

# Large-Scale Shill Bidder Detection in E-commerce

Michael Fire<sup>1</sup>, Rami Puzis<sup>1</sup>, Dima Kagan<sup>1</sup>, and Yuval Elovici<sup>1</sup>

Ben-Gurion University of the Negev  
{mickyfi,puzis,kagandi,elovici}@post.bgu.ac.il

**Abstract.** User feedback is one of the most effective methods to build and maintain trust in electronic commerce platforms. Unfortunately, dishonest sellers often bend over backwards to manipulate users' feedback or place phony bids in order to increase their own sales and harm competitors. The black market of user feedback, supported by a plethora of shill bidders, prospers on top of legitimate electronic commerce. In this paper we investigate the ecosystem of shill bidders based on large-scale data by analyzing hundreds of millions of users who performed billions of transactions, and we propose a machine-learning-based method for identifying communities of users that methodically provide dishonest feedback. Our results show that (1) shill bidders can be identified with high precision based on their transaction and feedback statistics; and (2) in contrast to legitimate buyers and sellers, shill bidders form cliques to support each other.

**Keywords:** Big Data · Cyber Security & Privacy · Fraud Detection · Data Science · Social Network Analysis.

## 1 Introduction

Electronic commerce (e-commerce) usage has increased sharply as e-commerce platforms have become interwoven into people's everyday lives as places to buy and sell products. The e-commerce worldwide sales to consumers is expected to pass the four trillion dollars mark in 2020, almost quadrupling itself compared to 2014 [12,11]. E-commerce platforms, such as Amazon,<sup>1</sup> Alibaba,<sup>2</sup> eBay,<sup>3</sup> Etsy,<sup>4</sup> and OnlineAuction,<sup>5</sup> already have hundreds of millions of active users [2,7,3,29]. In many e-commerce platforms, users can use the platforms to sell and buy various products from each other, from simple everyday products, such as a cup holder which costs only a few dollars, to more exotic products, such as a fighter jet which costs several million dollars [1]. Furthermore, more than 25 million

---

<sup>1</sup> <http://www.amazon.com/>

<sup>2</sup> <http://www.alibaba.com/>

<sup>3</sup> <http://www.ebay.com/>

<sup>4</sup> <http://www.etsy.com/>

<sup>5</sup> <http://www.onlineauction.com/>

individuals use e-commerce platforms as their primary or secondary source of income [15].

In many e-commerce platforms, after a user purchases a product, he or she can rate the product and write a review regarding the product’s features, such as the product’s quality [47]. Furthermore, many e-commerce platforms employ reputation systems, in which the buyer can leave feedback regarding the seller of the product [51]. The accumulated feedback on each seller, which in many cases is viewable to other users of the website, can assist other users in building trust towards the seller [51]. According to Lucking-Reile et al. [44] the feedback these sellers receive can have a measurable effect on their auction prices, where negative feedback has a much greater effect than positive feedback ratings. Moreover, in cases where sellers receive too much negative feedback from other users, the website operator can decide to revoke the seller’s selling privileges. For example, Facebook banned and reduced the number of ads for businesses who received too much negative feedback from buyers [10].

As in many other online platforms, such as search engines, online social networks, and online gaming platforms, the platform’s users can utilize dishonest techniques to manipulate the platform’s statistics in order to create profits [37]. Moreover, in many online platforms, the platform’s users can purchase this type of service from third-party providers, which in most cases violates the platforms terms of service. For example, the market of buying fake followers and fake retweets in Twitter is already a multimillion-dollar business [48]. The equivalent to Twitter’s fake followers in e-commerce are fake reviews. The fake review industry is flourishing; there are many paid reviewers and even people who receive free products in exchange for a positive review [4,13]. In e-commerce platforms dishonest users (referred to as shill bidders) can increase their product’s price in auctions by bidding on products with the intent to artificially increase the product’s price or desirability [28,54]. Additionally, shill bidders are also users who buy products in order to artificially improve a seller’s feedback or the product’s search standing [28]. Shill bidding is forbidden in many e-commerce platforms, such as eBay [28], tophatter [16], flippa [17], etc. Moreover, shill bidding in online auctions is illegal in some countries, such as the United States, and can be considered as wire fraud, a felony which can lead to a maximum penalty of up to four years in prison and a million dollars in fines [14].

In this paper we study the shill bidder ecosystem. The main contributions of this study are threefold. First, we offer generic algorithms for identifying shill bidders in e-commerce platforms. Moreover, we evaluate these algorithms on one the biggest e-commerce datasets in the world (referred to as the *e-commerce dataset*), which includes several billion buying transactions and several billion feedback interactions between the platform’s users. Second, we analyze the activities of over 187,244 identified shill bidders to better understand their characteristics (see Figure 1). Lastly, we investigate the ecosystem of shill bidders and offer methods on how this ecosystem’s properties can be utilized to better identify shill bidders. To the best of our knowledge, this study which aims to identify shill bidders is the largest of its kind to date.

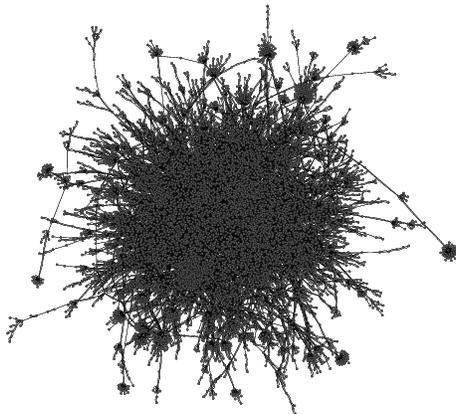


Fig. 1: Shill bidders feedback weighted directed graph (subgraph with 10,249 vertices and 21,859 links.)

The remainder of the paper is structured as follows: In Section 2 we give a brief overview of previous relevant studies related to e-commerce security and shill bidders. In this section, we also introduce several studies which used similar data mining algorithms as this study in other online platforms, such as Facebook and Twitter. Next, in Section 3 we present the methods and algorithms we developed for identifying shill bidders. Afterwards, in Section 4 we describe the e-commerce dataset. Additionally, in Section 4, we present the performed empirical evaluation process and the results of the shill bidder identification algorithms on the e-commerce dataset. Then, in Section 5, we present the performed empirical study of the shill bidders' characteristics. Furthermore, in this section we also present our analysis of the shill bidders' ecosystem. Lastly, in Section 6, we present our conclusions and also offer future research directions.

## 2 Related Work

### 2.1 E-commerce Platform Fraud

In the past decade, during the rise in the availability of the Internet, there was also a constant growth in popularity and volume of e-commerce [36]. However, the extreme popularity and the high anonymity of e-commerce attracted many fraudsters. According to Javelin's financial impact of fraud study [35], the 2017 online revenue loss estimation due to fraud stood at 40 billion USD.

In the past several years, many research efforts have been devoted to studying e-commerce platform's security and trust [57]; however, only a small portion of this work centered on the detection of shill bidders. Chakraborty et al. [22] defined shill bidding as an illicit participation of sellers in auctions designed to

increase the price at which a product sells. There are many types of shill bidding, and we are going to present two notable types as examples:

The first type is competitive shilling, which is the simplest kind of shill. Kauffman et al. [39] explains that it is used to make bidders pay more for a product. Basically the seller or his or her accomplices enter bids to make legitimate bidders pay more. The second type is reserve price shilling. Kauffman et al. [39] notes that this is a shilling that can be used to avoid paying auction house fees such as insertion fees or secret reserve fees. In addition, there are questionable techniques that stand on a thin line between tips or tricks and shill bidding. For example, Roth and Ockenfels [52] describe last-minute bidding (“sniping”), which is basically placing your bid exactly before the auction ends in order to pay a lower price and snatch the product. There is controversy about sniping and it is considered legal [43,53]. Moreover, Roth and Ockenfels [52] note that late bidding may be used by a dishonest seller who attempts to raise the price by using shill bidders.

The classical approach for tackling the shill-bidding problem is to use detection and prediction techniques. In 2005, Chau and Faloutsos [24] purposed a method for fraudster detection in online auctions. First, they created different combinations of user information-based features, for example, the number of items bought or sold in a certain time period. The features were extracted from a dataset of 43 fraudsters and 72 legitimate users from eBay. Next, they inserted the feature sets into a C5.0 decision-tree-based classification algorithm. In their best combination of features, they were able to detect malicious users with a precision of 82%, a True Positive (TP) rate of 83%, and a False Positive (FP) rate of 11%. In 2006, Chau et al. [25] presented an expansion of their first work, a novel 2-Level Fraud Spotting (2LFS) model. The 2LFS uses two types of features: user information features and network topography features. Afterwards, they used a Belief Propagation algorithm over a Markov Random Field to create the detection model. Finally, they evaluated their model on a manually labeled dataset of 55 users (graph with 55 nodes and 620 edges). In this work they demonstrated that social network information detection techniques can be applied in e-commerce fraud detection.

In 2007, Pandit [49] developed a model that represents users and transactions as a Markov Random Field. The model uses a Belief Propagation mechanism in order to detect fraudulent users. The detected fraudsters with a precision of 0.9 for the synthetic dataset; however, they were unable to present algorithm evaluation results on real-world datasets because they were unable to label them. In 2008, Beyene et al. [19] presented a model that represents eBay transactions and feedback in a graph. In their study they explained the difference between the eBay graph and a standard social network graph. First of all, they found that a rich club phenomenon [58] did not exist in the eBay feedback graph. Second, they showed that preferential attachment holds only partially. In addition, they discovered that when a user is sufficiently trustable, the exact number of positive reviews became less important. However, negative reviews significantly hurt the user’s credibility.

In 2011, Chang et al. [23] purposed a new two-stage phased model for early fraud detection in online auctions. The model’s first phase constructed behavior models based on users’ transaction histories. The first phase’s main was feature extraction in a way that significant behavioral differences between legitimate users and fraudsters could be determined. The second phase was fraud detection where the data of a suspicious account was inserted into the detection model in order to test if the seller is legitimate or not. They performed an Instance-Based Learning (IBL) algorithm. Moreover, the modeling here was based on phased modeling, which means that for each part of the data lifespan a behavior model is built, and also a separate detection model is built for each of the models. They were able to get an average recall rate of classifying accounts of over 93%.

In 2014, Tsang et al. [55] proposed a detection method based on supervised learning with generated synthetic data. In their study they proved that using synthetic data can have advantages over previous work using real data. They tested decision tree and neural network models with different kinds of features on synthetic and commercial datasets, where the largest dataset contained 58,162 users. They found the best results by using a decision-tree-based model with 0.999 and 0.977 detection precision for simple and complex shill cases, respectively.

In 2016 Majadi et al. [45] analyzed previous studies on shill bidders and found that the most common features used in past literatures were: first bidding, last bidding, bid increment, outbid time, bid frequency, affinity to the sellers, and winning ratio.

In 2017 Kaghazgaran et al. [38] analyzed fake-review properties. They found that fake reviewers give longer reviews and they tend write their reviews in bursts.

In 2018, Ganguly et al. [31] proposed an SVM-based method for shill bidder detection. They evaluated their method on a dataset that contained information on 149 auctions and 1024 bidders who bid on PDAs in eBay. In order to label the dataset, they used hierarchical clustering and manually labeled clusters that looked suspicious. Their classifier achieved AUC an of 0.86 using 10-fold cross validation.

## 2.2 Identifying Malicious Users in Online Platforms Using Supervised Learning

In addition to e-commerce, there are several other platforms that have risen during the internet era. These platforms have enormous numbers of users. For example, Facebook, the biggest online social network, has more than 2.23 billion monthly active users [6] and the user number is still growing. The similarity between these platforms is not only their size and popularity, but also the high resemblance in the nature of their malicious users. The e-commerce malicious users are very similar to spammers and fake profiles in social networks, as well as bots in online games. These malicious users have a goal that is different from benign users, and they perform non-standard actions with some hidden motives. A report [9] suggests that the current number of fake or clone accounts on Facebook is between 13% and 14%. Until recently, Twitter was considered a

controversial social network that was a safe haven for bots and fake profiles [18]. In the past year, Twitter started a full-scale campaign against fake users and bots and deleted about 6% of its users [5,8].

In 2012, Rahman et al. [50] developed the FRAppE application for socware identification. They used various application properties, such as the number of permissions required by the applications, as classifier features. Rahman et al. used Support Vector Machine (SVM) to build a classifier for socware detection. They detected socware with 99.5% accuracy and a low FP rate of 4.1%. In 2010, Wong [56] presented a method for spammer identification on Twitter. His method was based on the use of Twitter graph structure in order to build a classifier that was based on graph features. This method was proven to be fairly successful with 89% precision. Another interesting study was performed by Mitterhofer et al. [46]. They developed a method of bot detection in World of Warcraft. Their detection mechanism was based on the route that the character made in the game. They discovered that bots did repetitive actions like traveling the same route many more times than a benign user.

The general methodology of using supervised learning algorithms for classifying users in online platforms covers the following topics: feature extraction, ground truth classification, choice of a learning algorithm, choice of a training set, evaluation method, and performance metrics. When the data is imbalanced, e.g., malicious users are underrepresented in the data, a training set that reflects the real distribution of users may result in poorly trained classifiers. Tsang et al. [55] suggested that by changing the ratios of legitimate and malicious users we can improve the TP rate and the FP rate. They also mentioned that the training set size can outweigh the effects of the class imbalance. Chawla et al. [26] noted that measures like accuracy can be misleading, and they suggested using more accurate measures such as Receiver Operating Characteristic (ROC) curves. Another method that Chewla et al. suggested is to perform oversampling or undersampling in order to minimize the imbalance effect on the classification result. The main idea in oversampling/undersampling is to fit class distribution of the dataset to the class in the real world. In 2016 Hooi et al. [34] developed FRAUDAR, a system for identifying “camouflaged” malicious accounts. FRAUDAR uses density-based features in order to identify malicious vertices in bipartite networks.

In 2017 Kumar et al. [40] studied sockpuppetry in discussion communities. They discovered that sockpuppets behave differently from benign users; for instance, they have more clustered ego-networks and are more likely to interact with each other.

### 3 Identifying Shill Bidders Using Supervised Learning

In this study, we focus on developing generic classifiers that, can identify shill bidders in various e-commerce platforms. To cope with the challenge of identifying shill bidders in e-commerce platforms, we follow the regular methodology of using supervised learning algorithms for predicting the likelihood of a user to

be a skill bidder. User features employed in this study to identify skill bidders are common to many e-commerce platforms.

### 3.1 Feature Extraction

In order to construct classifiers for identifying skill bidders, we define features to be extracted from the transaction and feedback data of each e-commerce user. Next, we describe in detail each one of the features we extract for every e-commerce user.

Let be  $v$  an user in an e-commerce platform. For each user  $v$ , we can extract features based on  $v$ 's buying and selling transactions, such as  $v$ 's number of buying transactions, and based on  $v$ 's given and received feedback, such as the amount of feedback  $v$  gave. In this section, we describe in detail all the features we extracted and used during this study's experiments. We open this section with the formal definitions of the extracted features which were based on the users' buying and selling transactions. Afterwards, we introduce the formal definition of all the extracted features that were based on the users' feedback activities. Next, we present several personal user's features, such as the country the user's declared to live in. Lastly, we present the target class feature.

**Transaction Features.** Let  $G_T = \langle V, E_T \rangle$  be the directed multigraph that represents the buying and selling transactions between two users in the e-commerce platform, where  $V$  is the multigraph vertices set, which contains all the e-commerce users, and  $E_T$  is the multigraph's links set, which contains data on all the transaction interactions between e-commerce users. The links in the transactions multigraph are denoted by  $e_T := (u, v, p, d) \in E_T$ , where  $u, v \in V$  are two e-commerce users,  $p$  is the product which was purchased, and  $d$  is the purchase time and date. Each link  $e_T$  represents a buying transaction of a single product  $p$ , which a user  $u$  bought from a user  $v$  in time and date  $d$ . For each link  $e_T \in E_T$ , we also define the following three additional properties:

1.  $e_T^q$  the transaction's product quantity. Namely, the purchased quantity of the product in the transaction.
2.  $e_T^p$  the purchased product price in US dollars.
3.  $e_T^a := e_T^q \cdot e_T^p$  - the transaction total amount in US dollars.

Using these definitions, we define the features for each  $v \in V$  (see Table 4).

**Feedback Features.** Similar to the transactions-directed multigraph, we can also define the feedback-directed multigraph, which is based on the e-commerce users' feedback. Formally, let  $G_F = \langle V, E_F \rangle$  be the directed multigraph that represents the feedback activities between two users in the e-commerce platforms, where  $V$  is the multigraph vertices set, which contains all the e-commerce users, and  $E_F$  is the multigraph's links set, which contains data on all the feedback interactions between e-commerce users. The links in the feedback multigraph are denoted by  $e_F := (u, v, r, d) \in E_F$ , where  $u, v \in V$  are two e-commerce users,  $r \in \mathbb{Z}$  is the feedback rating, and  $d$  is a the feedback's time and date. Each link  $e_F$  represents a feedback interaction with rating of  $r$ , which a user  $u$  gave a user

Table 1: Features

| Name                               | Description   | Formula   |
|------------------------------------|---|---|
| Transaction Features               |   |   |
| <b>Buy-Trans-Num</b> ( $v$ )       | The total number of buying transactions which $v$ performed. With respect to $G_T$ 's topology, the Buy-Trans-Num( $v$ ) feature is equal to the out-degree of $v$ in $  \{(v, u, p, d) \in E_T \mid \exists u \in V \}  $ the multigraph $G_T$ . | $  \{(v, u, p, d) \in E_T \mid \exists u \in V \}  $                |
| <b>Sell-Trans-Num</b> ( $v$ )      | The total number of selling transactions which $v$ performed. With respect to $G_T$ 's topology, the Sell-Trans-Num( $v$ ) feature is equal to the in-degree of $v$ in the multigraph $G_T$ .   | $  \{(v, u, p, d) \in E_T \mid \exists u \in V \}  $                |
| <b>Unique-Sellers</b> ( $v$ )      | The distinct number of users which $v$ bought products from. With respect to $G_T$ 's topology, the Unique-Sellers( $v$ ) feature is equal to the number of vertices in the multigraph $G_T$ which are connected to $v$ by at least one out-link. | $  \{u \in V \mid \exists (u, v, p, d) \in E_T \}  $                |
| <b>Unique-Buyers</b> ( $v$ )       | The distinct number of users which $v$ sold products to. With respect to $G_T$ 's topology, the Unique-Buyers( $v$ ) feature is equal to the number of vertices in the multigraph $G_T$ which are connected to $v$ by at least one in-link.       | $  \{u \in V \mid \exists (v, u, p, d) \in E_T \}  $                |
| <b>Bidir-Trans-Users</b> ( $v$ )   | The distinct number of users which $v$ sold products to, and also bought products from.   | $  \{u \in V \mid \exists (u, v, p, d), (v, u, p, d) \in E_T \}  $  |
| <b>Max-Buy-Price</b> ( $v$ )       | The maximal paid amount in USD which $v$ paid to another user in a single buying transaction.   | $\max(\{e_p^b \mid \exists e_T := (v, u, p, d) \in E_T, u \in V\})$ |
| <b>Min-Buy-Price</b> ( $v$ )       | The minimal paid amount in USD which $v$ paid to another user in a single buying transaction.   | $\min(\{e_p^b \mid \exists e_T := (v, u, p, d) \in E_T, u \in V\})$ |
| <b>Max-Buy-Quantity</b> ( $v$ )    | The maximal number of products which $v$ bought from another user in a single buying transaction.   | $\max(\{e_p^b \mid \exists e_T := (v, u, p, d) \in E_T, u \in V\})$ |
| <b>Total-Buy-Quantity</b> ( $v$ )  | The total amount in USD which $v$ bought in products from other users.  | $\sum_{\{e_T := (v, u, p, d) \in E_T \mid u \in V\}} e_p^b$         |
| <b>Total-Buy-Amount</b> ( $v$ )    | The total amount in USD which $v$ bought in products from other users.  | $\sum_{\{e_T := (v, u, p, d) \in E_T \mid u \in V\}} e_p^b$         |
| <b>Max-Sell-Price</b> ( $v$ )      | The maximal amount in USD which $v$ received from another user in a single selling transaction.   | $\max(\{e_p^s \mid \exists e_T := (u, v, p, d) \in E_T, u \in V\})$ |
| <b>Min-Sell-Price</b> ( $v$ )      | The minimal amount in USD which $v$ received from another user in a single selling transaction.   | $\min(\{e_p^s \mid \exists e_T := (u, v, p, d) \in E_T, u \in V\})$ |
| <b>Max-Sell-Quantity</b> ( $v$ )   | The maximal number of products which $v$ sold to another user in a single buying transaction.   | $\max(\{e_p^s \mid \exists e_T := (u, v, p, d) \in E_T, u \in V\})$ |
| <b>Total-Sell-Quantity</b> ( $v$ ) | The overall number of products which $v$ sold to other users.   | $\sum_{\{e_T := (v, u, p, d) \in E_T \mid u \in V\}} e_p^s$         |
| <b>Total-Sell-Amount</b> ( $v$ )   | The total amount in USD which $v$ sold in products to other users.  | $\sum_{\{e_T := (v, u, p, d) \in E_T \mid u \in V\}} e_p^s$         |
| Feedback Features                  |   |   |
| <b>Gvn-Fdbk-Nmm</b> ( $v$ )        | The total amount of feedback a user $v$ gave to other users.  | $  \{(v, u, r, d) \in E_F \mid \forall u \in V \}  $                |
| <b>Rev-Fdbk-Num</b> ( $v$ )        | The total amount of feedback a user $v$ received from other users.  | $  \{(u, v, r, d) \in E_F \mid \forall u \in V \}  $                |
| <b>Gvn-Unique-Fdbk</b> ( $v$ )     | The number of unique users which received feedback from $v$ .   | $  \{u \in V \mid \exists (v, u, r, d) \in E_F \}  $                |
| <b>Rev-Unique-Fdbk</b> ( $v$ )     | The number of unique users, which gave feedback to $v$ .  | $  \{u \in V \mid \exists (u, v, r, d) \in E_F \}  $                |
| <b>Bidir-Fdbk-Users</b> ( $v$ )    | The distinct number of users which $v$ gave feedback to, and also received feedback from.   | $  \{u \in V \mid \exists (u, v, p, d), (v, u, p, d) \in E_F \}  $  |
| <b>Gvn-Pos-Fdbk</b> ( $v$ )        | The number of positive feedback ratings a user $v$ received from other users.   | $  \{(u, v, r, d) \in E_F \mid r > 0 \}  $                          |
| <b>Gvn-Neg-Fdbk</b> ( $v$ )        | The number of negative feedback ratings a user $v$ received from other users.   | $  \{(u, v, r, d) \in E_F \mid r < 0 \}  $                          |
| <b>Rev-Pos-Fdbk</b> ( $v$ )        | The number of positive feedback ratings a user $v$ received from other users.   | $  \{(u, v, r, d) \in E_F \mid r > 0 \}  $                          |
| <b>Rev-Neg-Fdbk</b> ( $v$ )        | The number of negative feedback ratings a user $v$ received from other users.   | $  \{(u, v, r, d) \in E_F \mid r < 0 \}  $                          |
| <b>Gvn-Fdbk-RSum</b> ( $v$ )       | The sum of the feedback ratings a user $v$ gave to other users.   | $\sum_{\{(u, v, r, d) \in E_F \mid u \in V\}} r$                    |
| <b>Rev-Fdbk-RSum</b> ( $v$ )       | The sum of the feedback ratings a user $v$ received from other users.   | $\sum_{\{(u, v, r, d) \in E_F \mid v \in V\}} r$                    |
| <b>Gvn-Fdbk-Avg</b> ( $v$ )        | The average of the feedback ratings a user $v$ gave to other users.   | $\frac{\text{Gvn-Fdbk-RSum}(v)}{\text{Gvn-Unique-Fdbk}(v)}$         |
| <b>Rev-Fdbk-Avg</b> ( $v$ )        | The average of the feedback ratings a user $v$ received from other users.   | $\frac{\text{Rev-Fdbk-RSum}(v)}{\text{Rev-Unique-Fdbk}(v)}$         |

$v$  in time and date  $d$ .<sup>6</sup> Using these feedback-directed multigraph definitions, we define features for each user  $v \in V$  (see Table 4).

**Users Details Features.** To construct our supervised learning classifiers, we also utilized the following users' details, which can be mainly extracted from the user's registration form, which exists in many e-commerce platforms:

1. **Birth-Year**( $v$ ) - the declared birth year of  $v$ .
2. **State**( $v$ ) - the declared state of  $v$ . For consistency, we converted the feature State( $v$ ) to an integer using CRC32 hash function [20].
3. **Active-Days**( $v$ ) - the number of days  $v$  was active in the e-commerce platform. In this study, we calculate this feature by calculating the number of days that passed between the date the user created a profile in the e-commerce platform and the last date the user performed a selling or buying transaction.

**Target class.** Every instance in the training set includes a binary target attribute that indicates whether the user was identified as a skill bidder. In this study, we assumed that a list of identified skill bidders (referred to as *skill bidders list*) is provided. These skill bidders are identified by the e-commerce platform experts and are used to train the supervised learning algorithms. A part of this

<sup>6</sup> It is worth to mentioning that some e-commerce platforms provide the sellers the opportunity to also rank the buyers. In this study, we treated feedback which was given from a seller to a buyer as the same as feedback from a buyer to a seller.

list is also used as a ground truth during evaluation of the algorithms as described below.

Although, there are unidentified skill bidders that do not appear in the skill bidders list, as we see next, the fraction such users is expected to be extremely low. Therefore, for the purpose of training supervised learning algorithms we assume that a randomly picked user is a benign user if it does not appear in the skill bidders list.

### 3.2 Selecting Users for the Training Set

In this study we assumed that the actual ratio between skill bidders and benign users in an e-commerce platform is unknown, yet there are indications that the ratio between benign users and skill bidders is much in favor of the benign users. For example, in other online platforms, such as online social networks, there is a clear indication that the ratio between the number of benign entities and the number of malicious entities is in favor of the benign entities. The official Facebook estimation is that approximately 3-4% of Facebook users are malicious users [9], and according to Rahman et al. [50] about 13% of applications in Facebook are malicious. To construct our classifiers, we choose to create a balanced training set with an equal number of benign users and skill bidders, similar to the methodology used by Guha et al. [32] to predict trust, by Leskovec et al. [41] to predict positive and negative links, and by Fire et al. [30] to identify fake users in Facebook.

### 3.3 Choosing a Supervised Learning Algorithm

To identify the supervised learning algorithm which yields the best classification results on our datasets, we trained the classifiers on all the users in the provided skill bidders list and an equal number of benign users. We then extracted for each user all the 31 features, which were described in Section 3.1. Afterwards, we used the constructed balanced training set and fed it to Weka [33], a popular suite of machine learning software. We used Weka’s OneR, C4.5 (J48) decision tree, K-Nearest-Neighbors (IBk; with  $K=3$ ), Naive-Bayes, Random-Forest, LogitBoost, Rotation-Forest, and Bagging implementations of the corresponding algorithms. For each of these algorithms, all of the configurable parameters were set to their default values.

We evaluated each classifier using the 10-fold cross validation method and calculated the True-Positive (TP) rate, False-Positive (FP) rate, F-Measure (FM) value, and the Area-Under-Curve (AUC) measure. These metrics assisted us in selecting the best supervised learning algorithm for identifying skill bidders.

### 3.4 Supervised Learning Algorithm Evaluation

To evaluate the precision of the selected supervised learning algorithm, which received the highest AUC in “the wild” on a real e-commerce imbalanced dataset, we performed the following steps:<sup>7</sup>

1. We created a balanced training set by randomly selecting 90% of the skill bidder users in the skill bidders list, and by randomly selecting an equal number of benign users.
2. The remaining 10% of the users in the skill bidders list, which were excluded in step 1, were utilized to construct several test sets having various imbalance rates. We created five test sets with 2, 5, 10, 20, and 100 benign users per single skill bidder.
3. For each user in the constructed training and testing sets, we extracted all the features, which were described in Section 3.1.
4. Using the balanced training set and the selected supervised learning algorithm, we constructed a skill bidder identification classifier.
5. For each user in each one of the five imbalanced testing sets, we used the constructed balanced classifier to predict the user likelihood of being a skill bidder.
6. For each imbalanced testing set, we calculated the classifier precision at the top  $k$  (*precision@k*). Namely, for each one of the five imbalanced datasets and for an integer  $k \in [1, n]$ , we calculated the percent of the top  $k$  users which received the highest likelihood of being skill bidders and were actually skill bidders.
7. Lastly, to reduce variability, we repeated steps 1 to 6 three times, and averaged the obtained precision at  $k$  results for each  $k$ .

## 4 Empirical Evaluation

### 4.1 The E-commerce Dataset

In this study, we used anonymized datasets provided by one of the largest e-commerce companies in the world to evaluate our skill-bidder detection algorithms. Additionally, in order to query the provided datasets, we used Hadoop infrastructure which includes several thousands of Hadoop nodes. The e-commerce dataset included information of several billions of buying and selling transactions, as well as several billions feedback transactions. Each feedback transaction included feedback ratings with one of three possible values: negative feedback (-1), neutral feedback (0), and positive feedback (+1). All of the transactions in the e-commerce dataset were actual transactions performed by over several hundred million platform users through the end of 2012. Furthermore, we were provided with a list of 187,224 user which were marked as skill bidders by the company’s

<sup>7</sup> An additional method for evaluating the classifier precision is to manually validate the classifiers results. However, due to privacy limitations, these types of methods were not available in this study.

proprietary algorithms. To select a benign users list for our training and testing sets throughout this study, we randomly selected from the e-commerce dataset a list of 500,000 seller-users which performed at least one sell transaction. We then removed from the benign users list all the users which also appeared in the shill bidders list.

## 4.2 Experiment Setup

We evaluated various supervised learning algorithms to construct classifiers which can identify which of the users are shill bidders. Using the provided 187,224 shill bidders and an additional 187,224 benign users from the e-commerce dataset, we constructed a balanced training set and evaluated various supervised learning algorithms (see Section 3.3). Furthermore, using the e-commerce dataset, we constructed five imbalanced testing sets with the following ratios:

1. *1 to 2* - with 18,722 shill bidders and 37,444 benign users.
2. *1 to 5* - with 18,722 shill bidders and 93,610 benign users.
3. *1 to 10* - with 18,722 shill bidders and 187,220 benign users.
4. *1 to 20* - with 15,000 shill bidders and 300,000 benign users.
5. *1 to 100* - with 3,200 shill bidders and 320,000 benign users.

We used these imbalanced datasets to evaluate the constructed classifiers' precision at  $k$  for  $k \in [1, 30, 1000]$ , using the methods described in Section 3.4.

## 4.3 Results

In this section, we present the results obtained by the method described in Section 3.

Table 2: Supervised Learning Classifiers Results on Balanced Training Set

| Classifier                | TP    | FP    | FN    | AUC   |
|---------------------------|-------|-------|-------|-------|
| <b>OneR</b>               | 0.800 | 0.252 | 0.780 | 0.774 |
| <b>Naïve-Bayes</b>        | 0.886 | 0.764 | 0.645 | 0.752 |
| <b>Decision-Tree(J48)</b> | 0.822 | 0.193 | 0.816 | 0.860 |
| <b>Random-Forest</b>      | 0.854 | 0.230 | 0.820 | 0.885 |
| <b>Bagging</b>            | 0.834 | 0.179 | 0.829 | 0.902 |
| <b>LogitBoost</b>         | 0.811 | 0.170 | 0.819 | 0.901 |
| <b>Rotation-Forest</b>    | 0.845 | 0.173 | 0.838 | 0.912 |

According to our evaluation results, among all tested supervised learning algorithms, the Rotation-Forest classifiers performed best on the e-commerce dataset, with an especially good AUC result of 0.912 and a TP rate of 0.845 (see Table 2). Therefore, we chose to construct shill the bidder identification classifiers using the Rotation-Forest algorithm and to evaluate the constructed

classifiers' precision at  $k$  for the five imbalanced testing sets, which were defined in Section 3.4. The results showing the classifiers' precision at  $k$  average values for all five testing sets are presented in Figure 2. According to the evaluation results, the developed shill bidders identification classifiers presented high precision at  $k$ , with precision at 1,000 of 1, 1, 0.999, 0.925, and 0.358, when the ratio between the number of shill bidder and the number of benign users was 1 to 2, 1 to 5, 1 to 10, 1 to 20, and 1 to 100, respectively (see Figure 2). These results indicate that the presented shill bidder identification algorithms can detect shill bidders far better than a random algorithm. Moreover, the presented algorithm gave very high precision rates when the percentage of the shill bidder in the e-commerce dataset was at least 5%.

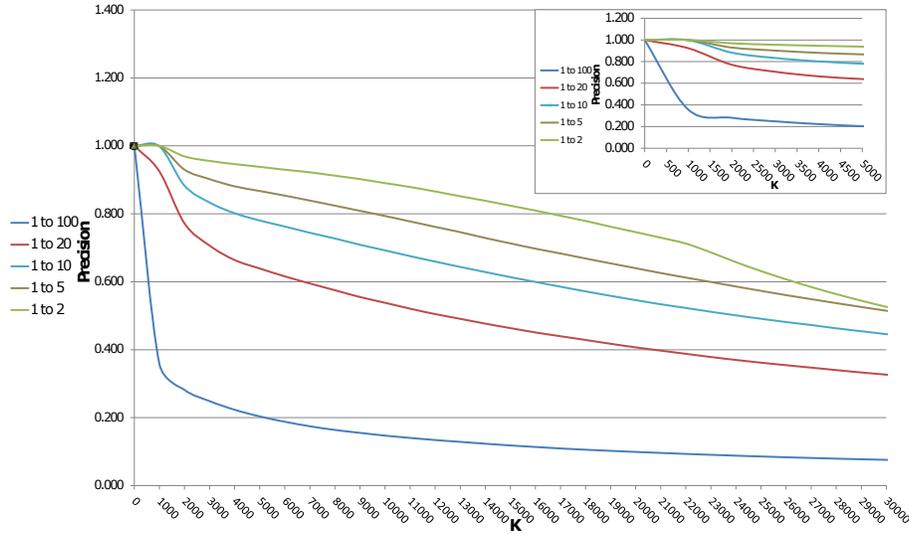


Fig. 2: Rotation-Forest classifier precision at  $k$  results - for five imbalance testing sets.

## 5 The Shill Bidder Ecosystem

In this study, we utilize the provided shill bidders list to empirically analyze the characteristics of shill bidders and to study the shill bidder ecosystem, i.e., the interactions of shill bidders with each other.

To analyze the shill bidders' characteristics using the e-commerce dataset, we calculated the average and median values for all the numeric features we defined in Section 3.1 for the users in the provided shill bidders list. Additionally, to better understand which features assist the supervised learning algorithms to

Table 3: Features’ Information Gain Scores

|                     | Info Gain Score |
|---------------------|-----------------|
| Min-Sell-Price      | 0.268           |
| Sell-Trans-Num      | 0.248           |
| Total-Sell-Quantity | 0.247           |
| Unique-Buyers       | 0.240           |
| State               | 0.240           |
| Total-Sell-Amount   | 0.196           |
| Rcv-Unique-Fdbk     | 0.188           |
| Rcv-Fdbk-Num        | 0.187           |
| Min-Buy-Price       | 0.180           |
| Rcv-Fdbk-RSum       | 0.162           |
| Rcv-Pos-Fdbk        | 0.162           |
| Gvn-Unique-Fdbk     | 0.154           |
| Gvn-Fdbk-Num        | 0.153           |
| Gvn-Pos-Fdbk        | 0.153           |
| Gvn-Fdbk-RSum       | 0.153           |
| Fdbk-Bi-Degree      | 0.150           |
| Total-Buy-Quantity  | 0.142           |
| Buy-Trans-Num       | 0.143           |
| Unique-Sellers      | 0.128           |
| Max-Sell-Price      | 0.116           |
| Total-Buy-Amount    | 0.120           |
| Rcv-Fdbk-Avg        | 0.112           |
| Gvn-Fdbk-Avg        | 0.089           |
| Max-Buy-Price       | 0.076           |
| Active-Days         | 0.079           |
| Max-Buy-Quantity    | 0.070           |
| Gvn-Neg-Fdbk        | 0.059           |
| Rcv-Neg-Fdbk        | 0.047           |
| Max-Sell-Quantity   | 0.016           |
| Birth-Year          | 0.006           |
| Trans-Bi-Degree     | 0.002           |

Table 4: Features Average and Median Values Ratio Between Shill Bidders and Random Seller-Users

|                     | AverageRatio | Median Ratio |
|---------------------|--------------|--------------|
| Buy-Trans-Num       | 2.689        | 4.421        |
| Sell-Trans-Number   | 3.341        | 16.75        |
| Unique-Sellers      | 2.642        | 4.094        |
| Unique-Buyers       | 3.506        | 13.625       |
| Bidir-Trans-Users   | 1.412        | 2            |
| Max-Buy-Price       | 1.663        | 2.168        |
| Min-Buy-Price       | 0.035        | 0.966        |
| Max-Buy-Quantity    | 1.866        | 1.5          |
| Total-Buy-Quantity  | 2.683        | 4.525        |
| Total-Buy-Amount    | 2.441        | 4.295        |
| Max-Sell-Price      | 2.693        | 2.676        |
| Min-Sell-Price      | 0.063        | 0.227        |
| Max-Sell-Quantity   | 1.373        | 1            |
| Total-Sell-Quantity | 3.36         | 16.875       |
| Total-Sell-Amount   | 3.583        | 9.874        |
| Gvn-Fdbk-Num        | 2.76         | 6.286        |
| Rcv-Fdbk-Num        | 2.964        | 6.463        |
| Gvn-Unique-Fdbk     | 2.921        | 6.054        |
| Rcv-Unique-Fdbk     | 3.112        | 6.222        |
| Bidir-Fdbk-Users    | 2.989        | 6.091        |
| Gvn-Pos-Fdbk        | 2.76         | 6.366        |
| Gvn-Neg-Fdbk        | 2.748        | inf          |
| Rcv-Pos-Fdbk        | 2.786        | 5.702        |
| Rcv-Neg-Fdbk        | 2.04         | inf          |
| Gvn-Fdbk-RSum       | 2.76         | 6.317        |
| Rcv-Fdbk-RSum       | 2.789        | 5.66         |
| Gvn-Fdbk-Avg        | 1.017        | 0.991        |
| Rcv-Fdbk-Avg        | 1.025        | 0.994        |
| Birth-Year          | 1            | 0.999        |
| Active-Days         | 1.387        | 1.545        |
| Birth-Year          | 1            | 0.999        |
| Active-Days         | 1.387        | 1.545        |

distinguish between shill bidders and benign, we used the balanced training set (see Section 4.2) to calculate the features importance using Weka’s Information Gain features’ selection algorithm. The results of the shill bidders’ characteristics analysis are presented in Tables 3 and 4.

From the shill bidders characteristics analysis results presented in Table 4, it can be observed that shill bidders behaved differently from the random seller-users in the following manner: First, on average, shill bidders are active for more days and perform far more selling and buying transactions than random sellers. These results may indicate that shill bidder users are, in general, active users, which perform many buying and selling transactions. Second, on average, shill bidders sell more products to unique buyers and buy more products from unique sellers than random seller-users. However, the shill bidders’ buying and selling minimum price was on average much less than the buying and selling minimum price of random-seller users. We believe that these results are due to the shill bidders’ attempts to maximum their profits and minimize their losses when the buy or sell feedback. Third, on average, shill bidders received more negative feedback than random seller-users. We assume that these results indicate that in the end many shill bidders utilize their obtained positive feedback to mislead other users in the platform, which in return gives the shill bidders negative feedback. Lastly, on average, shill bidders gave more negative feedback than random seller-users. We believe that this result may indicate that shill bidders are also being utilized to perform sybil attacks [42]. We hope to prove these assumptions in a future study. Additionally, we discovered that the shill bidders

had 798 unique State feature values, while the randomly selected seller-users had 1,435 unique State feature values. Furthermore, the most common State feature value among the skill bidders was the “default” value which appeared in the details of 134,979 skill bidders, while “default” state value appeared only in the details of 46,805 randomly selected seller-users.

From the features’ Information Gain scores results, which are presented in Table 3, it can be observed that the Min-Sell-Price, Sell-Trans-Num, and Total-Sell-Quantity features received the highest Information Gain scores. We believe that these features received the highest scores due to the skill bidders behavioral patterns. In many cases, skill bidders attempted to decrease their losses by selling cheap products using many transactions in order to collect a great deal of positive feedback, to spend as little money as possible.

To study the skill bidder ecosystem, we analyzed the feedback graph created by the skill bidders. The skill bidder feedback graph can assist us in understanding the “big picture” beyond the skill bidders and their interactions, and even assist us in understanding the skill bidders’ working methods. We defined the feedback graph as following: Given a list  $\hat{V} \subseteq V$  of e-commerce users, we can define  $\hat{V}$ ’s feedback graph to be a weighted directed graph, where each directed weighted link is defined to be the amount of feedback user  $u \in \hat{V}$  gave user  $v \in \hat{V}$ . Formally, the feedback graph is defined to be  $H_{\hat{V}} := \langle \hat{V}, E_{\hat{V}} \rangle$ , where each link  $e_{\hat{V}} \in E_{\hat{V}}$  in the graph defined as

$$e_{\hat{V}} := \{(u, v, w) | \exists u, v \in \hat{V} \text{ and } w = |\{(u, v, r, d) \in E_F\}|\}.$$

Using the provided skill bidders list  $B$ , we first constructed the feedback graph  $H_B := \langle B, E_B \rangle$  as explained above. We then calculated various graph properties using graph theory algorithms. Namely, we mostly used the igraph software package [27] to calculate the following graph’s properties:

1. number of vertices.
2. number of links.
3. maximum and minimum link weight.
4. average link weight.
5. number of bidirectional links ( $|\{(u, v, w) \in E_B | \exists (v, u, w) \in E_B\}|$ ).
6. density.
7. components number.
8. largest component size.

Additionally, we used the igraph implementation of the Bron-Kerbosch algorithm [21] to find the maximal cliques in the graph<sup>8</sup> and calculate their distribution. By identifying cliques, we can identify a group of skill bidders who worked together. Lastly, we used a random sample to select a list with an equal number of seller-users, and we compared the properties of the feedback graph created by the random sampled seller-users of those the feedback graph created by the skill bidders.

<sup>8</sup> The used igraph cliques detections algorithm implementation treats the directed graph as an undirected graph.

From the ecosystem analysis results, it can be observed that the skill bidder feedback graph is a relatively dense graph that spawns 1,805,199 directed links between 156,769 skill bidders, with an average of 1.31 feedback occurrences per link (see Table 5 and Figure 1). Additionally, 79.09% of the skill bidders are located in a single component. Furthermore, according to the maximal clique detection results, it can be noticed that in contrast to the random user feedback graph, the skill bidder feedback graph consists of many cliques (see Table 6). These results may indicate that many skill bidders work together and assist each other to receive positive feedback. Another alternative explanation is that skill bidders open several accounts and use them to boost their own reputation, or they sell feedback to other users.

Table 5: Feedback Graphs Properties

|  | <b>Shill Bidders<br/>Feedback Graph</b> | <b>Random Users<br/>Feedback Graph</b> |
|--|---|--|
| <b>Number of Users</b>                               | 187,224                                 | 187,224                                |
| <b>Number of Feedback<br/>Between Users</b>          | 2,391,312                               | 91,863                                 |
| <b>Number of Positive<br/>Feedback Between Users</b> | 2,373,993                               | 91,097                                 |
| <b>Number of Negative<br/>Feedback Between Users</b> | 8,786                                   | 341                                    |
| <b>Number of Non-Isolated Users</b>                  | 156,769                                 | 35,599                                 |
| <b>Number of Links</b>                               | 1,805,199                               | 67,383                                 |
| <b>Average Link Weight</b>                           | 1.31                                    | 1.35                                   |
| <b>Max Link Weight</b>                               | 914                                     | 219                                    |
| <b>Min Link Weight</b>                               | -11                                     | -8                                     |
| <b>Number of Bidirectional Links</b>                 | 1,675,411                               | 59,692                                 |
| <b>Density</b>                                       | $7.35 \cdot 10^{-5}$                    | $5.32 \cdot 10^{-5}$                   |
| <b>Component Number</b>                              | 8,123                                   | 9,309                                  |
| <b>Largest Component Size</b>                        | 148,072 (79.09%)                        | 21,443 (11.45%)                        |
| <b>Maximal Clique Size</b>                           | 7                                       | 3                                      |
| <b>Maximal Clique Number</b>                         | 895,844                                 | 37,080                                 |

Table 6: Maximal Cliques Size Distributions

| <b>Clique<br/>Size</b> | <b>Random<br/>Feedback Graph</b> | <b>Shill<br/>Feedback Graph</b> |
|------------------------|----------------------------------|---------------------------------|
| <b>3</b>               | 260                              | 66,892                          |
| <b>4</b>               | 0                                | 3,945                           |
| <b>5</b>               | 0                                | 803                             |
| <b>6</b>               | 0                                | 173                             |
| <b>7</b>               | 0                                | 23                              |

## 6 Conclusions

According to the presented method various user features are first extracted from user transactions (buying and selling) and from feedback activities (giving and receiving). Second, supervised learning algorithms are used to train a model for classifying users into shill bidders and legitimate accounts. We evaluated the algorithms using a real large-scale anonymized e-commerce dataset. This dataset includes over several billions of buying and selling occurrences performed by several hundred million users, as well as several billions of feedback occurrences they gave or received through the end of 2012. The dataset also includes a list of 187,224 users, which were marked as shill bidders in e-commerce platform systems and were used as ground truth for training and testing our algorithms. Evaluation results of the presented method show the area under the ROC curve (AUC) of up to 0.912 and precision at 1,000 of 0.999 when the ratio between the shill bidders and the benign users was 1 to 10 (see Section 4.3).

By analyzing the e-commerce dataset, we also empirically studied the characteristics of shill bidders, as well as the shill bidder ecosystem. As a result of this analysis, we discovered that, on average, shill bidders were more active, performed more selling and buying of transactions, and gave and received more feedback compared to randomly selected seller-users (see Table 4). Moreover, we also discovered that shill bidders gave more negative feedback compared to randomly selected seller-users. These results may indicate that shill bidders can be used to also perform sybil attacks [42]. These sybil attacks can be against targeted platform users, such as shill bidders competitors, and aim to damage the targeted users' reputations by giving them unjust negative feedback. Additionally, we discovered that the shill bidder feedback graph, which was constructed from all the feedback links between each two shill bidders (see Section 5), is a relatively dense graph with 1,805,199 links among 156,769 shill bidders (see Table 5 and Figure 1). Furthermore, in the shill bidders feedback graph we identified 66,892 cliques with at least 3 shill bidders, and 23 cliques with at least 7 shill bidders (see Table 6). These results indicate that many shill users collaborate to increase their overall reputations. Moreover, these results may indicate the existence of e-commerce bots which perform automatic buying and selling transactions. These bots also submit feedback to each other and to other users.

We believe that these observations regarding the shill bidder ecosystem can assist in improving the detection of shill bidders in following manner: First, we can utilize the results regarding the shill bidder graph structure and improve our shill bidder identification classifiers by extracting additional graph structure-based features, such as the number of cliques a user is member of and the number of shill bidders the user is connected to. Second, we can use the fact that many shill bidders are connected to each other (see Figure 1 and Table 5) and set our shill bidder identification classifier to identify shill bidders among users which are connected to several shill bidders, instead of applying the classifier on random sets of users chosen out of several hundred million e-commerce platform users. We believe that this technique of focusing the shill bidder identification classifier on these connected users can identify shill bidders with even higher

precision. Lastly, we can utilize the results that indicate the shill bidders tend to formulate relatively large cliques (see Table 6) and use various clique identification algorithms to identify large cliques in the feedback graph created from all the e-commerce users' feedback interactions. We believe that large cliques in this feedback graph have a high likelihood of containing shill bidders. We hope to verify these three assumptions in our future research.

The study presented here offers many additional future research directions to pursue. One possible research direction is to analyze not only the structured data which was extracted from the e-commerce users' feedback and transactions, but also to use Natural Language Processing (NLP) algorithms to analyze the users' content data, such as the feedback comments and the products information page. An additional possible research direction is to utilize various clustering algorithms to identify shill bidders. Another possible research direction is to construct classifiers, which utilize similar features to those presented in Section 3.1, to identify other types of malicious users, such as fraudsters who sell fictional products.

## References

1. 10 most expensive items ever listed on ebay - page 3 - cbs news. <https://www.cbsnews.com/media/10-most-expensive-items-ever-listed-on-ebay/3/>, (Accessed on 10/08/2018)
2. Alibaba group announces march quarter 2018 results and full fiscal year 2018 results. [https://www.alibabagroup.com/en/news/press\\_pdf/p180504.pdf](https://www.alibabagroup.com/en/news/press_pdf/p180504.pdf), (Accessed on 10/08/2018)
3. Amazon total active users 2013-2016 — statistic. <https://www.statista.com/statistics/476196/number-of-active-amazon-customer-accounts-quarter/>, (Accessed on 10/08/2018)
4. Amazon's fake review problem is now worse than ever, study suggests. <https://www.forbes.com/sites/emmawoollacott/2017/09/09/exclusive-amazons-fake-review-problem-is-now-worse-than-ever/#5a2358997c0f>, (Accessed on 10/07/2018)
5. Battling fake accounts, twitter to slash millions of followers - the new york times. <https://www.nytimes.com/2018/07/11/technology/twitter-fake-followers.html>, (Accessed on 10/07/2018)
6. Company info — facebook newsroom. <https://newsroom.fb.com/company-info/>, (Accessed on 10/07/2018)
7. ebay: number of users 2018 — statistic. <https://www.statista.com/statistics/242235/number-of-ebays-total-active-users/>, (Accessed on 10/07/2018)
8. Exclusive: Twitter is suspending millions of bots and fake accounts every day to fight disinformation - the washington post. [https://www.washingtonpost.com/technology/2018/07/06/twitter-is-sweeping-out-fake-accounts-like-never-before-putting-user-growth-risk/?noredirect=on&utm\\_term=.c3dfc8c59f7e](https://www.washingtonpost.com/technology/2018/07/06/twitter-is-sweeping-out-fake-accounts-like-never-before-putting-user-growth-risk/?noredirect=on&utm_term=.c3dfc8c59f7e), (Accessed on 10/07/2018)
9. Facebook - financials - sec filings details. <https://investor.fb.com/financials/sec-filings-details/default.aspx?FilingId=12512043>, (Accessed on 10/06/2018)

10. Facebook will ban sellers of shoddy products - wsj. <https://www.wsj.com/articles/facebook-will-ban-sellers-of-shoddy-products-1528794000>, (Accessed on 10/08/2018)
11. Global ecommerce 2019 - emarketer trends, forecasts & statistics. <https://www.emarketer.com/content/global-ecommerce-2019>, (Accessed on 03/03/2020)
12. Global retail e-commerce market size 2014-2021 — statista. <https://www.statista.com/statistics/379046/worldwide-retail-e-commerce-sales/>, (Accessed on 10/06/2018)
13. I get paid to write fake reviews for amazon — cracked.com. <http://www.cracked.com/personal-experiences-2376-i-get-paid-to-write-fake-reviews-amazon.html>, (Accessed on 10/07/2018)
14. Shill bidding. <https://www.nycriminallawyer.com/fraud-charge/auction-fraud/shill-bidding/>, (Accessed on 10/08/2018)
15. There are 168 million active buyers on ebay right now (infographic) - small business trends. <https://smallbiztrends.com/2018/03/ebay-statistics-march-2018.html>, (Accessed on 10/07/2018)
16. Tophatter : Faq. <http://help.tophatter.com/customer/portal/articles/458685-what-is-the-penalty-for-shill-bidding->, (Accessed on 10/07/2018)
17. What is shill bidding? – flippa help center. <https://support.flippa.com/hc/en-us/articles/202469674-What-is-Shill-Bidding->, (Accessed on 10/07/2018)
18. Why twitter is still teeming with millions of bots. <https://mashable.com/2017/10/16/twitter-bots-here-to-stay/#VN71yZjyYmqu>, (Accessed on 10/07/2018)
19. Beyene, Y., Faloutsos, M., Chau, D.H., Faloutsos, C.: The ebay graph: How do online auction users interact? In: INFOCOM Workshops 2008, IEEE. pp. 1–6. IEEE (2008)
20. Black, R.: Fast crc32 in software. ATM Document Collection **3** (1994)
21. Bron, C., Kerbosch, J.: Algorithm 457: finding all cliques of an undirected graph. Communications of the ACM **16**(9), 575–577 (1973)
22. Chakraborty, I., Kosmopoulou, G.: Auctions with shill bidding. Economic Theory **24**(2), 271–287 (2004)
23. Chang, W.H., Chang, J.S.: A novel two-stage phased modeling framework for early fraud detection in online auctions. Expert Systems with Applications **38**(9), 11244 – 11260 (2011). <https://doi.org/http://dx.doi.org/10.1016/j.eswa.2011.02.172>, <http://www.sciencedirect.com/science/article/pii/S0957417411003964>
24. Chau, D.H., Faloutsos, C.: Fraud detection in electronic auction. In: European Web Mining Forum at ECML/PKDD. pp. 87–97 (2005)
25. Chau, D.H., Pandit, S., Faloutsos, C.: Detecting fraudulent personalities in networks of online auctioneers. In: Knowledge Discovery in Databases: PKDD 2006, pp. 103–114. Springer (2006)
26. Chawla, N.V., Japkowicz, N., Kotcz, A.: Editorial: special issue on learning from imbalanced data sets. ACM Sigkdd Explorations Newsletter **6**(1), 1–6 (2004)
27. Csardi, G., Nepusz, T.: The igraph software package for complex network research. InterJournal, Complex Systems **1695**(5) (2006)
28. eBay: Shill bidding policy (2014), <http://pages.ebay.com/help/policies/seller-shill-bidding.html> [Online; accessed 10/07/2018]
29. Etsy: Etsy 2017 annual report. <https://investors.etsy.com/~media/Files/E/Etsy-IR/annual-report-proxy-materials/etsy-ar2017.pdf>, (Accessed on 10/08/2018)

30. Fire, M., Kagan, D., Elyashar, A., Elovici, Y.: Friend or foe? fake profile identification in online social networks. arXiv preprint arXiv:1303.3751 (2013)
31. Ganguly, S., Sadaoui, S.: Online detection of shill bidding fraud based on machine learning techniques. In: International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems. pp. 303–314. Springer (2018)
32. Guha, R., Kumar, R., Raghavan, P., Tomkins, A.: Propagation of trust and distrust. In: Proceedings of the 13th international conference on World Wide Web. pp. 403–412. ACM (2004)
33. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.: The weka data mining software: an update. ACM SIGKDD Explorations Newsletter **11**(1), 10–18 (2009)
34. Hooi, B., Song, H.A., Beutel, A., Shah, N., Shin, K., Faloutsos, C.: Fraudar: Bounding graph fraud in the face of camouflage. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 895–904. ACM (2016)
35. Javelin: 2017 financial impact of fraud study: Exploring the impact of fraud in a digital world. [https://s3.amazonaws.com/dive\\_static/paychek/Financial\\_Impact\\_of\\_Fraud\\_Study\\_FINAL.pdf](https://s3.amazonaws.com/dive_static/paychek/Financial_Impact_of_Fraud_Study_FINAL.pdf) (2017), (Accessed on 10/07/2018)
36. Jones, C.: Ecommerce is growing nicely while mcommerce is on a tear. Forbes (October 2013), <http://www.forbes.com/sites/chuckjones/2013/10/02/ecommerce-is-growing-nicely-while-mcommerce-is-on-a-tear/> [Online; accessed 10/07/2018]
37. Kabus, P., Terpstra, W.W., Cilia, M., Buchmann, A.P.: Addressing cheating in distributed mmogs. In: Proceedings of 4th ACM SIGCOMM Workshop on Network and System Support for Games. pp. 1–6. NetGames '05, ACM, New York, NY, USA (2005). <https://doi.org/10.1145/1103599.1103607>, <http://doi.acm.org/10.1145/1103599.1103607>
38. Kaghazgaran, P., Caverlee, J., Alfifi, M.: Behavioral analysis of review fraud: Linking malicious crowdsourcing to amazon and beyond. In: ICWSM. pp. 560–563 (2017)
39. Kauffman, R.J., Wood, C.A.: Running up the bid: detecting, predicting, and preventing reserve price shilling in online auctions. In: Proceedings of the 5th international conference on Electronic commerce. pp. 259–265. ACM (2003)
40. Kumar, S., Cheng, J., Leskovec, J., Subrahmanian, V.: An army of me: Sockpuppets in online discussion communities. In: Proceedings of the 26th International Conference on World Wide Web. pp. 857–866. International World Wide Web Conferences Steering Committee (2017)
41. Leskovec, J., Huttenlocher, D., Kleinberg, J.: Predicting positive and negative links in online social networks. In: Proceedings of the 19th international conference on World wide web. pp. 641–650. ACM (2010)
42. Levine, B., Shields, C., Margolin, N.: A survey of solutions to the sybil attack. University of Massachusetts Amherst, Amherst, MA (2006)
43. Lockhart, J.: <http://www.makeuseof.com/tag/can-you-really-win-almost-any-ebay-auction-by-sniping/>, [Online; accessed 10/07/2018]
44. Lucking-Reiley, D., Bryan, D., Prasad, N., Reeves, D.: Pennies from ebay: The determinants of price in online auctions. *The Journal of Industrial Economics* **55**(2), 223–233 (2007)
45. Majadi, N., Trevathan, J., Bergmann, N.: Analysis on bidding behaviours for detecting shill bidders in online auctions. In: Computer and Information Technology (CIT), 2016 IEEE International Conference on. pp. 383–390. IEEE (2016)

46. Mitterhofer, S., Platzer, C., Kruegel, C., Kirda, E.: Server-side bot detection in massive multiplayer online games. *IEEE Security and Privacy* **7**(3), 29–36 (2009)
47. Mudambi, S.M., Schuff, D.: What makes a helpful online review? a study of customer reviews on amazon. com. *MIS quarterly* **34**(1), 185–200 (2010)
48. Nicholas Confessore, Gabriel J.X. Dance, R.H., Hansen, M.: The follower factory - the new york times. <https://www.nytimes.com/interactive/2018/01/27/technology/social-media-bots.html>, (Accessed on 10/08/2018)
49. Pandit, S., Chau, D.H., Wang, S., Faloutsos, C.: Netprobe: a fast and scalable system for fraud detection in online auction networks. In: *Proceedings of the 16th international conference on World Wide Web*. pp. 201–210. ACM (2007)
50. Rahman, M.S., Huang, T.K., Madhyastha, H.V., Faloutsos, M.: Frappe: detecting malicious facebook applications. In: *Proceedings of the 8th international conference on Emerging networking experiments and technologies*. pp. 313–324. ACM (2012)
51. Resnick, P., Kuwabara, K., Zeckhauser, R., Friedman, E.: Reputation systems. *Communications of the ACM* **43**(12), 45–48 (2000)
52. Roth, A.E., Ockenfels, A.: Last minute bidding and the rules for ending second-price auctions: Theory and evidence from a natural experiment on the internet. Tech. rep., National bureau of economic research (2000)
53. Techopedia.com: <http://www.techopedia.com/definition/27959/auction-sniping/>, [Online; accessed 10/07/2018]
54. Trevathan, J., Read, W.: Detecting shill bidding in online english auctions. *Handbook of research on social and organizational liabilities in information security* pp. 446–470 (2009)
55. Tsang, S., Koh, Y.S., Dobbie, G., Alam, S.: Detecting online auction shilling frauds using supervised learning. *Expert Systems with Applications* **41**(6), 3027 – 3040 (2014). <https://doi.org/http://dx.doi.org/10.1016/j.eswa.2013.10.033>, <http://www.sciencedirect.com/science/article/pii/S0957417413008506>
56. Wang, A.H.: Don't follow me: Spam detection in twitter. In: *Security and Cryptography (SECRYPT)*, *Proceedings of the 2010 International Conference on*. pp. 1–10. IEEE (2010)
57. Zhang, Y., Bian, J., Zhu, W.: Trust fraud: A crucial challenge for china's e-commerce market. *Electronic Commerce Research and Applications* **12**(5), 299 – 308 (2013). <https://doi.org/http://dx.doi.org/10.1016/j.elerap.2012.11.005>, <http://www.sciencedirect.com/science/article/pii/S1567422312001202>, chinese E-Commerce
58. Zhou, S., Mondragón, R.J.: The rich-club phenomenon in the internet topology. *Communications Letters, IEEE* **8**(3), 180–182 (2004)