Multiple Behavioral Models

A Divide and Conquer Strategy to Fraud Detection in Financial Data Streams

Roberto Saia, Ludovico Boratto, and Salvatore Carta
Dipartimento di Matematica e Informatica, Università di Cagliari, Italy
{roberto.saia,ludovico.boratto,salvatore}@unica.it

Abstract

Any business that operates on the Internet and accepts payments through debit or credit cards, also implicitly accepts that some transaction may be fraudulent. The design of effective strategies to face this problem is challenging, due to factors such as the heterogeneity and the non-stationary distribution of the data, as well as the presence of an imbalanced class distribution, and the scarcity of public datasets. Differently from most of the state-of-the-art strategies, instead of using a single model based on the past transactions of the users, our approach uses a set of models (behavioral patterns) to evaluate a new transaction, built taking into account the behavior of the user in different temporal frames of her/his history. By using only the legitimate past transactions of the users, we can operate in a proactive manner, by detecting the fraudulent ones that have never occurred, also overcoming the data imbalance issue. We evaluate our strategy by comparing it with one of the most performing approaches at the state of the art (i.e., Random Forests), using a real-world credit card dataset.

Introduction

Fraud is one of the major issues related with the use of debit and credit cards, which are becoming the most popular way to conclude every financial transaction. The research of efficient ways to face this problem has become an increasingly crucial imperative in order to eliminate, or at least minimize, the related economic losses.

Open Problems

Considering that the number of fraudulent transactions is typically much smaller than that of the legitimate ones, the data distribution is highly *unbalanced*, reducing the effectiveness of many learning strategies. A problem also complicated by the *heterogeneity*, *scarcity*, and *non-stationary* distribution of the information that characterizes a typical financial transaction record, which generates an *overlapping* of the classes of user expense.

Almost all of the state-of-the-art strategies are based on the detection of the suspicious changes in the user behavior, detected by using a *single behavioral model* defined on the basis of the past transactions of the users, a trivial approach that in many cases generates false alarm.

Our Solution

Our strategy extends the canonical criteria, by integrating the ability to operate with heterogeneous data (i.e., numeric and non-numeric), as well as by adopting multiple behavioral models. Such approach reduces the problems related with the *imbalanced* class distribution of the data, as well as those related with their *heterogeneity* and *scarcity*.

It happens because we take into account all the parts of a transaction, thus extracting more information, contrasting the *overlapping* of the classes of expense. By generating *multiple behavioral models* of the users, made by dividing the sequence of transactions in several event-blocks, we also face the problem of the *non-stationary* distribution of data, effectively modeling anyway their behavior.

Differently from most of the canonical machine learning approaches at the state of the art (e.g., Random Forests), our models do not need to be trained with the fraudulent transactions, overcoming the problem of *data imbalance* that afflicts the machine learning approaches.

The level of reliability of a new transaction is evaluated by comparing its behavioral pattern to each of the behavioral patterns of the user.

Main Contributions

This work provides the following main contributions to the current state of the art:

- introduction of the **Absolute Variation Calculation**, a process able to convert the past transactions of a user into a set of absolute numeric variations, using a specific criterion for each type of data;
- definition of the **Transaction Determinant Field** (TDF) set, a series of distinct terms, extracted from a transaction field, able to give more weight to certain data;
- introduction of the **Event-block Shift Vector** ($EBSV$) operations, which store, in the behavioral patterns of a user, the average values of the variations measured in a defined event-block;
- definition of a **Discretization Process**, which applied to the $EBSV$ and TDF outputs allow us to adjust the sensitivity of the fraud detection system, generating as output a set of behavioral patterns;
- formalization of the **Transaction Evaluation** process, used to evaluate the reliability of a new transaction, by comparing its behavioral pattern with the set of patterns previously obtained.

Our Approach

The following block diagram introduces a high-level architecture of our approach, which takes as input the past transactions of the users, and gives as output a series of behavioral patterns that characterize their transaction history.

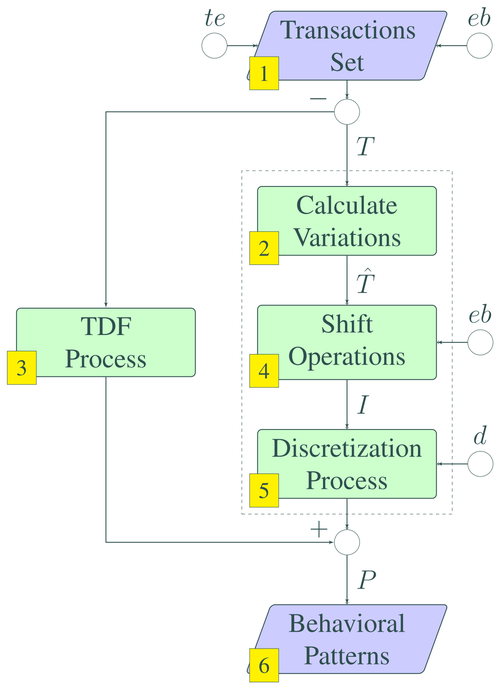

Block 1 - Transactions Set

It contains the past transactions of a user. The set **T** in output depends on the presence of a new transaction to evaluate in the input channel te: in absence of it, we have as output all transactions; otherwise we have only the $eb-1$ transactions (eb denotes the event-block size), followed by the te transaction to evaluate, because in this case we need to define only a single behavioral pattern.

Blocks 2 and 3 - Calculate Variations and TDF Process

As input it takes the set **T** without the fields $f \in F = \{f_1, f_2, \ldots, f_X\}$ chosen as *Transaction Determinant Field* (TDF), i.e., the parts of a transaction to which we decided to give more relevance during the fraud detection process, in accord with the operations of the block *TDF Process*. As output we have a new set **T** that reports the absolute numeric variations measured between each pair of contiguous transactions.

$$\hat{T} = |f_x^{(t_n)} - f_x^{(t_{n-1})}|$$

Block 4 - Shift Operations

The absolute variations in the set $\hat{\mathbf{T}}$ are here processed, in accord with the value in the input channel eb, that defines the size of the *Event-block Shift Vector* ($EBSV$). The result is a set **I** of behavioral patterns.

$$\hat{T} = [v_1, v_2, v_3, v_4, v_5] \Rightarrow c_1 = \frac{v_1+v_2+v_3}{|eb|}, c_2 = \frac{v_2+v_3+v_4}{|eb|}, c_3 = \frac{v_3+v_4+v_5}{|eb|}$$
$$\Downarrow$$
$$I = [c_1, c_2, c_3]$$

Block 5 - Discretization Process

This block converts the continuous values v_c present in the set **I** in discrete values v_d, according with the value of discretization in the input channel d.

$$v_d = \left\lceil \frac{v_c}{\left(\frac{max(f) - min(f)}{d}\right)} \right\rceil$$

Block 6 - Behavioral Patterns

The set **P** represents the output of the entire process, it is built by integrating the output of the *Discretization Process* with the output of the *TDF Process*. The level of reliability **R** of a new transaction (applied to the input channel te) is evaluated by comparing, through the *cosine similarity* ($cos(\Theta)$), its behavioral pattern $\hat{p}$ with all behavioral patterns in the set **P**.

$$R = sim(\hat{p}, P) = \frac{min(\cos(\Theta)) + max(\cos(\Theta))}{2}$$

Experimental Results

As we can observe in the following figure, the performance of our $EBSV$ approach reach those of the RF one, and this without training its models with the past fraudulent transactions (as occurs in RF). This result shows an important aspect, i.e., that $EBSV$ is able to operate in a proactive manner, by detecting fraudulent transactions that have never occurred in the past.

Conclusions and Future Work

In this paper we proposed a novel approach able to eliminate (or reduce) the threats connected with the frauds in the electronic financial transactions. Differently from the canonical state-of-the-art strategies, instead of exploiting a single model based on the past transactions of the users, we adopt multiple models, which take into account the behavior of the users in different temporal frames of her/his history.

These behavioral models do not consider the past fraudulent transactions, allowing us to operate in a proactive way, by detecting fraudulent transactions that have never occurred, with the advantage to solve the problem of data imbalance that afflicts the machine learning approaches.

A new transaction is evaluated by comparing its behavioral pattern to each of the behavioral patterns of the user, and the experiments show that the performance of the proposed approach are very close to those of the Random Forests, and this without training our models with the past fraudulent transactions.

A possible follow up of this work could be its development and evaluation in scenarios with different kind of financial transaction data, e.g., those generated in an E-commerce environment.

Acknowledgements

This work is partially funded by Regione Sardegna under project SocialGlue, through PIA - Pacchetti Integrati di Agevolazione Industria Artigianato e Servizi (annualit 2010), and by MIUR PRIN 2010-11 under project Security Horizons.