

Detecting Unusual User Behaviour to Identify Hijacked Internet Auctions Accounts^{*}

Marek Zachara¹ and Dariusz Pałka²

¹ AGH University of Science and Technology, Poland
mzachara@agh.edu.pl

² Pedagogical University of Cracow, Poland
dpalka@up.krakow.pl

Abstract. For over 15 years auction services have grown rapidly, constituting a major part of e-commerce worldwide. Unfortunately, they also provide opportunities for criminals to distribute illicit goods, launder money or commit other types of fraud. This calls for methods to mitigate this threat. The following paper discusses the methods of identifying the accounts of users participating in internet auctions that have been hijacked (taken over) by malicious individuals and utilised for fraudulent purposes. Two primary methods are described, monitoring users' activities (e.g. the number of auctions created over time) with EWMA and clustering similar auction categories into groups for the purpose of assessing users' sellers profiles and detecting their sudden changes. These methods, utilised together allow for real-time detection of suspicious accounts. The proposed models are validated on real data gathered from an auction web site.

Keywords: internet auctions, identity theft, anomaly detection.

1 Internet Auctions - Introduction

Since the launch of eBay in 1995, internet auctions have become an important part of the global marketplace. According to the eBay annual report, their income from transactions amounted to 7.7 billion dollars in 2009. Assuming an average fee for a transaction to be below 10%, the total sales through eBay would amount to around 100 billion dollars, compared to 135 billion dollars of total e-commerce retail sales in the US [19] during the same year. There are certainly other auction services beside eBay, but even considering only the numbers related to eBay (which is certainly the largest one), the importance of this transaction medium is obvious.

One of the primary reasons for the success of auction services is the low cost of entry. A person does not need any specific tools nor formalities to start selling their products (or information services). This results in a large number of both sellers and buyers registered with auction sites. A large user-base of

* This work is partially supported by NCBiR grant 0021/R/ID2.

the sellers means statistically high chances of users with weak passwords or otherwise vulnerable to hacking methods. A huge amount of buyers, on the other hand, provide an excellent opportunity to find those interested in illicit goods or susceptible to various scam methods. As a result, auction systems are an important medium for criminals, granting them means of expanding their illegal activities, including fraud and/or the provision of prohibited goods. Left unmitigated, this would constitute a serious threat to public security.

Although most readers are probably familiar with how internet auctions work, a brief explanation will be provided here for reference purposes. A person willing to sell an item, posts its description (often with photos) and an initial asking price at an auction site. Other users can view the offer, may ask additional questions and may also bid a certain sum for the item. Auctions usually end after a specific time (e.g. 14 days), with the item sold to the highest bidder. There might be other types of offers (e.g. a fixed price, multiple items, etc.), but in all cases the transaction is concluded between two registered users of the auction service.

After each transaction the parties have a chance to evaluate it by posting their comments and ratings of the other party. Such ratings for each user are usually aggregated into an overall reputation rating (e.g. a person with 96 positive and 3 negative 'comments' would have a rating of 93).

The reputation system. i.e. the method of calculating the reputation rating and the actual numbers are vital to an auctions system. Contrary to traditional sales scenario, where both parties meet in person and the goods are exchanged for money at the same time, purchases made over the internet usually take much longer, with money often being paid up-front and the goods delivered after a few days. Buyers are therefore likely to make the buying decision based on their trust that sellers will keep their part of the bargain. This trust is likely to be higher if a lot of other users have already concluded transactions with this particular seller, and were satisfied with them, which would be reflected in the sellers reputation rating. Similarly, the seller is more likely to offer e.g. a CoD option to a buyer with a good reputation standing. Although the reputation rating is valuable to every user, it is vital to sellers, as it will directly affect their business and profit.

The reputation system and the 'snowball effect' of an increasing number of buyers and sellers using the auction systems for their needs have motivated many merchants who were selling their products via their own web service to integrate with an auction system and use it as their primary sales channel, resulting in such a large volume of trade as mentioned at the beginning of the article.

1.1 Auction-Related Fraud

The volume of transactions made via internet auctions make it a valuable target for criminals and abusers. According to [9], auction frauds can be split into three major categories:

- *Pre-auction fraud*, which includes misrepresentation, the sale of illegal goods or triangulation. The former two are not specific to auctions or e-commerce

in general, while the later (triangulation) is the sale of goods purchased with stolen credit card for cash - leaving the fraudster with cash and transferring the risk of seizure to the recipient [9].

- *In-auction fraud*, which is used to disrupt competitors' sales (e.g. by placing a high bid via a fraudulent account with no intention of buying the item, or by inflating the price by bidding on one's own items.
- *Post-auction fraud*, consisting mainly of non-delivery of the purchased item, the delivery of a broken or inferior item or stacking the buyer with additional fees.

More details about auction fraud can be found in [11] and [9]. However, of all the possible options, the most profitable to a fraudster are the ones which include upfront payment and non-delivery of the item, or the delivery of an item inferior to the offered one. Unfortunately for a fraudster, the reputation system does not allow this scenario to be exploited for long, as negative feedback from the buyers will soon warn other users and effectively prevent the fraudster from using his/her account with the auction system for this purpose. On the other hand, having access to an account with a high reputation allows for a larger numbers of buyers to be attracted to the fraudster's offer, allowing him/her to gather more money before negative feedback starts pouring in. Developing the means to abuse or circumvent the reputation system is therefore vital to a fraudster. It is a broad subject discussed e.g. in [18], but can be narrowed down to two most often used methods:

- *Building up a fraudulent reputation*, often utilizing a 'Sybil Attack' [2] where positive feedback for a specific account is generated via dummy accounts controlled by the fraudster
- *Gaining access to a legitimate account*, and exploiting it for own purposes (e.g. fraudulent offers/sales), leaving the original account owner with unhappy customers and, potentially, a legal struggle.

Of these two methods, the first one is more deterministic, although it requires a certain amount of effort and time to reach the stage when the fraudster can execute his/her schema, after which the account is basically unusable and a new one needs to be prepared. The second method is less reliable, as it depends on certain circumstances, often outside the fraudster's control (e.g. carelessness of a certain user or the auction system operator), but provides the fraudster with an account that can be utilized on the spot and with possible less risk as the original user will be the primary target of the claims.

1.2 Existing Fraud Prevention and Detection Techniques

It was not long after eBay launched that fraudsters noticed the new options it provided. An initial analysis of auction fraud and its prevention appeared as early as 2000 [3]. By 2006, online auction fraud was the most often reported offence in Australia, according to a government report [21].

So far, most of the research focus has been applied to identifying the fraudulent accounts that were used to build up a reputation score based on the distribution of accumulated feedback in time [4], decision trees [5] or belief propagation

and Markov random fields [25]. Also, there are proposals to utilize non-technical methods (i.e. social groups and their collective expertise) to combat some specific forms of auction fraud [7]. There is, however, substantially less interest in identifying hacked or stolen auction accounts. Although the issue (also named an 'identity theft') is very important to financial industry, as outlined in [17],[23], there is little specific research related to auction accounts, even though, as will be demonstrated in this article, this specific environment provides opportunities to utilize various techniques based on specifically available data.

2 An Overview of the Proposed Method

In this paper we propose a multi-model approach to detecting anomalies in the behaviour of sellers participating in the internet auctions. For each seller a different behaviour model is created, which is next constantly matched against the current profile (offers and transactions performed). The model consists of a number of features and procedures which are used to evaluate the users' behaviour.

The primary task of the model is to assign a probability value to the current behaviour of the seller. This probability value reflects the probability of the occurrence of the given feature value with regards to an established seller profile. The assumption is that feature values with a sufficiently low probability indicate potentially abnormal behaviour, which in turn may be the result of an account hijacking by a malicious individual. Based on the model outputs, the user's behaviour may be reported as abnormal. This decision is reached by calculating a number of anomaly scores. The current user's behaviour is reported as anomalous if at least one of these anomaly scores is above the corresponding detection threshold. This approach shares some concepts with intrusion detection systems (IDS) [16], however, it operates on different types of data and behaviour models, as IDS operates on the network traffic level - detecting anomalies in network packets.

Similar multi-model approaches were successfully used for detecting potential attacks on web applications [15], [13]. Sample models of the seller's behaviour are described in the following section.

The real data about users' activities presented in this article have been gathered by the authors by monitoring Polish largest auction service (allegro.pl). This service consistently hosts over 1 million active auctions at any given time, and has an important advantage over eBay from the research point of view, as it allows for the retrieval of users' history (past auctions).

3 EWMA of the User's Activity

The proposed model is based on measuring the total number of items offered for auction in all categories on any given day. To restrict the model sensitivity to temporary fluctuations in the number of items offered daily, the model utilizes an exponentially weighted moving average. This average ($S(t)$) is calculated according to a recursive formula:

$$S(t) = \begin{cases} \alpha \cdot y(t-1) + (1-\alpha) \cdot S(t-1) & \text{if } t > 2 \\ y(1) & \text{if } t = 2 \end{cases} \quad (1)$$

Where:

- t is discrete time (the number of the day), in which we calculate the average number of auctions; the mean is calculated from the initial time $t = 2$
- $y(t)$ is users' activity (e.g. the number of items offered by a seller) on the day t
- α is the smoothing constant (filter factor)

Additionally, the variance is calculated recursively:

$$V(t) = \alpha \cdot (y(t) - S(t-1))^2 + (1-\alpha) \cdot V(t-1) \quad (2)$$

Where:

- $V(t)$ is the variance at the moment t

Applying Chebyshev's inequality

$$P(|x - E(x)| > \varepsilon) < \frac{V(x)}{\varepsilon^2} \quad (3)$$

and substituting $E(x) = S$ and $\varepsilon = |y(t) - S|$, we obtain:

$$P(|y - S| > |y(t) - S|) < \frac{V(t)}{|y(t) - S|^2} \equiv P(y(t)) \quad (4)$$

The Chebyshev's inequality imposes an upper bound on the probability, i.e. that the difference between the value of a random variable x and $E(x)$ exceeds a certain threshold ε , for an arbitrary distribution with variance $V(x)$ and mean $E(x)$. The inequality is very useful because it can be applied to various arbitrary distributions with finite variance.

The formula (4) calculates the probability value $P(y(t))$ if the amount of user's activity (e.g. the number of items put up for auctions) at any given time $y(t)$ exceeds the current value of $S(t)$. If the number of items is smaller then or equal to $S(t)$, it is assumed that $P(y(t)) = 1$. The value of $P(y(t))$ is the value returned by this model.

Figure 1 illustrates a typical scenario, with varying but consistent user's activity over time. Although the activity is changing substantially, the value of the $dV(t)/dt$ function does not reach significant levels.

In another scenario, illustrated in Fig. 2 the user's activity includes a significant peak at a certain time (around 40th day). This is promptly signalled as a suspicious activity by the change in variation exceeding the value of 10. The proposed model proves also its usefulness in Fig. 3, when an activity of a specific user is illustrated. This user apparently puts up items for sale in weekly 'batches'. As can be seen in this figure, the model does not alert of a suspicious activity in this case, which is a desired outcome, as such behaviour is consistent and unsurprising.

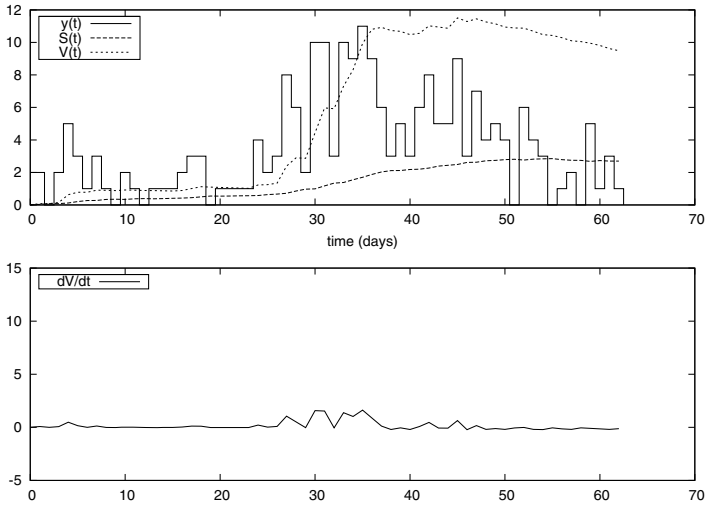


Fig. 1. Non suspicious activity of a selected user. The values of moving average, variance and variance's derivative are presented. The values calculated for $(\alpha = 0.02)$.

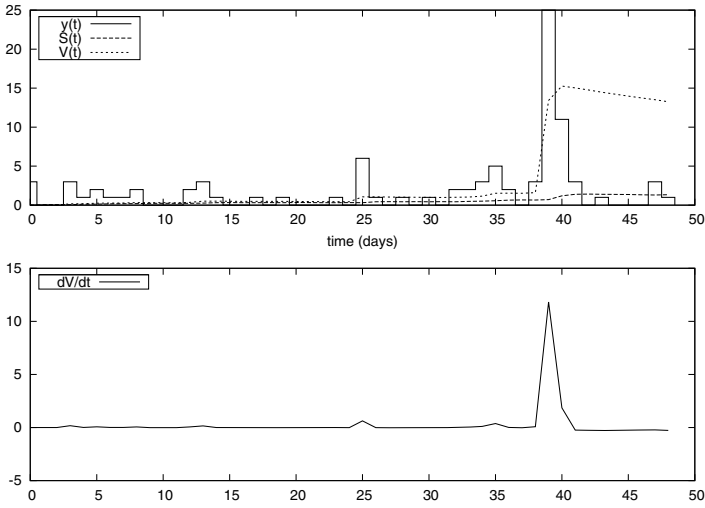


Fig. 2. Example of suspicious activity $(\alpha = 0.02)$

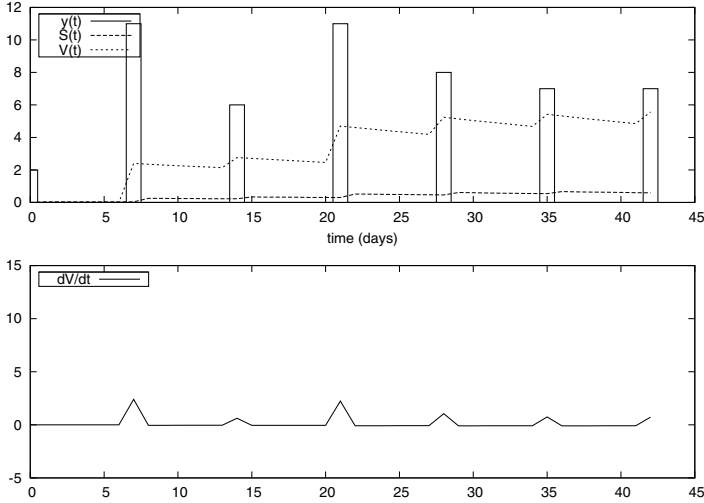


Fig. 3. Insensitivity of the detection to periodical activity ($\alpha = 0.02$)

4 'Thematic' Category Clusters

Although the proposed model of user's activity performs up to the expectations, it is usually better to have multiple detection systems (at least two) for the confirmation of a suspicious case.

Another criterion of the suspicious seller's behaviour (which might indicate a takeover of an account) is a sudden change of the types of items provided by the seller. Since all auction services allow the sellers to assign the offered item with a category (from a provided list), a sudden change in the number of items offered (or transactions) per categories is a possible warning sign. For example, a user who so far has sold items mostly in the categories *for children* \rightarrow *toys* and *books* \rightarrow *comics* suddenly starts to sell in the category *jewellery for men* and *jewellery for women*.

In order to detect such changes in the profile of categories for a given seller, it is necessary to cluster all categories of an auction service into thematic groups. By 'thematic' we mean groups that are likely to share similar items across several categories. Such clusters are likely to group together the already mentioned *jewellery for men* and *jewellery for women* as well as e.g. *books* \rightarrow *guidebooks* and *car* \rightarrow *manuals*.

This clusterization allows to build and observe sellers' activity profiles within given thematic categories. Unfortunately, the hierarchy of categories offered by auction services often does not suit this purpose, as similar items can be offered in distant categories (according to the hierarchy tree).

In order to create useful clusters of categories, they were grouped on the basis of similarity of the names of items present. This is done as follows:

In given time intervals (one month in the existing implementation), the names of all objects offered in all categories are acquired. For each category pair the probability is calculated using the formula:

$$s(c_a, c_b) = \frac{\sum_{i=1}^n \max_{1 \leq j \leq m} \tilde{f}(p_{ca}(i), p_{cb}(j))}{n} \quad (5)$$

where

- n the number of auctions in the category c_a
- m the number of auctions in the category c_b
- $p_{ca}(i)$ - the name of the object with the number and in the category c_a
- $p_{cb}(j)$ - the name of the object with the number and in the category c_b

next, the similarity factor is calculated:

$$\tilde{f}(p_{ca}(i), p_{cb}(j)) = \begin{cases} 0 & \text{if } f(p_{ca}(i), p_{cb}(j)) < 0.5 \\ f(p_{ca}(i), p_{cb}(j)) & \text{if } f(p_{ca}(i), p_{cb}(j)) \geq 0.5 \end{cases} \quad (6)$$

$$f(p_{ca}(i), p_{cb}(j)) = \frac{1 - L_{dist}(p_{ca}(i), p_{cb}(j))}{\max(|p_{ca}(i)|, |p_{cb}(j)|)} \quad (7)$$

where

- $L_{dist}(p_{ca}(i), p_{cb}(j))$ - the Levenshtein distance between name $p_{ca}(i)$ and $p_{cb}(j)$
- $|p_{ca}(i)|$ - size (number of characters) of name $p_{ca}(i)$
- $|p_{cb}(j)|$ - size (number of characters) of name $p_{cb}(j)$

The similarity \tilde{f} shown in equation (6), represents the percentage distance between names (i.e. the minimum number of edits needed to transform one name into another divided by the length of the longest name multiplied by 100%). If it exceeds 50%, the value of similarity \tilde{f} is assigned the value of 0 to limit the influence on the similarity of the category $s(c_a, c_b)$ of the objects significantly differing in names (the suggested cut off threshold at 50% is arbitrary, but has proven to be a reasonable value).

Before calculating the Levenshtein distance [14] L_{dist} between the names of the items, p_{ca} and p_{cb} are normalized:

- all 'marketing' marks used by sellers in order to attract buyers such as: '#', '!', '*', are removed
- white spaces and the following signs “,;.-” are concatenated to a single space
- all letters are transformed to lower case.

Such normalization of names is necessary to achieve a meaningful distance between the names, as sellers tend to utilize numerous ways of modifying the names in order to stand out with their offers. As can be observed, due to the

way of defining the similarity S between categories, $0 \leq s(c_a, c_b) \leq 1$ as well as self-similarity of categories $s(c_a, c_a) = 1$.

On the basis of the similarity s between categories the symmetrical similarity measure is defined as:

$$s_{sym}(c_a, c_b) = \frac{s(c_a, c_b) + s(c_b, c_a)}{2} \tag{8}$$

On the basis of the symmetrical similarity measure s_{sym} , an undirected graph is built which represents the similarity between categories. In this graph the vertices represent given edges, and edges represent the similarity between given categories. The weight of the edges connecting vertices c_a and c_b equals $s_{sym}(c_a, c_b)$. If $s_{sym}(c_a, c_b) = 0$, the edge is discarded.

During the next step, the graph constructed undergoes a clusterization in order to group thematically similar categories together. The clusterization algorithm used is a recursive spectral algorithm described in [12]. This algorithm was chosen because of its many advantages, including its speed and the fact that it can be successfully applied in a variety of contexts [1], [8], [20], [22], [10], [24].

The specific algorithm used in the reference implementation was based on [6] and is described in (Algorithm 1).

Algorithm 1. Clustering of the categories

Input: Matrix $n \times n$ containing weights of undirected weighted graph representing categories similarity

Output: A tree whose leaves are the row indexes of A representing clusters

1. Initialize
 - Let $R^2 \in \mathfrak{R}^{n \times n}$ be a diagonal matrix whose diagonal entries are the row sums of AA^T
 2. Compute Singular Vector
 - Compute the second largest right singular vector v' of the matrix $A^T R^{-1}$
 - Let $v = R^{-1}v'$
 3. Cut
 - Sort v coordinates so that $v_i \leq v_{i+1}$
 - Find the value t that minimizes the conductance of the cut:
 $(S, T) = (\{v_1, \dots, v_t\}, \{v_{t+1}, \dots, v_n\})$
 - Let A_S, A_T be the submatrices of A whose rows are those in S, T
 4. Normalize
 - Adjust the selfsimilarities
$$A_{ii}^2 := A_{ii}^2 + \begin{cases} \sum_{j \in T} A_{(i)} \cdot A_{(j)} & \text{if } i \in S \\ \sum_{j \in S} A_{(i)} \cdot A_{(j)} & \text{if } i \in T \end{cases}$$
 5. Recurse
 - Recurse steps 2-4 on the submatrices A_S and A_T
-

The conductance of a cut $(S, V \setminus S)$ is calculated as follows:

$$cond(S, V \setminus S) = \frac{d(S, V \setminus S)}{\min(d(S), d(V \setminus S))} \tag{9}$$

where

- $d(A, B) = \sum_{i \in A, j \in B} A_{(i)} \cdot A_{(j)}$
- $d(A) = d(A, V)$
- $A_{(i)}$ is i -th row vector in matrix A

The results of the clusterization can be seen in Fig. 4, which illustrates how all activities of two users really belong to one primary specific cluster of categories, with some marginal activity in other category clusters.

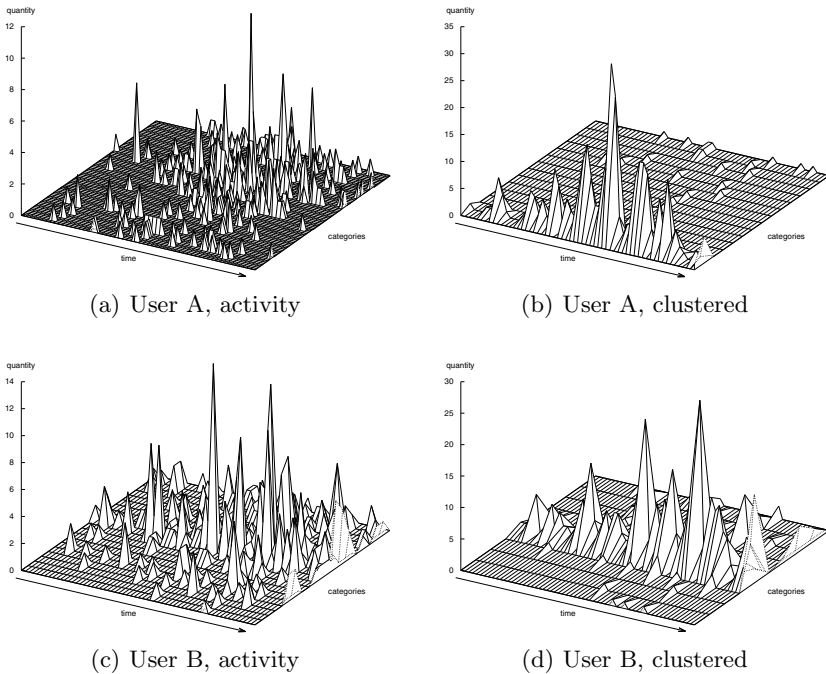


Fig. 4. Illustration of User’s activity (the quantity of daily transactions) for two different accounts. The graphs on the right illustrate activity aggregated into ‘thematic’ clusters.

5 Detecting Unusual Activities

After the clusterization into thematic category groups, the probability of a certain number of offers appearing in a given group on a given day is calculated. The probability is calculated in the same way as the EWMA model described above. The probability of correct (non anomalous) behaviour yielded by

this model $P(y(t))$ is described as the minimum of probabilities in particular clusters:

$$P(y(t)) = \min(P_c(y(t))) \quad (10)$$

where

- $c \in C$ (set of all clusters)

After calculating the probability of non anomalous behaviour at a given time t using particular models expressed as $P_m(t)$, it is possible to calculate the following parameters:

$$\begin{aligned} anomaly_score_w &= \sum_{m \in M} w_m \cdot (1 - P_m) \\ anomaly_score_{max} &= \max(1 - P_m) \end{aligned} \quad (11)$$

The first one represents a weighted sum of anomalous behaviour calculated by each model, while $(1 - P_m)$ denotes the probability of anomalous behaviour according to the model m , and w_m represents the weights associated with this model. The second parameter specifies a maximum probability of anomalous behaviour yielded by all models.

Finally, it is possible to select thresholds k_w and k_{max} respectively for calculated anomaly scores in such a way that after exceeding them, the system will report a possibility of unauthorized usage of the suspicious account. The thresholds needs to be adjusted manually in order to minimize the number of false positive alerts while preserving the sensitivity of the system to anomalous behaviour.

6 Conclusion

The models proposed in this paper for the assessment of the user's activity behaviour have proven very effective against the provided set of data. The data used for validating the models were gathered by daily retrieval of all the auctions from their web site for a period of one month. There were several millions of auctions retrieved during that time. Unfortunately, due to legal and privacy concerns, we were not able to receive data on real accounts taken over by criminals, so the model was validated with the data manually reviewed which deemed to be suspicious (e.g. Fig. 2). The clusterization of the categories has also proved to yield extraordinary results, with significant portion of users having most of their transactions in just a few (or even one) primary category groups. Interestingly, with the total number of groups equal to approximately a quarter of all categories, some groups consisted of over 200 categories, while the others were single-membered.

The most computationally expensive part of the proposed process is the grouping of categories, which can fortunately be done quite rarely (e.g. once a month) and off-line. Other algorithms are lightweight and can easily be utilized for a real-time monitoring on any scale of users. Implementing such solutions will not

eliminate the possibility of fraudulent use of a hijacked account, but will at least greatly limit the benefits, as an alert can be risen very quickly and the suspicious account suspended for evaluation. As has been mentioned before, auction fraud is a considerable aspect of public security, therefore, its mitigation is of interest to both auction service providers and security forces (e.g. police).

Although the proposed model proves to be effective, it can be further enhanced with other detection factors (e.g. the assessment of the value of items offered instead of their number). This may further improve its ability to distinguish anomalies in users' behaviour.

References

1. Alpert, C., Kahng, A., Yao, Z.: Spectral partitioning: the more eigenvectors the better. *Discrete Applied Mathematics* 90, 3–26 (1999)
2. Beranek, L.: Auditing Electronic Auctions Systems. *ISACA OnLine Journal* 4 (2010), <http://www.isaca.org/Journal/Past-Issues/2010/Volume-4/Pages/default.aspx>
3. Boyd, C., Mao, W.: Security Issues for Electronic Auctions. Technical Report, Hewlett Packard (2000)
4. Chang, J.S., Chang, W.H.: An Early Fraud Detection Mechanism for Online Auctions Based on Phased Modeling. In: *Proceedings of Joint Conferences on Pervasive Computing (JCPC)*, Taipei, pp. 743–748 (2009)
5. Chau, D., Faloutsos, C.: Fraud Detection in Electronic Auction. In: *Proceedings of EWMMF 2005: European Web Mining Forum*, Porto (2005)
6. Cheng, D., et al.: On a recursive spectral algorithm for clusterin from pairwise similarities. MIT LCS Technical Report MIT-LCS-TR-906 (2003)
7. Chua, C., Wareham, J.: Fighting Internet Auction Fraud: An assessment and proposal. *IEEE Computer* 37(10), 31–37 (2004)
8. Dhillon, I.: Co-clustering documents and words using bipartite spectral graph partitioning. In: *Knowledge Discovery and Data Mining*, pp. 269–274 (2001)
9. Dong, F., Shatz, S., Zu, H.: Combating Online in-Auction Fraud: Clues, Techniques and Challenges. *Computer Science Review* 3(4), 245–258 (2009)
10. Fowlkes, C., et al.: Spectral Grouping Using the Nyström Method. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 214–225 (2004)
11. Gavish, B., Tucci, C.: Reducing Internet Auction Fraud. *Communications of the ACM* 51(5), 89–97 (2008)
12. Kannan, R., Vempala, S., Vetta, A.: On clusterings: good, bad and spectral. In: *Proceedings of the 41st Annual Symposium on Foundations of Computer Science, California*, pp. 367–380. IEEE Computer Society (2000)
13. Kruegel, C., Vigna, G., Robertson, W.: A multi-model approach to the detection of web-based attacks. *Computer Networks* 48, 717–738 (2005)
14. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady* 10, 707–710 (1966)
15. Pałka, D., Zachara, M.: Learning Web Application Firewall - Benefits and Caveats. In: Tjoa, A.M., Quirchmayr, G., You, I., Xu, L., et al. (eds.) *ARES 2011*. LNCS, vol. 6908, pp. 295–308. Springer, Heidelberg (2011)
16. Pietro, R., Mancini, L. (eds.): *Intrusion Detection Systems*. Springer (2008) ISBN: 978-0-387-77265-3

17. Putting an End to Account-Hijacking Identity Theft. Federal Deposit Insurance Corporation (2004)
18. Reichling, F.: Effects of Reputation Mechanisms on Fraud Prevention in eBay Auctions. Thesis, Stanford University (2004)
19. Quaterly Retail E-commerce Sales (2009),
<http://www.census.gov/retail/mrts/www/data/pdf/09Q4.pdf>
20. Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(8), 888–905 (2000)
21. The risk of criminal exploitation of online auctions. Australian Institute of Criminology (2007)
22. Weiss, Y.: Segmentation using eigenvectors: a unifying view. In: *Proceedings of IEEE International Conference on Computer Vision*, pp. 975–982 (1999)
23. Wheeler, R., Aitken, S.: Multiple algorithms for fraud detection. *Knowledge-Based Systems* 13, 93–99 (2000)
24. Xiang, T., Gong, S.: Spectral clustering with eigenvector selection. *Pattern Recognition* 41(3), 1012–1029 (2008)
25. Zhang, B., Zhou, Y., Faloutsos, C.: Toward a Comprehensive Model in Internet Auction Fraud Detection. In: *Proceedings of Hawaii International Conference on System Sciences*, pp. 79–87. IEEE Computer Society (2008)