COMPUTERS & SECURITY 53 (2015) 175–186

Available online at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/cose

Computers & Security

BankSealer: A decision support system for online banking fraud analysis and investigation

Michele Carminati [a,*], Roberto Caron [a], Federico Maggi [a], Ilenia Epifani [b], Stefano Zanero [a]

[a] Politecnico di Milano, Dipartimento di Elettronica, Informazione e Bioingegneria, Italy
[b] Politecnico di Milano, Dipartimento di Matematica, Italy

ARTICLE INFO

Article history:
Received 20 December 2014
Received in revised form
18 March 2015
Accepted 2 April 2015
Available online 14 April 2015

Keywords:
Internet banking
Fraud detection
User profiling
Decision support system

ABSTRACT

The significant growth of online banking frauds, fueled by the underground economy of malware, raised the need for effective fraud analysis systems. Unfortunately, almost all of the existing approaches adopt black box models and mechanisms that do not give any justifications to analysts. Also, the development of such methods is stifled by limited Internet banking data availability for the scientific community. In this paper we describe BankSealer, a decision support system for online banking fraud analysis and investigation. During a training phase, BankSealer builds easy-to-understand models for each customer's spending habits, based on past transactions. First, it quantifies the anomaly of each transaction with respect to the customer historical profile. Second, it finds global clusters of customers with similar spending habits. Third, it uses a temporal threshold system that measures the anomaly of the current spending pattern of each customer, with respect to his or her past spending behavior. With this threefold profiling approach, it mitigates the under-training due to the lack of historical data for building well-trained profiles, and the evolution of users' spending habits over time. At runtime, BankSealer supports analysts by ranking new transactions that deviate from the learned profiles, with an output that has an easily understandable, immediate statistical meaning.

Our evaluation on real data, based on fraud scenarios built in collaboration with domain experts that replicate typical, real-world attacks (e.g., credential stealing, banking trojan activity, and frauds repeated over time), shows that our approach correctly ranks complex frauds. In particular, we measure the effectiveness, the computational resource requirements and the capabilities of BankSealer to mitigate the problem of users that performed a low number of transactions. Our system ranks frauds and anomalies with up to 98% detection rate and with a maximum daily computation time of 4 min. Given the good results, a leading Italian bank deployed a version of BankSealer in their environment to analyze frauds.

* *Corresponding author.*
E-mail addresses: michele.carminati@polimi.it (M. Carminati), roberto.caron@mail.polimi.it (R. Caron), federico.maggi@polimi.it (F. Maggi), ilenia.epifani@polimi.it (I. Epifani), stefano.zanero@polimi.it (S. Zanero).
http://dx.doi.org/10.1016/j.cose.2015.04.002
0167-4048/

1. Introduction

The popularity of Internet banking has led to an increase of frauds, perpetrated through cyber attacks, phishing scams and malware campaigns, resulting in substantial financial losses (Wei et al., 2013; Bolton and David). In 2013, Kaspersky Lab[1] detected 28.4 million attacks using financial malware, with a 27.6% increase over 2012. The number of users targeted in attacks involving financial malware also rose by 18.6% to 3.8 million. A similar trend characterizes online banking frauds which increased 30% in 2012–2013.[2]

Internet banking frauds are difficult to analyze and detect because the fraudulent behavior is dynamic, spread across different customer profiles, and dispersed in large and highly imbalanced datasets (e.g., web logs, transaction logs, spending profiles). Despite the importance of the problem, the development of new online banking fraud decision support systems is made difficult by the limited availability of transactions and fraud datasets, due to privacy concerns. As a consequence, only a limited amount of research deals with fraud detection in online banking. Commercial systems do exist, but they offer limited insight in their inner workings due to obvious intellectual property concerns. We noticed that most existing approaches build black box models that are not very insightful for analysts in the subsequent manual investigations, making the process less efficient. In addition, systems based on baseline profiling are not adaptive, and do not take into account cultural and behavioral differences that vary from country to country. Instead of focusing on *pure detection* approaches, we believe that more research efforts are needed toward systems that *support* investigations. Cooperating with a leading security company which helps banks build fraud detection systems and processes, we had the unique opportunity to work on a real-world, anonymized dataset of Internet banking transactions.

In this paper we present a detailed description of BankSealer (Carminati et al., 2014), a decision support system for online banking fraud analysis and investigation that automatically ranks frauds and anomalies in transactions. Most of the development was driven by the analysis of the dataset itself. BankSealer uses a combination of advanced data mining, statistical, and mathematical techniques to automatically rank transactions on the basis of the risk of being fraudulent. During a training phase, it builds a local, global, and temporal profile for each user. The local profile models past user behavior to evaluate the anomaly of new transactions by means of a novel algorithm that uses the Histogram Based Outlier Score (HBOS). The global profiling clusters users according to their transactions features via an iterative version of Density-Based Spatial Clustering of Applications with Noise (DBSCAN), and compute the anomaly with the Cluster-Based Local Outlier Factor (CBLOF). The temporal profile aims to model transactions in terms of time-dependent attributes. For this, we design a series of thresholds and measure the anomaly in terms of the percentage gap from the thresholds once they are exceeded. We handle the concept drift of the scores with an exponential decay function that assigns lower weights to older profiles.

We tested the BankSealer on real-world data, injecting a realistic set of attacks (e.g., credential stealing, banking trojan activity, and frauds repeated over time) built in collaboration with domain experts. Our system ranked fraud and anomalies with up to 98% detection rate.

In summary, our main contributions are:

1. An in-depth analysis of a real-world online banking dataset, in which we highlight the aforementioned challenges and the importance of dealing with dataset scarcity in this research field.
2. A general framework for online semi-supervised outlier-detection based on a combination of different models to discover different types of frauds. Our approach has a score with a clear statistical meaning, is adaptive to non-stationary sources and can deal with concept drift and data scarcity.
3. An almost exhaustive evaluation through a set of realistic attacks and in a real-world setting, thanks to the deployment to a large national bank.

2. Online banking fraud detection: goals and challenges

Our goal is to support the analysis of (novel) frauds and anomalies. Hence, we do not want to focus on a classifier but provide the analysts with a ranked list of transactions, along with the risk score. The rationale behind this design decision is that analysts must investigate reported alerts in any case: therefore, the focus is on collecting and correctly ranking evidence that support the analysis of fraudulent behavior, rather than just flagging transactions.

From a literature review (described in Section 6) and a real-world dataset obtained from a large national bank (described in Section 3), we found peculiar characteristics that make the analysis of this data particularly challenging. First and foremost, the distribution of attributes values is imbalanced and skewed (non-symmetric), which makes it difficult to approximate with most common statistical distributions, and unusable with most statistical methods to explain or predict trends and outliers. A second troublesome characteristic is the prevalence of users who perform a low number of transactions – an issue not considered in previous literature. Finally, the system must adopt a simple design and must be able to handle the high load of transactions avoiding high computational and spatial complexity.

Given the scarcity of labeled datasets, such a system must be able to work in an unsupervised or semi-supervised fashion (we can assume that no fraud exists in this dataset, as indicated by our collaborators). This conflicts with the requirement of the system being able to provide "readable" evidence to corroborate each alert. These peculiarities have remarkable implications for the typical statistical and data mining methods used in the outlier detection field.

[1] Kaspersky Lab – Financial cyber threats in 2013 – Available at http://goo.gl/8iaDcU.

[2] Symantec – Internet security threat Report 2013 – Available at http://goo.gl/hDgafz.