate transactions;
- definition of an algorithm able to to classify a new transaction as *legitimate* or *fraudulent*, on the basis of the previous comparison process.

The rest of the paper is organized as follows: Section 2 presents the background and related work of the scenario taken into account; Section 3 provides a formal notation, makes some assumptions, and defines the faced problem; Section 4 describes the proposed approach; Section 5 provides details on the experimental environment, on the used datasets and metrics, as well as on the adopted strategy and selected competitor, presenting the experimental results; some concluding remarks and future work are given in the last Section 6.

2 BACKGROUND AND RELATED WORK

The main task of a fraud detection system is the evaluation of a new financial transaction with the aim to classify it as *legitimate* or *fraudulent*, by using the information gathered in the past (i.e. value of the features that compose each transaction and if it was a fraud or not). This section provides a general overview of the context taken into account in this paper, starting with the introduction of the most used strategies and approaches, continuing with the description of the outstanding problems, and concluding with the presentation of core concepts on which the proposed approach is based, giving some details about the state-of-the-art approach chosen to evaluate its performance.

2.1 Strategies and Approaches

Operative Strategies. The fraud detection approaches operate by following a *supervised* or *unsupervised* strategy (Phua et al., 2010).

A *supervised* strategy works by exploiting the previous *fraudulent* and *non-fraudulent* transactions gathered by the system, using them in order to define a model able to classify the new transactions in a specific class (i.e., *legitimate* or *fraudulent*). It is evident how such strategy needs a series of examples concerning both classes, and how its effectiveness is limited to the recognition of known patterns.

An *unsupervised* strategy operates instead by analyzing the new transactions in order to evaluate when they present anomalies in their values, compared to the typical range of values that characterizes the context taken into account. It is an inefficient strategy because a fraudster can operate in order to avoid that the transaction presents anomalies in its values, and for this reason the definition of effective *unsupervised* strategies represents a hard challenge.

Operative Approaches. The most common way to operate in order to detect fraudulent events in a financial data stream related to a credit card activity is the adoption of a *static approach* (Pozzolo et al., 2014). It operates by dividing the data stream into blocks of equal size, training its model by using only a limited number of initial and contiguous blocks. A different modality is instead adopted by the *updating approach* (Wang et al., 2003), where at each new block, the model is updated by training it by using a certain number of latest and contiguous blocks. The *forgetting approach* (Gao et al., 2007) represents another possible operative way. By following this approach, the user model is updated when a new block appears, by using only the legitimate transactions present in the last two blocks, but by using all the previous fraudulent transactions. The models generated by these approaches can be directly exploited in order to evaluate the future blocks, or they can be used to define a bigger model of evaluation.

The main problems related to the aforementioned approaches are the following: the *static approach* is not able to model the users behavior; the *updating approach* is not able to operate with small classes of data; the *forgetting approach* presents a high computational complexity. In addition, these approaches have to face the common issues described in the next Section 2.2.

2.2 Open Problems

Data Scarcity Issue. The task of the researchers working in this area is complicated by the scarcity of public real-world datasets. This mainly happens due to the restrictive policies adopted by those working in this field, which do not allow them to release information about their business activities for privacy, competition, or legal issues. Not even a release in anonymous form of the data is usually taken into account by many financial operators, since also in anonymous form such data can provide precious information about their customers, and thus they could reveal potential vulnerabilities of the related e-commerce infrastructure.

Non-adaptability Issue. This problem concerns the inability of the fraud detection models to correctly classify the new transactions, when their features give rise to different patterns (wrt the patterns used to define the evaluation model). Both the *supervised* and *unsupervised* fraud detection approaches are affected by this problem (Sorournejad et al., 2016), which leads toward misclassifications, due to their inability to detect new legitimate or fraudulent patterns.

Data Heterogeneity Issue. Pattern recognition represents a very important branch of the machine learning, since it can be used to solve a large number of real-world problems. The effectiveness of these processes is nevertheless jeopardized by the heterogeneity of the involved data. Such problem happens due to incompatibility between similar features resulting in the same data being represented differently in different datasets (Chatterjee and Segev, 1991).

Data Unbalance Issue. Another important issue that the approaches of fraud detection have to face is the unbalanced distribution of data during the training of their evaluation models. This means that the information available to train an evaluation model are typically composed by a large number of legitimate cases and a small number of fraudulent ones, a data configuration that reduces the effectiveness of the classification approaches (Japkowicz and Stephen, 2002; Brown and Mues, 2012). A common strategy adopted in order to face this problem is the artificial balance of data (Vinciotti and Hand, 2003), made by performing an *over-sampling* or *under-sampling* process: in the first case the balance is obtained by duplicating some of the transactions that are less in number (usually, the fraudulent ones), while in the second case it is obtained by removing some of the transactions that are in greater number (usually, the legitimate ones).

Cold-start Issue. The cold-start problem arises when the set of data used to train an evaluation model does not contain enough information about the domain taken into account, making it impossible to define a reliable model (Donmez et al., 2007). In other words, it happens when the training data are not representative of all the involved classes of information (Attenberg and Provost, 2010) (i.e., in our case, legitimate and fraudulent). The approach presented in this paper faces this problem by training its evaluation model by using only a class of data (i.e., the legitimate cases). It represents a side effect of the adopted proactive approach that allows a system to operate without the need to have previous fraudulent transactions as examples, with all the advantages that derive from it.

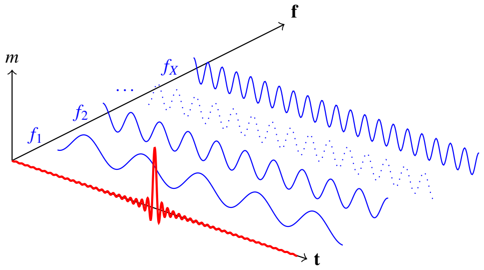

Figure 1: Time and Frequency Domains.

2.3 Proposed Approach and Competitor

Frequency Spectrum Evaluation. The idea that composes the core of this work is to perform the evaluation process in the frequency domain, by defining the evaluation model in terms of frequency components (spectral pattern). Such operation is performed by considering the sequence of values assumed by the transaction features as a *time series*, moving its analysis from the canonical domain to the frequency one. In order to perform this operation we use the *Discrete Fourier Transform* (*DFT*), whose formalization is shown in Equation 1, where i is the imaginary unit.

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

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

The result is a set of sinusoidal functions, as shown in Figure 1, each of them related to a specific frequency component. We can return to the original time domain by using the *inverse Fourier transform* reported in Equation 2. In the context of the presented approach we use the *Fast Fourier Transform* (*FFT*) algorithm in order to perform the Fourier transformations, since it allows us to rapidly compute the *DFT* by factorizing the input matrix into a product of sparse (mostly zero) factors. This is a largely used algorithm because it is able to reduce the computational complexity of the process from $O(n^2)$ to $O(n \log n)$, where n denotes the data size.

Random Forests Approach. The *Random Forests* (Breiman, 2001), approach represents one of the most common and powerful state-of-the-art techniques for data analysis, because in most of the cases it outperforms the other ones (Lessmann et al., 2015; Brown and Mues, 2012; Bhattacharyya et al., 2011). It consists in an ensemble learning method for classification and regression based on the construction of a number of randomized decision trees during the training phase. The conclusion are inferred by averaging the obtained results and this technique can be used to solve a wide range of prediction problems, with the advantage that it does not need any complex configuration.

3 PRELIMINARIES

This section provides the formal notation adopted in this paper and some basic assumptions, as well as the definition of the faced problem.

3.1 Formal Notation

Given a set of classified transactions $T = \{t_1, t_2, \ldots, t_N\}$, and a set of features $V = \{v_1, v_2, \ldots, v_M\}$ that compose each $t \in T$, we denote as $T_+ = \{t_1, t_2, \ldots, t_K\}$ the subset of legitimate transactions (then $T_+ \subseteq T$), and as $T_- = \{t_1, t_2, \ldots, t_J\}$ the subset of fraudulent ones (then $T_- \subseteq T$). We also denote as $\hat{T} = \{\hat{t}_1, \hat{t}_2, \ldots, \hat{t}_U\}$ a set of unclassified transactions. It should be observed that a transaction only can belong to one class $c \in C$, where $C = \{legitimate, fraudulent\}$. Finally, we denote as $F = \{f_1, f_2, \ldots, f_X\}$ the frequency components (spectrum) obtained as result of 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 τ 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 assumed by each distinct feature $v \in V$ that characterize the transactions in the set T_+ (i.e., the past legitimate transactions), and this sequence of values represents the *time series* taken into account in the *DFT* process. Their fundamental period τ starts with the value assumed by the feature in the oldest transaction of the set T_+ and it ends with the value assumed by the feature in the newest transaction, thus we have that $sr = |T_+|$; the sample interval *si* is instead given by the fundamental period τ divided by the number of acquisition, i.e., $si = \frac{\tau}{|T_+|}$. The frequency-domain representation, obtained by the *DFT*, process gives us information about the magnitude and phase of the signal at each frequency.

Denoting as x the output of the process, it represents a series of complex numbers, where x_r is the real part and x_i is the imaginary one (i.e., we have that $x = (x_r + ix_i)$). The magnitude can be calculated by using $|x| = \sqrt{(x_r^2 + x_i^2)}$ and that the phase can be calculated by using $\varphi(x) = \arctan\left(\frac{x_i}{x_r}\right)$. In our approach we take into account only the frequency magnitude.

3.3 Problem Definition

Denoting as Ξ the process of comparison between all the spectral patterns related to the *time series* extracted from the set T_+ and all the spectral patterns related to the *time series* extracted from the set of unevaluated transactions $\hat{T}$ (taken one at a time), our approach is aimed to classify as *legitimate* or *fraudulent* each transaction $\hat{t} \in \hat{T}$. Given a function $EVAL(\hat{t}, \Xi)$ able to verify the classification correctness of a transaction $\hat{t} \in \hat{T}$ made by using our approach, which returns a boolean value β (*0=misclassification*, *1=correct classification*), our objective can be formalized in terms of maximization of the results sum, as shown in Equation 3.

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

4 PROPOSED APPROACH

The implementation of the proposed approach was performed through the following three steps, which will be explained later:

1. **Data Definition**: definition of the *time series* in terms of sequence of values assumed by the transaction features;
2. **Data Processing**: conversion of the *time series* in the frequency spectrum by recurring to the *DFT* process;
3. **Data Evaluation**: formalization of the algorithm able to classify a new transaction as *legitimate* or *fraudulent* on the basis of a spectrum comparison process.

4.1 Data Definition

Formally, a *time series* is a series of data points stored by following the time order and, in most of the cases, it is a sequence of discrete-time data measured at successive equally spaced points in time. In the context of our approach, we considered as *time series* (*ts*) the sequence of values $v \in V$ assumed by the features of the transactions in T_+ and $\hat{T}$, as shown in Equation 4. 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.

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

$$ts(T_+) = (v_{1,1}, v_{2,1}, \ldots, v_{K,1}), (v_{1,2}, v_{2,2}, \ldots, v_{K,2}), \cdots, (v_{1,M}, v_{2,M}, \ldots, v_{K,M})$$
$$ts(\hat{T}) = (v_{1,1}, v_{2,1}, \ldots, v_{U,1}), (v_{1,2}, v_{2,2}, \ldots, v_{U,2}), \cdots, (v_{1,M}, v_{2,M}, \ldots, v_{U,M})$$

4.2 Data Processing

In this step, we move the *time series* of the transactions to the frequency domain by a *DFT* process performed through the *FFT* approach introduced in Section 5. In a preliminary study we compared the transaction representations in the time domain (*time series*) to those in the frequency domain (*frequency spectrum*). Without going deep into the merits of the formal characteristics of a Fourier transformation but by limiting our analysis at the context taken into account, in our approach we want to exploit the following two properties:

1. **Phase Invariance:** the first property demonstrates that there are not variations in the spectral pattern in case of value translation[2]. More formally, it is one of the *phase properties* of the Fourier transform, i.e., a shift of a *time series* in the time domain leaves the magnitude unchanged in the frequency domain (Smith et al., 1997). It means that the representation in the frequency domain allows us to detect a specific pattern, regardless the position of the values assumed by the transaction features that originate it;
2. **Amplitude Correlation:** the second property instead proves the existence of a direct correlation between the values assumed by the features in the time domain and the corresponding magnitudes assumed by the spectral components in the frequency domain. More formally, it is the *homogeneity property* of the Fourier transform (Smith et al., 1997), i.e., when the amplitude is altered in one domain, it is altered by the same entity in the other domain[3]. This ensures that the proposed approach is able to differentiate identical spectral patterns on the basis of the values assumed by their transaction features.

4.3 Data Evaluation

The evaluation of a new transaction $\hat{t} \in \hat{T}$ is performed by comparing its spectral pattern $F(\hat{t})$ (i.e., the series of values $f \in F$) to that of each previous legitimate transactions $t_+ \in T_+$. This is done by using the *cosine similarity* metric described in Section 5, as shown in Equation 5, where Δ represents the similarity value in terms of *cosine similarity*, α a threshold value experimentally defined in Section 5, and c is the resulted classification.

$$\Delta = cos(F(t), F(\hat{t})), \; with \; c = \begin{cases} \Delta \geq \alpha, & \text{legitimate} \\ \Delta < \alpha, & \text{fraudulent} \end{cases} \tag{5}$$

The final classification of a new transaction $\hat{t}$, which take into account all the comparisons (Equation 5) between the transaction $\hat{t}$ and all the transactions in T_+, is performed by using the Algorithm 1. It takes as input the set T_+ of past legitimate transactions, a transaction $\hat{t}$ to evaluate, and the threshold value α to use in the spectral pattern comparison process. It returns a boolean value that indicates the $\hat{t}$ classification (i.e., *true*=legitimate or *false*=fraudulent). It should be noted that, if it needs, the time complexity of the algorithm can be reduced by parallelizing the process over several machines (e.g., by exploiting large scale distributed computing models).

Algorithm 1: Transaction evaluation.

Input: T_+=Legitimate previous transactions, $\hat{t}$=Unevaluated transaction, α=Threshold value

Output: β=Classification of the transaction $\hat{t}$

```
procedure TRANSACTIONEVALUATION($T_+$, $\hat{t}$)
    $ts1 \leftarrow getTimeseries(\hat{t})$
    $sp1 \leftarrow getSpectrum(ts1)$
    for each $t_+$ in $T_+$ do
        $ts2 \leftarrow getTimeseries(t_+)$
        $sp2 \leftarrow getSpectrum(ts2)$
        $cos \leftarrow cos + getCosineSimilarity(sp1, sp2)$
    end for
    $avg \leftarrow \frac{cos}{|T_+|}$
    if $avg > \alpha$ then
        $\beta \leftarrow true$
    else
        $\beta \leftarrow false$
    end if
    return $\beta$
end procedure
```

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

[3] Scaling in one domain corresponds to scaling in the other domain

5 EXPERIMENTS

This section reports information about the experimental environment, the used datasets and metrics, the adopted strategy, and the results of the performed experiments.

5.1 Environment

The proposed approach was developed in Java, where we use the *JTransforms*[4] library to operate the Fourier transformations. The state-of-the-art approach (i.e., *Random Forests*) and the metrics to evaluate it have been implemented in *R*[5], by using *randomForest*, *DMwR*, and *ROCR* packages. The *RF* parameters have been tuned by searching those that maximize the performance. For reasons of reproducibility of the *RF* experiments, the *R* function *set.seed()* has been used in the code to fix the seed of the random number generator.

5.2 DataSet

The real-world dataset used for the evaluation of the proposed approach is related to a series of credit card transactions made by European cardholders[6]. In more detail, this dataset contains the transactions carried out in two days of September 2013, for a total of *492* frauds out of *284,807* transactions. It should be observed how it represents an highly unbalanced dataset (Pozzolo et al., 2015), considering that the fraudulent cases are only the *0.0017%* of all the transactions.

5.3 Metrics

Cosine Similarity. The *cosine similarity* (*Cosim*) between two non-zero vectors $\vec{v_1}$ and $\vec{v_2}$ is calculated in terms of cosine angle between them, as shown in the Equation (6). It allows us to evaluate the similarity between two spectral patterns by comparing the vectors given by the magnitude of their frequency components.

$$Cosim(\vec{v_1},\vec{v_2}) = cos(\vec{v_1},\vec{v_2}) = \frac{\vec{v_1}\cdot\vec{v_2}}{\|\vec{v_1}\|\cdot\|\vec{v_2}\|} \quad (6)$$

Accuracy. The *Accuracy* metric reports the number of transactions correctly classified, compared to the total number of them. Given a set of transactions $\hat{T}$ to be classified, it is calculated as shown in Equation 7, where $|\hat{T}|$ stands for the total number of transactions, and $|\hat{T}^{(+)}|$ for the number of those correctly classified.

$$Accuracy(\hat{T}) = \frac{|\hat{T}^{(+)}|}{|\hat{T}|} \quad (7)$$

Sensitivity. The *Sensitivity* metric measures the number of transactions correctly classified as *legitimate*, providing an important information, since it allows us to evaluate the predictive power of our approach in terms of capability to identify the legitimate transactions. Given a set of transactions $\hat{T}$ to be classified, the *Sensitivity* is calculated as shown in Equation 8, where $|\hat{T}^{(TP)}|$ stands for the number of transactions correctly classified as *legitimate* and $|\hat{T}^{(FN)}|$ for the number of *legitimate* transactions wrongly classified as *fraudulent*.

$$Sensitivity(\hat{T}) = \frac{|\hat{T}^{(TP)}|}{|\hat{T}^{(TP)}|+|\hat{T}^{(FN)}|} \quad (8)$$

F-score. The *F-score* is considered an effective performance measures for unbalanced datasets (Pozzolo et al., 2015). It represents the weighted average of the *Precision* and *Recall* metrics and it is a largely used metric in the statistical analysis of binary classification, returning a value in a range $[0,1]$, where 0 is the worst value and 1 the best one. Given two sets $T^{(P)}$ and $T^{(R)}$, where $T^{(P)}$ denotes the set of performed classifications of transactions, and $T^{(R)}$ the set that contains the actual classifications of them, this metric is defined as shown in Equation 9.

$$\begin{aligned} &F\text{-}score(T^{(P)},T^{(R)}) = 2\cdot\frac{Precision\cdot Recall}{Precision+Recal} \\ &\text{with} \\ &Precision(T^{(P)},T^{(R)}) = \frac{|T^{(R)}\cap T^{(P)}|}{|T^{(P)}|} \\ &Recall(T^{(P)},T^{(R)}) = \frac{|T^{(R)}\cap T^{(P)}|}{|T^{(R)}|} \end{aligned} \quad (9)$$

AUC. The *Area Under the Receiver Operating Characteristic* curve (*AUC*) is a performance measure used to evaluate the effectiveness of a classification model (Faraggi and Reiser, 2002). Its result is in a range $[0,1]$, where 1 indicates the best performance. Given the subset of previous legitimate transactions T_+ and the subset of previous fraudulent ones T_-, the formalization of the *AUC* metric is reported in the Equation 10, where Θ indicates all possible comparisons between the transactions of the two subsets T_+ and T_-. The result is obtained by averaging over these comparisons.

$$\Theta(t_+,t_-) = \begin{cases} 1, & if\, t_+ > t_- \\ 0.5, & if\, t_+ = t_- \\ 0, & if\, t_+ < t_- \end{cases} \qquad AUC = \frac{1}{|T_+|\cdot|T_-|}\sum_{1}^{|T_+|}\sum_{1}^{|T_-|}\Theta(t_+,t_-) \quad (10)$$

[4] https://sourceforge.net/projects/jtransforms/

[5] https://www.r-project.org/

[6] https://www.kaggle.com/dalpozz/creditcardfraud

5.4 Strategy

Cross-validation. In order to reduce the impact of data dependency, improving the reliability of the obtained results, all the experiments have been performed by using the *k-fold cross-validation* criterion, with *k=10*. Each dataset is divided in *k* subsets, and each *k* subset is used as test set, while the other *k-1* subsets are used as training set. The final result is given by the average of all results.

Threshold Tuning. Before starting the experiments we need to identify the optimal threshold α to use in the evaluation process, according to the Equation 5. In order to maintain a proactive approach, we perform this operation by using only the previous legitimate transactions, calculating the average value of the *cosine similarity* related to all pairs of different transactions $t_+ \in T_+$.Through this process we want to obtain the average value of *cosine similarity* measured between the frequency-domain representation (spectral pattern) of all pairs of previous legitimate transactions, so that we can use it to identify the potential fraudulent transactions. The results indicate *0.901* as optimal threshold (the recalculation frequency of that value depends on the context in which we operate).

5.5 Competitor

We compare our approach to *Random Forests* one. It is implemented in *R* language, by using the *randomForest* and *DMwR* packages. The *DMwR* package allows *Random Forests* to manage the class imbalance problem through the *Synthetic Minority Oversampling Technique* (*SMOTE*) (Bowyer et al., 2011). It represents a very popular sampling technique able to create new synthetic data by randomly interpolating pairs of nearest neighbors. The combined use of *Random Forests* and *SMOTE* allows us to verify the performance of our approach compared to one of the best solutions for fraud detection at the state of the art.

5.6 Results

The first set of experiments, the results of which are shown in Figure 2, was aimed to evaluate the capability of the spectral pattern to model a class of transactions (in our case, the legitimate one), compared to that of a canonic approach. We compared a million of transaction pairs (t'_+, t''_+) (with $t'_+, t''_+ \in T_+$ and $t'_+ \neq t''_+$), measuring the *cosine similarity* of the relative *time series* and spectral patterns. We observe how, since the beginning, the spectral pattern is able to effectively represent the class of legitimate transactions, since the variation of the *cosine similarity* is much smaller than that measured by comparing the same *time series*, providing a more stable modelization of the transaction class.

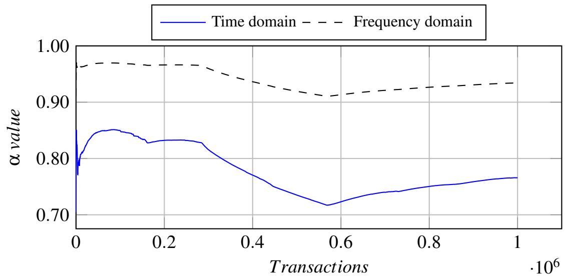

Figure 2: Model Evaluation.

The second set of experiments was instead focused on the evaluation of the proposed approach in terms of *Accuracy*, *Sensitivity*, and *F-score*. The results show in Figure 3 indicate that the *FD* performance are similar to those of *Random Forests*, although it does not use any previous fraudulent transaction to train its model, adopting a pure proactive strategy. It means that our approach is able to operate without training its model with both classes of transactions (legitimate and fraudulent).

The last set of experiments was aimed to evaluate the performance of the *FD* approach in terms of *AUC*. This metric is aimed to evaluate the predictive power of a classification model and the results in Figure 3 show how our model achieves performance close to those of *RF*, also considering that we do not exploit previous fraudulent transactions during the model training.

Summarizing, we can observe how the spectral representation of the transactions faces the *non-adaptability* and *heterogeneity* issues, thanks to its stability. We can also observe how the proactive strategy followed by our approach is able to reduce/overcome the *data imbalance* and *cold-start* issues, since only a class of transactions is used. Such proactivity allows a real-world fraud detection system to operate even in the absence of previous cases of fraudulent cases, with all the obvious advantages that derive from it.

6 CONCLUSIONS AND FUTURE WORK

Credit card fraud detection systems play a crucial role in our e-commerce age, where an increasing number of transactions takes place through this powerful instrument of payment, with all the risks that it involves. More than wanting to replace the existing state-of-the-art solutions, the approach presented in this paper

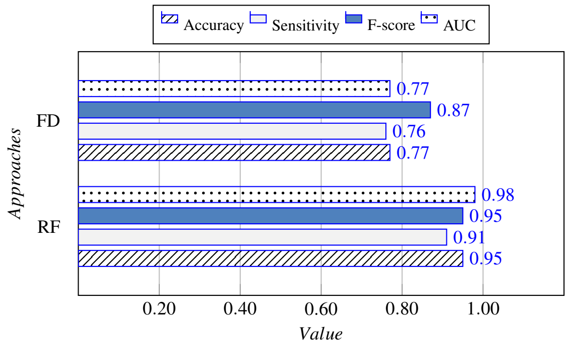

Figure 3: Performance.

wants to introduce a novel frequency-domain-based model that allows a fraud detection system to operate proactively. The results obtained are interesting, since it is necessary to consider that the state-of-the-art competitor taken into account (i.e., *Random Forests*), in addition to using both classes of transactions to train its model also preprocesses the dataset by using a balancing technique (i.e., *SMOTE*). It should be noted that the credit card context taken into account is only one of the possible scenarios, since the proposed approach can be used in any context characterized by financial electronic transactions.

A possible future work could take into account the definition of an hybrid approach of fraud detection that combines the characteristics of the canonical non-proactive approaches with those of our proactive approach.

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

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

Bhattacharyya, S., Jha, S., Tharakunnel, K. K., and Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3):602–613.

Bowyer, K. W., Chawla, N. V., Hall, L. O., and Kegelmeyer, W. P. (2011). SMOTE: synthetic minority oversampling technique. *CoRR*, abs/1106.1813.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

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.

Faraggi, D. and Reiser, B. (2002). Estimation of the area under the roc curve. *Statistics in medicine*, 21(20):3093–3106.

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.

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

Lessmann, S., Baesens, B., Seow, H., and Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, 247(1):124–136.

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.

Pozzolo, A. D., Caelen, O., Johnson, R. A., and Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. In *IEEE Symposium Series on Computational Intelligence, SSCI 2015, Cape Town, South Africa, December 7-10, 2015*, pages 159–166. IEEE.

Smith, S. W. et al. (1997). The scientist and engineer's guide to digital signal processing.

Sorournejad, S., Zojaji, Z., Atani, R. E., and Monadjemi, A. H. (2016). A survey of credit card fraud detection techniques: Data and technique oriented perspective. *CoRR*, abs/1611.06439.

Vinciotti, V. and Hand, D. J. (2003). Scorecard construction with unbalanced class sizes. *Journal of Iranian Statistical Society*, 2(2):189–205.

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.