

Cluster-based Analysis and Recommendation of Sellers in Online Auctions

Mikołaj Morzy and Juliusz Jezierski

Institute of Computing Science
Poznań University of Technology, Piotrowo 2, 60-965 Poznań, Poland
{Mikolaj.Morzy, Juliusz.Jezierski}@put.poznan.pl

Abstract. The expansion of the share of online auctions in electronic trade causes exponential growth of theft and deception associated with this retail channel. Trustworthy reputation systems are a crucial factor in fighting dishonest and malicious users. Unfortunately, popular online auction sites use only simple reputation systems that are easy to deceive, thus offering users little protection against organized fraud. In this paper we present a new reputation measure that is based on the notion of the density of sellers. Our measure uses the topology of connections between sellers and buyers to derive knowledge about trustworthy sellers. We mine the data on past transactions to discover clusters of interconnected sellers, and for each seller we measure the density of the seller's neighborhood. We use discovered clusters both for scoring the reputation of individual sellers, and to assist buyers in informed decision making by generating automatic recommendations. We perform experiments on data acquired from a leading Polish provider of online auctions to examine the properties of discovered clusters. The results of conducted experiments validate the assumptions behind the density reputation measure and provide an interesting insight into clusters of dense sellers.

1 Introduction

By the year 2006, 63% of online population will have engaged in e-commerce activities. It is estimated that already in 2006 e-commerce transactions will account for 18% of all global sales. Meanwhile, auctions are triumphantly coming back in the form of customer-to-customer (C2C for short) e-commerce model. Today, more than 250 online auction sites enable customers to trade all possible types of goods using a variety of bidding protocols. More than 1.3 millions of transactions are committed daily on online auction sites. eBay, the global leader in the online auction market, reports 95 millions of registered users and 5 millions of transactions each week. At any given point in time there are approximately 12 millions of items posted on eBay. Examination of the latest financial data published by eBay reveals an astonishing development: for the second quarter of 2005 eBay reported net revenues of \$1.09 billion (40% increase year on year), operating income of \$380 million (49% increase year on year), and net income of \$290 million (53% increase year on year).

Huge success of online auctions can be attributed to many reasons. Bidders are not constrained by time, bids are placed 24/7 and potential users are given enough time to search and bid for interesting items. The Internet removes geographical constraints on users as they do not have to physically attend an auction. Large number of sellers and buyers reduces selling costs and influences prices of offered goods. Last but not least, many users describe their bidding experiences as similar to gambling. Apart from offering new and unprecedented possibilities, online auctions provide opportunities for dishonest participants to commit fraud [13]. The lack of physical contact between involved parties decreases the degree of trust exposed by users. According to a recent Eurobarometer poll, 73% of customers who do not participate in e-commerce, refrain from doing so motivated by concerns about the security of payment, delivery issues, and warranty terms. This fear is caused mainly by the growing number of complaints regarding online auctions. American Federal Trade Commission reports that 48% of all complaints concerning e-commerce involved fraud committed in online auctions, with the total loss of \$437 million in one year. National Consumers League reveals that 63% of complaints about Internet fraud concerned online auctions, with an average loss of \$478 per person. Online fraud can occur during bidding process and after bidding ends. Popular fraudulent practices include bid shielding and bid shilling. Bid shielding consists in providing artificially high bid for an item, thus discouraging other bidders from competing for an item. At the last moment, the shielder withdraws the bid, so the winner of an auction becomes the second highest bid cooperating with the shielder. Bid shilling consists in using a false bidder identity to drive up the price of an item on behalf of the seller. After the bidding process is over, potential fraud includes refraining from paying (bidder) and sending no merchandise or sending merchandise of lower quality and inconsistent with the offer (seller). These types of fraud are dangerous from the economical point of view, because they undermine the trust that users develop toward the online auction site.

One of the mechanisms to build trust between anonymous participants of online auctions are reputation systems [11]. Reputation is perceived by auction participants as a fundamental issue in developing a successful customer-to-customer relationship [10]. Furthermore, reputation of sellers has an economically observable and statistically significant effect on price of items [5]. Unfortunately, simple reputation systems used by online auction sites do not protect participants from malicious users. Typically, the reputation of a participant is measured by the number of committed auctions, where each auction is judged by the second party as “positive”, “neutral”, or “negative”. Such simple schema is both unreliable and fraud-prone, because dishonest users can easily deceive the system into assigning unfairly high reputation ratings. A seller can create a set of virtual bidders who will “win” seller’s auctions and provide the seller with additional positive feedback points. This technique is known as “ballot stuffing” and it biases the entire system, because unearned reputation allows the seller to obtain more bids and higher prices from other users [7, 12]. In order to better disguise this fraudulent practice, a seller could create a network of auctions be-

tween virtual bidders, turning them into a clique and inflating their reputation. Another possibility is to use virtual bidders to provide artificially negative feedbacks to seller’s competitors. This technique is referred to as “bad-mouthing”. Bad-mouthing is more difficult to implement, because it requires to actually win a competitor’s auction. Nevertheless, if the gain of driving a competitor out of the market exceeds the investment cost, bad-mouthing can be beneficial. One thing that should be stressed is the fact, that sellers and buyers are exposed to different types of risk. Sellers can postpone the shipment of an item until the payment is delivered, so the sellers are not threatened financially. On the other hand, buyers pay before receiving an item, unless using a trusted third-party, such as PayPal. The reputation of buyers has little importance for sellers, whereas the reputation of sellers is of crucial importance to buyers, who have to decide upon participating in an auction solely based on seller’s reputation.

In this paper we introduce a new measure of reputation of sellers in online auctions. We draw inspiration from social network analysis. We mine the topology of links between auction participants to discover clusters of densely connected sellers. We evaluate the usefulness of discovered clusters in assessing the reputation of sellers and in providing automatic recommendations based on discovered clusters. Our original contribution includes the definition of the density reputation measure, the concept of using dense clusters for generating automatic recommendations, and the experimental evaluation of the proposed solution. The paper is organized as follows. In Sect. 2 we present the related work on the subject. Section 3 introduces the density reputation measure and presents the idea of clusters of densely connected sellers. The procedure for automatically generating valid recommendations based on discovered clusters is also explained. The properties of the new density measure are examined using thorough experiments, and the results of the experiments are presented in Sect. 4. We conclude the paper in Sect. 5 with a summary of the future work agenda.

2 Related Work

An anonymous, heterogeneous, and geographically distributed environment for commercial transactions requires an efficient mechanism for building trust between participants. Reputation systems [11] allow users to develop long-term business relationships and receive financial benefit for their past honest behavior. Most auction sites use the reputation system developed by eBay, where credibility is expressed as the number of positive feedbacks minus the number of negative feedbacks received by a user [5, 10]. This simple mechanism suffers from several deficiencies, as pointed out in [6]. Feedbacks issued by users are subjective, lack transactional and social context, and contain highly asymmetric information. Neutral feedbacks are very rare, the spectrum for positive feedbacks is very broad, and negative feedbacks occur only when the quality of service becomes unacceptable, otherwise users refrain from posting a negative feedback in the fear of retaliation.

In recent years several new solutions have been proposed that aim at overcoming at least some of the deficiencies of feedback-based models. An interesting proposal was formulated in [1] where the authors develop a complaint-only trust model. Although originally developed for peer-to-peer environment, this highly decentralized model can be successfully used in online auctions. Another model originating from peer-to-peer environment is PeerTrust [14]. PeerTrust is a complex model consisting of many parameters, such as feedback in terms of satisfaction, number of transactions, credibility of feedback, transaction context, and community context. Method presented in [9] and further investigated in [8] does not use feedbacks to compute the reputation of participants. Instead, it uses a recursive definition of credibility and performs data mining to discover credibility estimation for each participant. A solution presented in [3] tries to prune false feedbacks and accepts only feedbacks that are consistent with the majority of feedbacks received by a given user. The need for a trusted third party is advocated in [2]. The authors propose to introduce a trusted judge that could authorize, identify, and manage the reputation of auction participants. An efficient method for assessing the level of trust between any two individuals based on a small amount of explicit trust/distrust statements per individual is presented in [4]. An interesting comparison of typical fraudulent behavior in online auctions with the abuse of customers by pay-per-call industry in the 1990s is presented in [13]. In the opinion of the author, the ability of online auction business to self-regulate is limited and not adequate to solve the problem, so legislation must be introduced to guarantee sufficient customer protection.

3 Density Reputation Measure

The main drawback of all feedback-based reputation systems is the fact that the reputation estimation for a given user is strongly influenced by the reputation of users directly involved in auctions with the user. This allows dishonest participants to artificially inflate their reputation estimates. Therefore, we propose a new reputation measure for sellers. Our density reputation measure computes the reputation of a given seller based on the reputation of other “similar” sellers.

Given a set of sellers $S = \{s_1, s_2, \dots, s_m\}$. Two sellers s_i and s_j are linked if there are at least *min_buyers* who committed an auction with sellers s_i and s_j , and the closing price for each auction was at least *min_value*. The number of such buyers is called the *strength* of the link and is denoted by $link(s_i, s_j)$. The *neighborhood* $N(s_i)$ of a seller s_i consists of sellers $\{s_j\}$, such that the seller s_i is linked with s_j , given user-defined thresholds *min_buyers* and *min_value*, $N(s_i) = \{s_j \in S : link(s_i, s_j) > 0\}$. The cardinality of the neighborhood $N(s_i)$ is called the *density* of the neighborhood, $density(s_i) = |N(s_i)|$. The rationale behind user-defined thresholds is the following: *min_buyers* selects sellers with significant number of sales, and *min_value* prunes low-value transactions. The density measure can be interpreted as follows: a buyer who buys from sellers s_i and s_j acknowledges the quality of both sellers. Unexperienced buyers are unlikely to link many sellers, these are rather experienced buyers who are used

to link sellers. In this way the density measure discards unreliable information from unexperienced buyers. The fact that two sellers are linked indicates that either they trade similar and popular goods (such as music or books), or that their offer is complementary (like bicycles and bicycle add-ons). Obviously, a link between two sellers may be coincidental and may not bear any useful information. Nevertheless, high density of a seller is a good indicator of seller's trustworthiness. Another important issue is the type of a cluster to which a seller is linked. Density reputation measure discovers natural groupings of sellers around product categories, thus allowing to automatically generate meaningful recommendations.

The density reputation measure does not consider the strength of the link between any two sellers, only the density of a given seller's neighborhood. In order to distinguish between strongly and weakly linked sellers we introduce another reputation measure, called score, defined as

$$score(s_i) = \sum_{s_j \in N(s_i)} density(s_j) * \log_{min_buyers} link(s_i, s_j)$$

The score measure uses the density of each seller in the neighborhood of the current seller and multiplies it by the strength of the link between the two sellers. The logarithm is used to reduce the impact of very strong links between sellers.

The density reputation measure is very resistant to fraud and manipulation. Let us consider a malicious seller trying to enter the cluster of reliable sellers. Linking to a single trustworthy seller requires to create *min_buyers* and investing at least *min_buyers*min_value* in winning auctions of a trustworthy seller. Still, this links only to a single seller and places the cheater in the outskirts of the cluster. In order to receive higher density the cheater has to repeat this procedure several times. We attribute this feature of the density reputation measure to the fact that it uses other sellers to rate a current seller, rather than using information from buyers. We believe that it is much more difficult for cheaters to manipulate other sellers than to create virtual bidders and use them to inflate cheater's reputation.

The density reputation measure is used to provide users with automatic recommendations. When opening a page containing a given item, a user is presented with a set of top *n* dense sellers, who belong to the same cluster as the seller of the given item. Let *R* denote a set of target *n* sellers. Let *d(s_i, s_j)* denote the distance between the sellers *s_i* and *s_j* defined as the length of the shortest path connecting sellers *s_i* and *s_j* in the graph. The group density of the set *R* of sellers is defined as

$$density(R) = \frac{\sum_{s_r \in R} density(s_r)}{\sum_{(s_p, s_q) \in R \times R} d(s_p, s_q)}$$

When displaying top *n* sellers as a recommendation for currently selected seller *s_i* we are trying to find the set *R(s_i)* of sellers who are characterized by high group density and who are close to a given seller *s_i*,

$$R(s_i) = \arg \max_R \frac{\text{density}(R)}{\sum_{s_r \in R} d(s_i, s_r)}$$

Therefore, using the recommendation system the user gains access to auctions of reliable sellers who trade goods that are similar to the searched item. Most notably, the recommendation depends on neither textual descriptions provided by sellers nor category assignments of items. This is an important feature of the recommendation system, because it allows to generate description-independent and taxonomy-independent suggestions.

4 Experimental Results

The data have been acquired from www.allegro.pl, Polish leader of online auctions. The dataset consists of 440 000 participants, 400 000 auctions, and 1 400 000 bids. The number of participants is greater than the number of auctions, because for each participant their highest bid is stored, whether it was the winning bid or not. Therefore, we have data on some participants who did not win any auction. Analyzed dataset is a small subset of the original data and it has been created using the following procedure: 10 000 sellers have been selected, and for this seed set all their auctions from a period of six months and participants of these auctions have been collected. Analogously, 10 000 buyers have been selected randomly and a similar procedure has been applied to this seed set. Altogether, complete information on 20 000 participants was available. Data were stored and preprocessed using Oracle 10g database.




Fig. 1. Pairs and dense sellers (a)




Fig. 2. Pairs and dense sellers (b)

Figure 1 presents the number of linked pairs of sellers and the number of dense sellers when increasing the value of the *min_buyers* threshold. As can be seen, even for small threshold value the number of pairs and the number of dense sellers becomes manageable. Figure 2 presents analogous results for varying the values of the *min_price* threshold.




Fig. 3. Number of clusters




Fig. 4. Size of the biggest cluster

Figure 3 presents the changes in the number of discovered clusters when varying *min_price* and *min_buyers* thresholds. As can be seen, the *min_price* threshold has stronger impact on the number of discovered clusters, except when *min_buyers* = 1. The number of clusters is relatively small and for most combinations of both thresholds the space of sellers is dominated by a few clusters (usually with one main cluster being significantly bigger than the others). Figure 4 presents the size of the biggest discovered cluster for a given combination of *min_price* and *min_buyers* thresholds. When no thresholds are set, almost all sellers are assigned to a single cluster. Interestingly, this result suggests that Milgram’s concept of six degrees of separation applies also to the online auction environment (when discovering the borders of each cluster we never used more than five iterations to identify all members of the cluster). Another thing to notice is the fact, that the *min_price* threshold has very little impact on the size of the biggest cluster when more than two links are used to connect sellers. For realistic settings of both thresholds the size of the biggest cluster becomes relatively small, which makes this approach suitable for automatic recommendation generation. We believe that only the most trustworthy and reliable sellers are left in the clusters, thus making respective clusters a valid source of meaningful recommendation.

Two examples of density distribution are presented in Fig. 5 (no limits on *min_buyers* and *min_price*) and Fig. 6 (*min_buyers*=2 *min_price*=\$20). When no thresholds are defined, two clusters of sellers are visible. The majority of sellers are characterized by the density from the range $\langle 1, 500 \rangle$, but there is also a small group of very densely connected sellers, and their density is $\langle 3200, 3500 \rangle$. Average density is 1217 and 8493 sellers (85% of the entire population) turned out to be dense. When thresholds are set, the average density drops to 5.9 and the number of dense sellers is 885 (8.8% of the entire population). One might argue that the *min_price* threshold is set too prohibitively, but the average price of items in the mined dataset is close to \$30, so we rather believe, that the algorithm really discovers the set of most important and credible sellers.

An interesting question is how does the new density measure relate to traditional reputation rating computed as the aggregation of positive and negative




Fig. 5. Density distribution (a)




Fig. 6. Density distribution (b)




Fig. 7. Rating distribution (a)




Fig. 8. Rating distribution (b)

feedbacks. The average rating distribution with respect to density is presented in Fig. 7 ($min_buyers=3$, $min_price=0$) and Fig. 8 ($min_buyers=2$, $min_price=\$30$). In general, higher density is a good indicator of high rating, but this relationship is not linear, specially when min_price threshold is set to prune out low value transactions. Fig. 9 ($min_buyers=2$, $min_price=0$) shows the projection of average rating vs density. Many high rated sellers have low density, which is even more visible when min_price is set. Sellers with high ratings are usually trading popular goods that are not expensive, so min_price threshold is punishing them. Similar analysis of average rating vs score is presented in Fig. 10 ($min_buyers=3$, $min_price=0$). These figures reveal a shift along the x-axis. This suggests that the sellers with low density and high rating have much higher average strength of the link than densely connected sellers.

The distribution of average price of offered items with respect to density is depicted in Fig. 11 ($min_buyers=3$, $min_price=0$) (on the figure prices are given in Polish zloty). Surprisingly, there is no clear evidence that higher density has any impact on the closing price reached by sellers. Finally, Fig. 12 ($min_buyers=4$, $min_price=0$) presents the distribution of average number of sales with respect to density. This time it is easily noticeable that highly dense sellers enjoy much larger volume of sales. This fact, more than the distribution of average price of




Fig. 9. Projection of density on rating




Fig. 10. Projection of score on rating

items, convinces us, that density is a good predictor of future performance of a participant of an online auction.




Fig. 11. Density and average price




Fig. 12. Density and average number of auctions

5 Conclusions

In this paper we have presented a new density reputation measure for sellers in online auctions. Our measure considers the network of interconnections between participants and mines the topology of the network to derive useful knowledge about users. Discovered clusters of densely connected sellers can be used as a predictive of future performance of a user, thus providing additional insight into the data. In addition, discovered clusters can be used to generate description-independent and taxonomy-independent recommendations for participants of online auctions. We believe that the density of a seller can be successfully used as an indicator of seller's reliability. Main advantages of the proposed solution include resistance to manipulation, ability to discover complex fraudulent activities, and practical usability proved by experiments. The support exhibited by

our commercial partners encourages us to follow the work in this area of research. Our future work agenda includes other models of user reputation, efficient use of negative and lacking feedbacks, and thorough investigation of the properties of clusters of sellers.

References

1. K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *CIKM '01: Proceedings of the tenth international conference on Information and knowledge management*, pages 310–317, New York, NY, USA, 2001. ACM Press.
2. S. Ba, A. B. Whinston, and H. Zhang. Building trust in online auction markets through an economic incentive mechanism. *Decision Support Systems*, 3(35):273–286, 2003.
3. M. Chen and J. P. Singh. Computing and using reputations for internet ratings. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 154–162, New York, NY, USA, 2001. ACM Press.
4. R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 403–412, New York, NY, USA, 2004. ACM Press.
5. D. Houser and J. Wooders. Reputation in auctions: Theory, and evidence from ebay. Technical report, University of Arizona, 2001.
6. R. A. Malaga. Web-based reputation management systems: Problems and suggested solutions. *Electronic Commerce Research*, 4(1), 2001.
7. M. I. Melnik and J. Alm. Does a seller's ecommerce reputation matter? evidence from ebay auctions. *The Journal of Industrial Economics*, L(3), September 2002.
8. M. Morzy. New algorithms for mining the reputation of participants of online auctions. In *WINE 2005, 1st Workshop on Internet and Network Economics, 15-17 December 2005, Hong Kong*. Springer Verlag, 2005.
9. M. Morzy, M. Wojciechowski, and M. Zakrzewicz. Intelligent reputation assessment for participants of web-based customer-to-customer auctions. In *AWIC 2005, 3rd International Atlantic Web Intelligence Conference, June 6-9 2005, Lodz, Poland*, pages 320–326. Springer Verlag, 2005.
10. P. Resnick and R. Zeckhauser. *Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System*, volume 11 of *Advances in Applied Microeconomics*, pages 127–157. Elsevier Science, 2002.
11. P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12), 2000.
12. P. Resnick, R. Zeckhauser, J. Swanson, and K. Lockwood. The value of reputation on ebay: A controlled experiment. Technical report, School of Information, University of Michigan, 2003.
13. J. M. Snyder. Online auction fraud: Are the auction houses doing all they should or could to stop online fraud? *Federal Communications Law Journal*, 52(2), 2000.
14. L. Xiong and L. Liu. A reputation-based trust model for peer-to-peer ecommerce communities [extended abstract]. In *EC '03: Proceedings of the 4th ACM conference on Electronic commerce*, pages 228–229, New York, NY, USA, 2003. ACM Press.