Knowledge-Based Systems 71 (2014) 389–408

Contents lists available at ScienceDirect

Knowledge-Based Systems

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

SPAN: Finding collaborative frauds in online auctions

Sidney Tsang *, Yun Sing Koh, Gillian Dobbie, Shafiq Alam

Department of Computer Science, The University of Auckland, New Zealand

ARTICLE INFO

Article history:
Received 1 March 2014
Received in revised form 13 August 2014
Accepted 16 August 2014
Available online 23 August 2014

Keywords:
Online auction fraud
Fraud detection
Anomaly detection
Markov random fields
Belief propagation

ABSTRACT

Fraud is an ongoing concern for online auction websites. Current methods to detect or prevent fraud have been limited in several ways, making them difficult to apply in real world settings. Firstly, existing methods cannot adapt to changes in the behaviour of fraudulent users over time: new models must be continuously constructed as they gradually lose accuracy. In addition, each method can only be used to detect a specific type of fraud. Secondly, existing methods are generally poor at identifying collaborative frauds. And thirdly, method training and evaluation has been limited by the quality of available datasets. We propose an algorithm named SPAN (Score Propagation over an Auction Network), for detecting users committing collaborative fraud that addresses these problems. SPAN is a two phase method that first applies anomaly detection on multiple 2-dimensional feature subspaces to generate an initial anomaly score for each user, then applies belief propagation to revise those scores to identify suspicious groups of users. We report extensive experimental results using synthetic data which shows that SPAN performs well across three different types of fraud, and outperforms a previous approach for collaborative fraud detection called 2-Level Fraud Spotting. Experiments on a real dataset shows that SPAN behaves reasonably, and can identify groups of users that appear anomalous.

1. Introduction

Online auction sites such as eBay[1] and TradeMe[2] allow goods and services to be bought and sold online anonymously. The most common type of online auction is the English auction, where bids are placed in ascending order, publicly observable, and the winner is the final bidder with the highest bid [15]. In 2013, there were 128 million active users in eBay and auction volume of more than $22 billion USD [11].

The anonymity and simplicity of creating multiple aliases allow unsuspecting users to be exploited by dishonest users. This exploitation can take many forms, including shilling, non-delivery, misrepresentation, or the sale of stolen goods [9]. Dishonest users will also disguise themselves to avoid detection by imitating legitimate behaviours [7], making fraudulent behaviours difficult to define. Previous work has noted that users often appear to behave irrationally [16], and previous attempts at clustering users into predefined types according to their bidding behaviour have failed to label the majority of users [20]. The range of potential fraudulent behaviour together with the number and range of legitimate behaviours makes it difficult to differentiate between fraudulent and legitimate users. The class imbalance in auction data, where the number of legitimate actions outnumber the fraudulent, makes the accurate classification of users non-trivial [4].

There have been a range of methods proposed to detect fraudulent users and transactions in online auctions. However, to the best of our knowledge, all previously proposed methods have several limitations in common. First, each method is only able to identify one type of fraud. The same method will be less effective when used to detect other types of fraud, or even variations of the same type of fraud. In addition, strategies to commit fraud change as users using them are found and removed over time. As the strategies evolve, the detection method will gradually lose accuracy. Secondly, existing methods are not very good at detecting collaborative frauds since they do not make use of all available information. While some previous methods, such as that proposed by Lin et al. [14], make use of features derived by modelling the auction network as a graph, each user is still considered individually when determining whether they are fraudulent. The only exception is 2-Level Fraud Spotting (2LFS) by Chau et al. [8], which we discuss later. Thirdly, the quality of datasets used in previous work has been limited by dataset size or label accuracy. For real datasets, this is due to a lack of ground-truth, and often a limited dataset size. For synthetic datasets, this is due to the lack of

* Corresponding author.
E-mail addresses: stsa027@aucklanduni.ac.nz (S. Tsang), ykoh@cs.auckland.ac.nz (Y. Koh), gill@cs.auckland.ac.nz (G. Dobbie), sala038@aucklanduni.ac.nz (S. Alam).

[1] http://www.ebay.com/.
[2] http://www.trademe.co.nz/.

http://dx.doi.org/10.1016/j.knosys.2014.08.016
0950-7051/

evidence verifying that they emulate real data. As a result, models trained or evaluated using those datasets may not perform well in a real world setting.

In this work, we propose a novel approach which avoids the limitations listed above. Our approach consists of an anomaly detection phase and a belief propagation phase. Firstly, by using an anomaly detection method, our approach identifies users who behave sufficiently differently from the majority of users as suspicious (potentially fraudulent), and can detect users with different fraud strategies, including previously unknown strategies. The assumption is that the majority of users is legitimate. Secondly, the belief propagation phase allows groups of suspicious users to be found during the belief propagation phase. This is based on the assumption that users who interact with many other suspicious users are also likely to be suspicious. Thirdly, by using a validated synthetic data generator, we avoid the problems associated with unknown ground-truth in real data, and the problem of limited generalisability when using synthetic data. However, we also evaluate our method using real data to ensure that our approach does in fact, identify anomalous groups of users.

The main contributions of our paper are as follows:

- We propose a novel approach for detecting collaborative fraud in online auctions. The approach, which we name SPAN (Score Propagation over an Auction Network), contains two phases. In Phase 1, the anomaly scoring phase, the anomaly score of each user is calculated using a set of features describing that user. Specific to our approach is that outlier detection is performed, not in the whole feature space, but in carefully selected two-dimensional subspaces; this has several advantages, as described in Section 4.1. In Phase 2, the score propagation phase, the anomaly score for each user from Phase 1 is revised depending on their interactions with other users. This additional phase improves the overall accuracy of SPAN and allows groups of collaborating fraudulent users to be identified.
- We improve our previous auction simulation to generate data that accurately models the network features of legitimate users.
- We implement three types of collusive frauds described in previous literature. Combined with the auction simulation, we create multiple sets of synthetic data containing each fraud type, which is used to evaluate our proposed algorithm.

The paper is structured as follows: Section 2 describes related work in auction fraud detection, and briefly, in outlier detection. Section 3 gives background information important for understanding the paper. Section 4 describes the two phases of our proposed approach, and its time complexity. Section 5 defines the fraud types contained in the generated synthetic data used in evaluation. Section 6 presents evaluation results for SPAN and 2LFS under different conditions. Section 7 presents a case study on a real dataset from an online auction site. Sections 8 and 9 conclude the paper and present future work.

2. Related work

The following section describes related work. Section 2.1 describes the previous work in online auction fraud. Section 2.2 describes an anomaly detection method in graphs, called Oddball, on which the SPAN algorithm is partially based.

2.1. Auction fraud detection

The approach used in previous work to create fraud detection methods in online auctions generally follows three main steps: (1) define the behaviour of interest, (2) identify features that differentiate between legitimate and fraudulent users or auctions, and (3) develop a fraud detection algorithm based on the selected set of features. Below, we describe previous work in terms of these three steps.

2.1.1. Fraudulent behaviours

We concentrate on three types of fraud most commonly investigated in previous literature: shilling fraud, reputation manipulation, and non-delivery fraud. Other types of online auction fraud have been described in detail by Dong et al. [9].

Shilling fraud occurs when a user submits bids to a collaborating sellers' auction, and has been investigated by Kauffman and Wood [13], Trevathan and Read [21], Xu et al. [24], Tsang et al. [23]. There are three types of shilling fraud: competitive shilling, buy-back shilling and reserve-price shilling, each with a different purpose [9]. In competitive shilling, the goal is to maximise the amount a legitimate auction winner will pay by repeatedly outbidding legitimate users while avoiding winning accidentally. In buy-back shilling, the goal is to prevent an item from being sold to a legitimate user below value. In reserve-price shilling, the goal is to reduce the total amount of auction fees paid.

Reputation manipulation occurs when a user attempts to increase their positive reputation score to appear trustworthy. This can be achieved, for example, by using multiple accounts to create and complete fraudulent auctions, then posting positive feedback from each account in the transaction. Another method is to legitimately buy or sell very low value items, and gain positive feedback from legitimate users. The benefits of a positive reputation for sellers has been shown by Resnick and Zeckhauser [18]. This type of fraud has been investigated by Chau et al. [8], Gregg and Scott [10], You et al. [26], and Lin et al. [14].

Non-delivery fraud occurs when an auction is successfully completed, but after the winner sends payment, the item is never delivered. Non-delivery fraud is often committed after reputation manipulation when the account appears trustworthy. This type of fraud has been investigated by Chang and Chang [7], and Almendra [2].

2.1.2. Feature selection

Features used to differentiate between fraudulent and legitimate users can be divided into user level and network level features. User level features describe the behaviour of individual users, such as the number of auctions they participate in, or the average value and frequency of their bids. Chang and Chang [7] lists a comprehensive set of user-level features along with a brief description for each. Network level features describe the relationships between users. Network features will be discussed in greater detail in Section 4.

The vast majority of previous work uses only user-level features [13,21,9,24,23,10,26,7,2], and do not make use of network-level features to detect fraudulent users or auctions. Since users are considered individually, those proposed methods are not very effective for identifying users committing fraud collaboratively. The only exception is the work by Chau et al. [8], and the slight extension by Pandit et al. [17], where the network is used to identify groups of fraudulent users, and to a limited extent by Lin et al. [14], which uses the network to derive a feature as part of the set of inputs to a neural network.

2.1.3. Proposed methods

The methods that have been proposed to reduce fraud can be divided into broad categories of detection and prediction. Methods in fraud prediction include work by Kauffman and Wood [13], Xu et al. [24], Chang and Chang [7], Almendra [2]. Kauffman et al. constructed a probit model to predict the auctions in which