

RESEARCH

Open Access



TRIVIA: visualizing reputation profiles to detect malicious sellers in electronic marketplaces

Johannes Säng^{*} and Günther Pernul

^{*}Correspondence:
johannes.saenger@wiwi.uni-
regensburg.de
Department of Information
Systems, University of Regensburg,
Universitätsstraße 31, 93047
Regensburg, Germany

Abstract

Reputation systems are an essential part of electronic marketplaces that provide a valuable method to identify honest sellers and punish malicious actors. Due to the continuous improvement of the computation models applied, advanced reputation systems have become non-transparent and incomprehensible to the end-user. As a consequence, users become skeptical and lose their trust toward the reputation system. In this work, we are taking a step to increase the transparency of reputation systems by means of providing interactive visual representations of seller reputation profiles. We thereto propose TRIVIA - a visual analytics tool to evaluate seller reputation. Besides enhancing transparency, our results show that through incorporating the visual-cognitive capabilities of a human analyst and the computing power of a machine in TRIVIA, malicious sellers can be reliably identified. In this way we provide a new perspective on how the problem of robustness could be addressed.

Keywords: Reputation systems, Robustness, Attacks, Visual analytics

Introduction

Trust is an important mechanism to reduce risk perception and has therefore been identified to be a key factor for the success of electronic marketplaces [1]. As such environments usually involve strangers whose trustworthiness is unknown, reputation systems have become an integral part providing a valuable method to identify honest sellers and punish malicious actors. Typically, transaction partners are encouraged to leave feedback (numerical ratings or textual reviews) after each transaction denoting their satisfaction. Reputation systems collect all evidence, aggregate the referrals and give an overview of past behavior in a reputation profile. In a recent study, Diekmann et al. [2] could show that sellers with better reputation have an increased number of sales and obtain higher prices. Thus, reputation systems not only support buyers in deciding whom to trust, but also encourage trustworthy participation.

However, these reputational incentives do not always lead to honest behavior. One reason in focus of this work are multiple vulnerabilities current reputation systems suffer from. Malicious sellers who exploit the weaknesses can perform effective cheating such as providing unfair ratings or taking advantage of missing context data. Considering these

security gaps, cheating is much more profitable than acting honest [3]. For that reason, the security and robustness of reputation system has become a strongly discussed topic that many works focus on [3–7].

The red queen and the increasing non-transparency

In 1973, Leigh van Valen proposed the Red Queen Hypothesis denoting the idea of a constant arms race between co-evolving species [8]. He argued that an organism must constantly adapt and evolve to survive in an ever changing environment. This phenomenon's name was derived from Lewis Carroll's "Through the Looking-Glass" [9] where the Red Queen tells Alice that "*It takes all the running you can do to keep in place!*".

This effect can also be noticed in IT security. Security systems in general or trust and reputation systems in specific must constantly adapt themselves to get at the threats arising from the changing environment. While most early reputation systems could be easily manipulated, current systems have been gradually improved to cope with a variety of attacks. As a consequence of this constant improvement, reputation systems (particularly in research) have become quite sophisticated. Especially statistical solutions and machine learning solutions are highly complex. These advanced computation methods, however, are non-transparent and incomprehensible to the end-user as a recent study could demonstrate [10]. A numerical value as only outcome of the reputation assessment does in fact not reveal many details about the content of the input data (which reviews were considered, which not? To what extent do they affect the final reputation value?). Particularly, if the computation process is not replicable, many users become skeptical. Higher transparency could therefore notably increase the user experience as well as the users' trust toward the reputation system.

Approach and methodology

In this work, we are taking a step to enhance transparency of reputation systems and involve the user in the computation process by providing an interactive visual representation of seller reputation profiles. For this reason, we make use of visual analytics [11]. We develop different techniques to visually detect malicious sellers. As the robustness against attacks plays an important role in the trust management research community, this work provides a promising perspective on an alternative approach of how this problem could be addressed.

The remainder of this paper is organized as follows: first, we describe the research background and the related work relevant with respect to our approach ("Section Background and related work"). We thereby identify the challenges and motivate our work. At the same time, the research gap and objectives are clarified. In the following "Section Preliminary considerations for designing TRIVIA", we discuss preliminary considerations for designing TRIVIA (Trust and Reputation Visual Analytics) - a visual analytics cockpit to depict seller reputation profiles. Here, we introduce a range of visualization and interaction techniques suitable for reputation data. Based on this, we propose the conceptual design of TRIVIA and describe how this concept was implemented in a software tool in "Section TRIVIA: a visual analytics tool to detect malicious sellers in electronic marketplaces". To demonstrate the proper functioning of TRIVIA, we provide three case studies in which the interactive visualization of a seller profile is used to detect malicious behavior ("Section Case studies"). We set up a test-bed and simulated different seller attacks.

Results show that using TRIVIA, all simulated attacks can be reliably and transparently identified. Finally, we discuss the benefits and shortcomings of our approach, sum up our findings and provide our plans for future work in “Section Discussion and conclusion”.

Background and related work

Online trust and reputation systems

The literature on online trust and reputation systems (TRS) has been growing rapidly within the last two decades. Since the early work of Stephen Marsh [12] on the formalization of trust as a computational concept and the paper of Resnick et al. [13] on the use of reputation systems to facilitate trust in internet transactions, lots of works have been published in this context.

Although the computation models have notably changed since, the generic process of reputation systems stayed the same. According to [14], it can be divided into three phases, namely “collection”, “computation” and “presentation” as depicted on Fig. 1. In the first phase, the reputation system gathers evidence about the past behavior of a trustee and prepares it for subsequent computing. While personal experience is the most reliable source of evidence, it is usually not sufficiently available. Thus, data from other actors needs to be collected. In the computation phase the collected evidence is filtered, weighted and aggregated to calculate a single or several reputation values. The computation is considered being the heart of the reputation system. The output values are finally presented in a reputation profile in the third phase.

Most of the work proposed in the last years focuses on the improvement of the computation phase. The primary objective is to increase the robustness and accuracy of existing approaches by introducing novel methods to filter, weight or aggregate referrals. Only little research has been carried out regarding the presentation of reputation profiles. The output of most reputation systems is still a numerical value. Obviously, a numerical value as sole output cannot convey any information about the input data (e.g. which reviews were used? How do single reviews affect the final value?), thus leading to a cognitive gap. Hammer et al. [10] conducted a user-centric study of reputation metrics in which they show that more than half of the participants criticized the lack of transparency. Marsh et al. [15] also criticized the resulting complexity in a recent article: “*There is a noticeable trend to more and more complex models, using deeper mathematical techniques and*

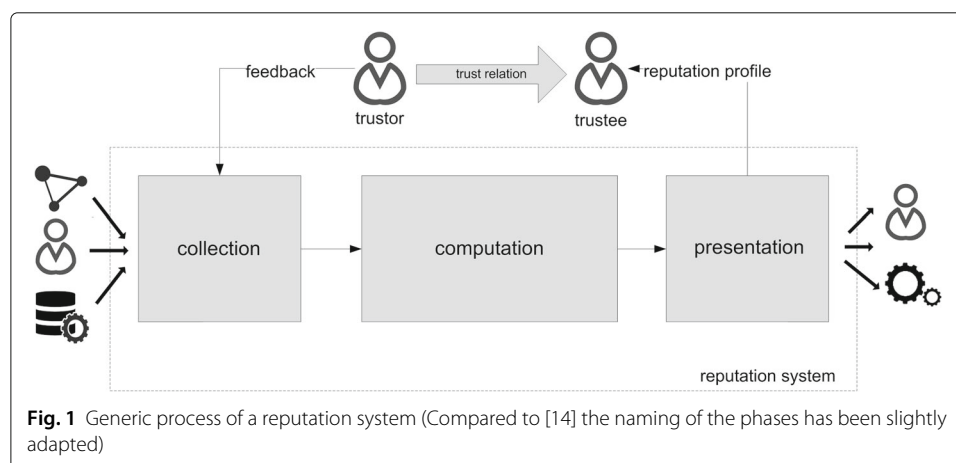


Fig. 1 Generic process of a reputation system (Compared to [14] the naming of the phases has been slightly adapted)

constructs [...] the increasing complexity of the models does not advance the field or its understanding or applicability in systems where it can be at its best”.

In this work, we try to take another path by not developing an even more sophisticated computation model, but by enhancing the presentation of reputation profiles. In order to increase the overall transparency we involve the user in the evaluation process and let the user detect coherence and reveal attacks. Through his continuously increasing experience and cognitive capabilities, a human analyst might have an advantage over analyses carried out by machines only. Marsh et al. [15] even argue that *“Humans, however, are much better at spotting anomalies in trust reasoning and reputation”*. Nevertheless, the support through automated computer analyses will be indispensable. For that purpose, we make use of a field that has emerged in the recent years called visual analytics.

Visual analytics

Visual Analytics (VA) is an interdisciplinary, fast-growing research area combining automated analysis techniques with interactive visualizations. Incorporating the visual-cognitive capabilities of a human analyst and the computing power of a machine, an effective understanding, reasoning and decision making is achieved [11].

Thus far, VA has not played a role in trust and reputation except two works that address the interactive visual detection of context-based attacks on reputation systems [16] as well as unfair ratings [17]. Approaches making use of visual representations only created static visualizations of relationship graphs or reputation scores [18, 19]. Further work regarding the visual communicating of information trustworthiness and quality was proposed by Nurse et al. [20].

Besides enhanced transparency, we think that advanced presentation might also increase the robustness by clearly depicting the underlying evidence and involving the user in the evaluation process. To elaborate on this idea, we need to take a look at attacks and robustness of reputation systems first.

Attacks and robustness of reputation systems

Common vulnerabilities of reputation systems as well as attacks to exploit these weaknesses have been identified and analyzed in multiple works. Most early papers on this topic introduced single attacks or vulnerabilities and at the same time proposed mechanisms to cope with those problems. Dellarocas [4], for instance, analyzed unfair ratings by buyers such as unfairly high ratings (“ballot stuffing”) and unfairly low ratings (“bad-mouthing”) as well as discriminatory seller behavior. He furthermore introduced techniques to reduce the effects of these attacks on reputation values and thus increased the robustness of reputation systems towards these attacks.

More detailed surveys of attacks and extensive studies on the robustness of common reputation systems are still young. One important work that analyzed the vulnerabilities of reputation systems in electronic marketplaces was proposed by Kerr and Cohen [21]. They identified several problems of the eBay reputation system such as value imbalance, reputation lag and re-entry. Based on these findings, they later derived a list of attacks on reputation systems in electronic marketplaces [3]. They could show that all these attacks and their combinations are very practical since every system tested was vulnerable to multiple attacks. In their scenario, cheating was much more profitable than honesty. Later, Jøsang and Goldbeck [5] revealed further seller attacks that had been described in

literature and summarized all misbehavior of advisors (the party providing a rating) to an unfair rating attack. Additionally, they discussed the practical robustness of current reputation systems arguing that it is also important to consider the ease and complexity of an attack's implementation since the identified attacks only give a theoretical perspective.

A recent work of Zhang et al. [6] elaborated on the term of unfair ratings as a catchall phrase for advisor misbehavior. They could identify seven classes of advisor attacks and evaluated the robustness of the existing trust models for handling unfair ratings. Equally to Kerr and Cohen, they could demonstrate that there is no single trust model that is robust against all the investigated attacks. Further recent studies did not only focus on the robustness but set up a detailed catalog of security problems of reputation systems [1, 7, 22, 23].

Based on the literature above, we identified the most suitable attacks for sellers in electronic markets. Table 1 gives an overview of selected seller attacks. Here, we excluded all kinds of unfair rating attacks (e.g. ballot stuffing, bad mouthing) which are common examples for advisor attacks. Collusion and Sybil Attack have been listed as potential seller attacks for the sake of completeness, although mainly being implemented as advisor

Table 1 Overview of selected seller attacks in electronic marketplaces

Attack	Description
Playbook	Playbooks are sequences of actions that allow the maximum outcome/profit for the player under specific auxiliary conditions. A simple example would be an on-off attack where a seller acts honestly to build high reputation by selling high quality products. After a particular time the seller changes his behavior and starts selling low quality products (under low production costs) and "milks" the high reputation. This sequence can be run through again and again. Overall, the seller stays unsuspecting [42]. This attack is particularly effective if the reputation metric "forgets" negative ratings.
Value Imbalance Attack	Sellers making use of the value imbalance problem build up high reputation selling cheap products while cheating on the expensive ones. In contrast to playbooks, this attack is non periodical but the malicious seller has to keep a balance between good and bad behavior to keep his reputation at a certain level for the whole time. Zhang et al. [43] extended the term of value imbalance to transaction context imbalance where not only the value but also the product type or the time can be used for asymmetrical allocation of good service.
Reputation Lag Attack	Usually, there is a time-lag between the advance payment and the delivery. Since referrals or ratings are normally made after the product is received, a malicious seller can exploit that time-frame by selling many low quality products before being rated badly for the first time.
Proliferation Attack	In a proliferation attack, the seller offers the same product from several accounts or channels to increase the probability that a buyer chooses his product instead of buying from a different seller offering the same product. Although often named as an attack on reputation systems, the "malicious" sellers do not really cheat on their buyers nor do they manipulate the reputation system. However, they get an advantage of their competitors. (The proliferation attack can be considered as a subset of the Sybil attack)
Re-entry Attack	Performing a re-entry attack (often referred to as whitewashing or newcomer attack), the malicious seller opens an account, cheats on the buyers and leaves the community to open a new account whenever his reputation is damaged. In re-entry attacks, the actor does never have to behave good. This type of attacks particularly exploits systems where a registration without any proof of identity is possible.
Collusion	If multiple actors coordinate their behavior to gain an advantage over the rest, this is called collusion. The purpose of collusions can be various, e.g. unfairly increase/decrease the reputation of an actor, discriminate groups or run coordinated playbooks.
Sybil Attack	In contrast to collusions where the accounts are created by multiple individuals, the Sybil attack is performed by one attacker who creates a number of accounts (pseudonyms). Due to the greater influence, the attacker can easily manipulate reputation values.

attacks. To get a holistic view on security problems, please refer to the aforementioned works.

Research gap: interactively visualizing seller reputation profiles to detect attacks

As the robustness of reputation systems is still a topical issue and current approaches lack a certain degree of transparency, the objectives of this work are as follows:

- We want to increase the transparency of reputation systems through depicting all input data in integrated visual representations of the reputation profile (in TRIVIA)
- We want to enhance the robustness of reputation systems through involving the user in the evaluation process (in TRIVIA)

Due to the variety of different application areas and attacks on reputation systems, we here focus on malicious sellers in eBay-like electronic marketplaces. The insights gained from this work might also be transferred to advisor attacks and other application areas.

Preliminary considerations for designing TRIVIA

Before describing the conceptual design of TRIVIA in the following section, we here give a basic overview of preliminary considerations we made. Adapting the visual analytics process, the conceptual design of a visual analytics software can be described within two essential blocks – namely “models” that are created using automated data analysis and “visualization & interaction techniques” provided for the user. The selection of appropriate models and visualizations, however, cannot be made *ex ante* but depends on the structure of the data provided. Thus, the first step designing TRIVIA was to range the raw data according to data types. At the same time, generic logical information blocks within the raw data had to be determined. The outcomes of these analyses represent the basis for the choice of suitable models and visualization techniques in the subsequent step.

Data classification

Depending on the application area, the input data of reputation systems can vary strongly. Figure 2 depicts a generic transaction and review process in electronic marketplaces. To derive information blocks that are representative for such environments, we analyzed this process by the example of an eBay-like reputation system.

Starting point is a transaction between a buyer and a seller (most systems do not allow feedback without a prior transaction). This transaction is carried out under a

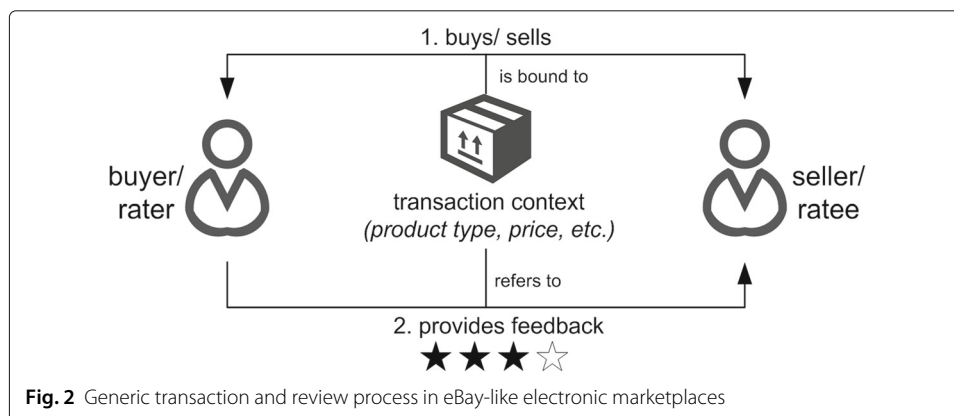


Fig. 2 Generic transaction and review process in eBay-like electronic marketplaces

transaction-specific context including, for instance, the product type, the price, the payment method or the transaction time. In the second step, the buyer fills the role of the rater or advisor who provides feedback toward the ratee (seller). This feedback refers to the same transaction context. For this scenario, we identified the three information blocks “*feedback*”, “*transaction context*” and “*actors*”.

As the selection of a suitable visualization technique depends on the data type of the input, the attributes of each information block need to be further analyzed. A meaningful classification of data types for visual analytics is still under discussion in literature. Based on the taxonomies of Shneiderman [24] and Keim [25], we distinguish the classes 1-dimensional, 2-dimensional, multi-dimensional, text/hypertext, networks, hierarchies, and software and algorithms. Table 2 depicts the resulting data classification.

Visualization techniques

The best choice of visualization techniques for specific data types is hard to measure and a quite subjective estimate. In this work, we considered some of the most suitable visualization techniques commonly used to illustrate data of the identified types. We here point out why each visualization is beneficial for reputation assessment and how they need to be adapted to perfectly fit our requirements.

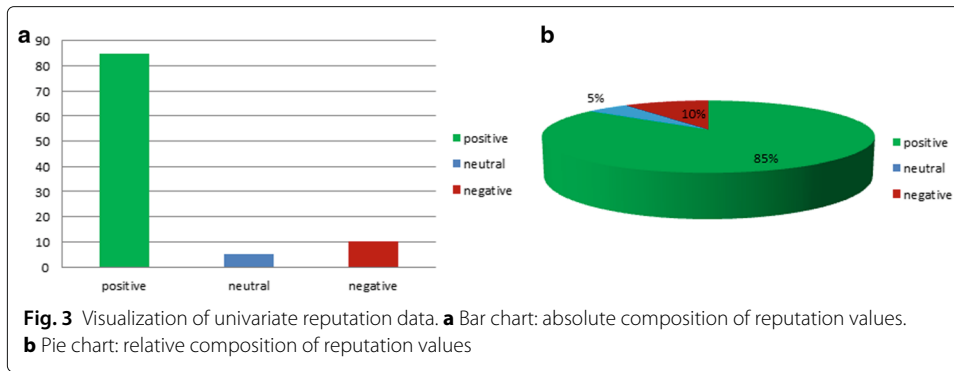
As we address the visual detection of seller attacks using TRIVIA, we assume all advisors to be honest. Thus, we only focused on the information blocks *feedback* and *transaction context* in this work.

Feedback: rating

The 1-dimensional rating can be considered being the most important information of a referral. Data is called 1-dimensional or univariate, if only one attribute changes in regard to one or several independent attributes [26]. There are many broadly known visualization techniques to depict univariate data. A very intuitive and easy to understand visualization that gives an overview of aggregated values is a simple *bar chart* revealing the absolute composition of positive (green), neutral (blue) and negative (red) ratings as depicted on Fig. 3a. A second well-known example to depict the relative composition is a *pie chart* as shown on Fig. 3b.

Table 2 Classification of common reputation data in electronic marketplaces

Information block	Data	Data type	Comment
Feedback		Multi-dimensional	<i>Typical feedback in electronic marketplaces involves both a rating and a textual review</i>
	Rating	1-dimensional	
	Review	Text	
	Time	1-dimensional	
Transaction context		Multi-dimensional	<i>Product type, price and time as context attributes are exemplary chosen</i>
	Product type	1-dimensional	
	Price	1-dimensional	
	Time	1-dimensional	
Actors		Network	<i>Through direct ratings, a uni-directed referral graph is created</i>
	Rater/advisor	Multi-dimensional	
	Ratee	Multi-dimensional	



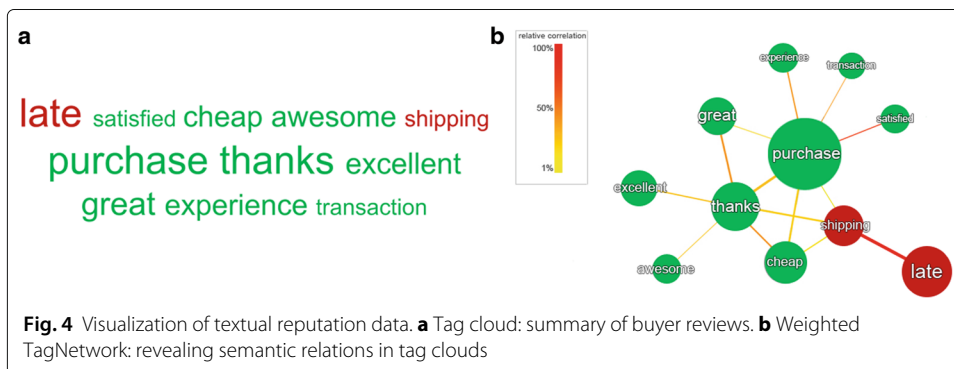
A visualization depicting the composition of input data can particularly be valuable as additional information to a numerical value since most metrics cannot convey concrete information about the input data. It can furthermore mitigate effects such as the “inaccurate equations” described by Tavakolifard and Almeroth [7].

Feedback: review

In eBay-like reputation systems, the rater is usually encouraged to leave a short text describing the buyer experience. These reviews are insightful for potential buyers as numerical ratings cannot reveal any details about the reasons for good or bad ratings. With a growing amount of feedback, however, a user may easily lose track of the bottom line. Using a visualization of textual comments in TRIVIA, vast numbers of textual reviews can be summarized in one picture.

Texts are a typical examples for unstructured data. A visualization technique applied to texts that has become very popular in the last decade is the *tag cloud*. The classical tag cloud is a text-based visual representation of a number of tags varying the font size in accordance with the tags’ importance [27]. To create a tag cloud, the text has to be parsed and filtered by auxiliary words, numbers and punctuation. Inflected words have furthermore to be reduced to their stem (stemming). Similar words are then grouped by the associated rating and counted. To additionally include the sentiment, we colored the tags according to the rating it is related to. The result is a colored tag cloud that gives an overview of buyer reviews as shown on Fig. 4a.

A shortcoming of the tag cloud, however, is that this visualization technique cannot expose relations between single tags. To close this gap, Hassan et al. [28] introduced the



weighted TagNetwork as a combination of the classical tag cloud and a weighted node-edge graph. Thereby, the size of a node grows according to the tags importance. By this means, significant tags can be identified immediately. An edge is created whenever two tags are used in the same review. While the width of an edge indicates the absolute frequency of an association, the color of an edge illustrates the relative frequency. In this way, the semantic context can easily be deduced. Figure 4b illustrates a weighted Tag-Network. Here it becomes obvious that “shipping” is strongly related to “late”. The color moreover indicated the sentiment the review is associated with. For further information on techniques for visual text analytics please refer to Alencar et al. [29].

Feedback: time

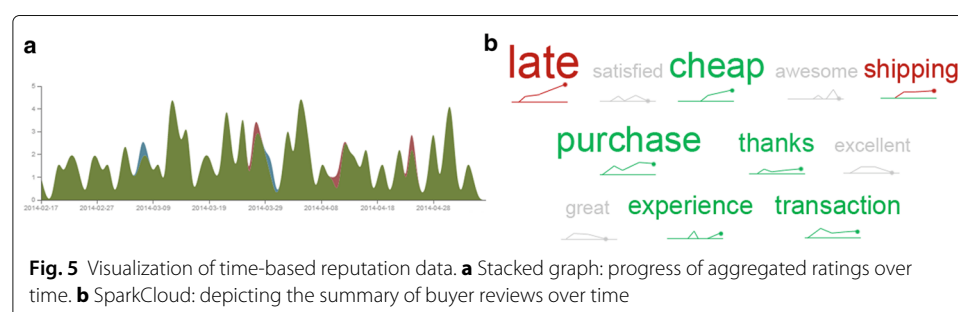
Due to the dynamic nature of trust, it has been widely recognized that time is a crucial factor for the importance of feedback. Old referrals might not be as relevant for reputation scoring as new referrals [30]. Though many computation models include a “forgetting factor”, it is still an open challenge to find the ideal trade-off between unlimited memory and a too small scope [7]. Both settings can lay the reputation system open to attacks. Visualizing the progress of feedback over time in TRIVIA, however, this problem can be alleviated through letting the user decide.

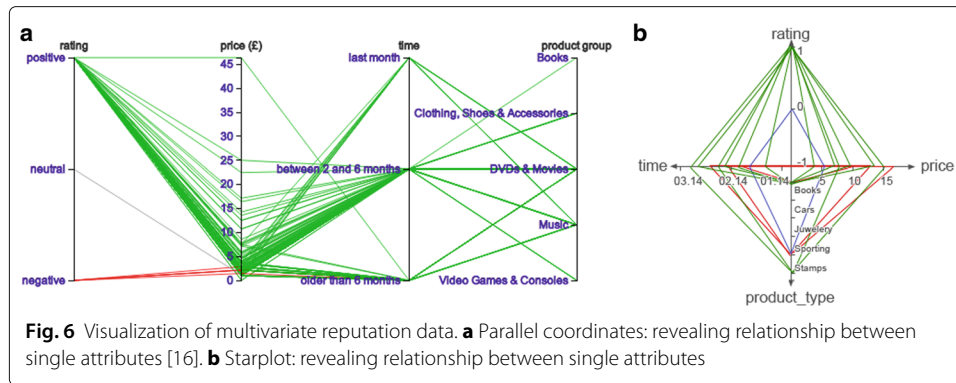
A common technique to depict multiple time-dependent variables over time is a *stacked graph* [31]. In stacked graphs, the development graphs of the single rating values $r = \{-1,0,1\}$ are placed on top of each other. In this manner, both the history of the entire activities and the composition of positive, neutral and negative ratings can be clearly depicted (Fig. 5a).

Besides the time-dependent visualization of ratings, changes in textual reviews should also be noticed. Referring to the tag cloud, a more advanced visualization technique that depicts the evolution of texts over the course of time is the *SparkCloud* introduced by Lee et al. [32] (Fig. 5b). In contrast to the tag cloud, the SparkCloud can reveal time-sensitive relations such as decreasing quality (e.g. the word “late” has lately become important). Here again, we colored the tags according to the rating the review is attributed to. To find more details on visualizing time-orientated data, please refer to Aigner et al. [31].

Transaction context: product type, price & time

The second information block we wanted to visualize in TRIVIA denotes the product type, price and time of the transaction the feedback refers to. Although transaction context has been recognized to play an important role for trust evaluation in literature [33, 34], only a very limited number of trust models formalizes the influence of context





for reputation assessment [35]. The non-consideration of the transaction context, however, might lead to the value imbalance problem, where a malicious seller can build a high reputation selling cheap products while cheating on the expensive ones [16]. To cope with this problem, an interactive visualization of transaction context can provide a remedy.

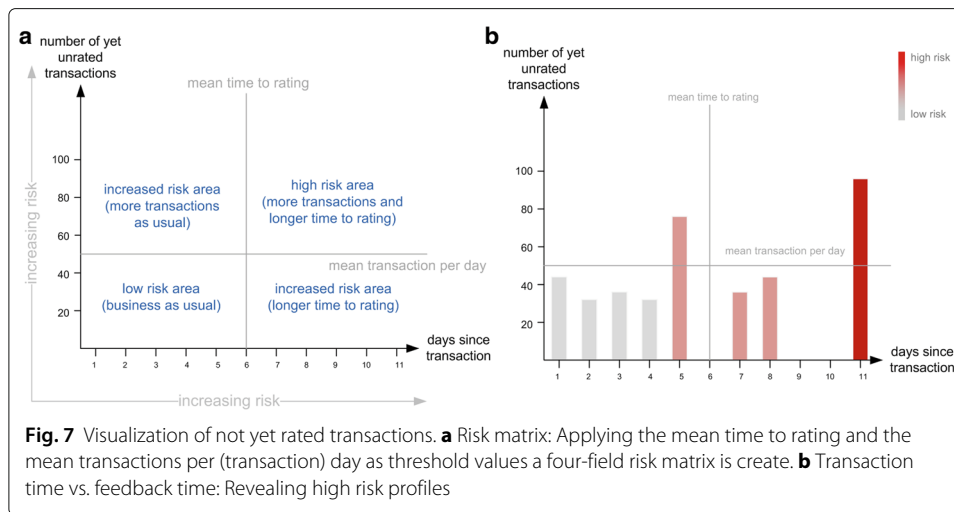
There is wide a range of visualization techniques suitable for multi-dimensional data sets. Here, we considered two techniques that have become quite popular due to their easy to understand and intuitive presentation. The first one - *parallel coordinates* - projects n dimensions to a 2-dimensional visualization space. Thereto, n axes are laid out in parallel side-by-side. A point in the n -dimensional space is then represented as a polyline with vertices on the parallel axes [36]. Parallel coordinates have proven themselves to be a perfect tool for revealing correlations between single attributes. Here, we provide an integrated view on ratings and its corresponding context as depicted on Fig. 6a. For a detailed description of parallel coordinates as a technique to depict transaction context in trust and reputation systems including case studies based on real-life data, please refer to [16]. A second visualization technique to depict multivariate data is the *star plot* (Fig. 6b). In contrast to parallel coordinates, the axes are laid out in a star pattern.

Comparing the two information blocks *feedback* and *transaction context*, it is apparent that both blocks contain a timestamp. The relation of these attributes led us to the idea of including not yet rated transactions.

Not yet rated transactions (reputation lag)

There is usually a time frame between the payment, the service provision and the providing of feedback. This “reputation lag” has been identified as a problem in many works (see “Reputation Lag Attack” - Table 1). Sellers could misuse this period by selling many low quality products or even not deliver before being rated badly for the first time. To cope with this issue, transactions that have not yet been rated could be involved in TRIVIA. For sake of simplicity, we assumed the delivery and the provision of feedback to happen at the same time. Since this is a quite specific application case, we here designed a chart for this particular scenario.

To get an idea of how long it usually takes to rate a transaction, the mean time to rating for the specific seller is of interest. Malicious sellers could try to increase this time-frame to maximize their payoff. Furthermore, it is helpful to know if there are many customers currently waiting for their delivery. This number of transactions per (transaction) day is the second lever to maximize the profit. Thus, we depict the number of transactions per



day on the number of days since the begin of the transaction for yet unrated transactions by means of a bar chart. The visualization is additionally extended by the mean time to rating and the mean transactions per (transaction) day leading to a four-field matrix (Fig. 7a). Figure 7b depicts the resulting bar chart.

Interaction techniques


Visual analytics is characterized through interactive user exploration of visual data representations. Interaction techniques therefore played a crucial role for the design of a TRIVIA. Heer and Shneidermann [37] classify interactive dynamics for visual analysis in “data and view specification”, “view manipulation” and “process and provenance”. We here considered a range of these basic interaction techniques and denote how they could be applied to evaluate reputation data.

Data and view specification

The first class “data and view specification” comprises basic interaction techniques such as *selecting* proper visualizations, *ordering* the output to expose patterns, *deriving* models from the source data or *filtering* all input data to focus on interesting aspects [37]. While this class provides quite fundamental tools, filtering seems to be particularly suitable for visually evaluating reputation data, since Sänger and Pernul [14] identified filtering to be one of the basic operations in the computation process of a common reputation system. If the user is particularly interested in feedback for one product type, for instance, he should be able to filter all other referrals. eBay, for example, has already implemented such an interactive filter to focus on positive, neutral or negative ratings in specific time-frames.

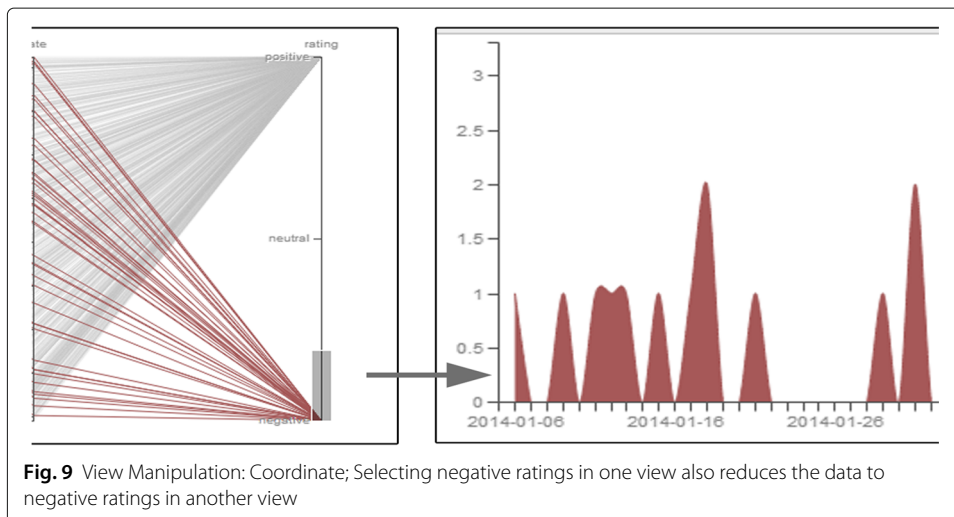
View manipulation

Once the input data and views have been specified, the human analyst should be able to “dive” into the data through manipulating TRIVIA’s visualizations. Here, common manipulation techniques can be subclassified to *selection*, *navigation*, *coordination* and *organization* [37]. *Selection* features support a user in pointing at specific parts of a visualization, highlighting them and in this way revealing relation between single attributes. These techniques are particularly helpful for the integrated visualization of ratings and corresponding context attributes to reveal relation (Fig. 8).



According to Shneiderman’s [24] often referred to navigation mantra “Overview first, zoom and filter, then details-on-demand,” a visual investigation should furthermore support the *navigation* through the data. This also holds good for reputation assessment. One example is the time-frame considered. Though the most actual feedback might be the most valuable, the whole history as well as single segments should be contemplated to avoid attacks based on time. A zooming feature could enable the user to analyze periods of varying length. *Coordination* as the third subclass of features addresses the integration of multiple views to reveal correlations between attributes and information blocks. One example for coordination is the interaction techniques “link & brush”. The idea of link & brush is to combine various visualizations to overcome the shortcomings of a single technique. Interactive changes made in one visualization (highlighting of specific parts or filtering) are then automatically reflected in the other visualizations [38] (Fig. 9). In this way, a human analyst can gain more information than evaluating the visualizations independently.

The fourth subclass *organization* concerns the challenge of managing different views. With trust being a quite subjective mechanism, it strongly depends on the user, which information of the provided feedback is considered to be the most important. Thus, an



analyst using TRIVIA should be able to organize the multiple views to meet his personal requirements. Focusing on textual reviews, for instance, all other views could be minimized.

Process and provenance

State-of-the-art visual analytics software is not only limited to visualizing and manipulating, it also supports an iterative process of knowledge discovery and interpretation. Thereto, findings can be *recorded*, *annotated* and *shared*. Furthermore, a human analyst should be *guided* through the application [37]. As these techniques mostly relate to the general analysis processes, we have not considered any of these for TRIVIA.

TRIVIA: a visual analytics tool to detect malicious sellers in electronic marketplaces

Based on the preliminary considerations made, we firstly propose the conceptual design of TRIVIA in this section. Secondly, we show how the concept was implemented in a software tool.

Conceptual design

According to the visual analytics process, the conceptual design of a visual analytics tool can be described within the two central building blocks “models” and “visualizations and interactions techniques”. Based on the data classification of a generic transaction and review process made above (Section “Data classification”), we selected five visualization techniques (and their inherent models) for TRIVIA, namely a *bar chart* to depict the distribution of positive, neutral and negative ratings, a *tag cloud* to sum up the content of all textual reviews, a *stacked graph* to illustrate the progression of all ratings over time, a *parallel coordinates* visualization to reveal correlation between ratings and context attributes as well as the visualization of *not yet rated transactions* to prevent the reputation lag. For the sake of completeness, we furthermore provide a list of all feedback to additionally increase the transparency. Besides, we added four interaction techniques to involve the user in the reputation assessment. As the parallel coordinates visualization combines data of both information blocks transaction context and feedback (rating), it constitutes the central part of our software. Figure 10 gives an overview of the schematic concept.

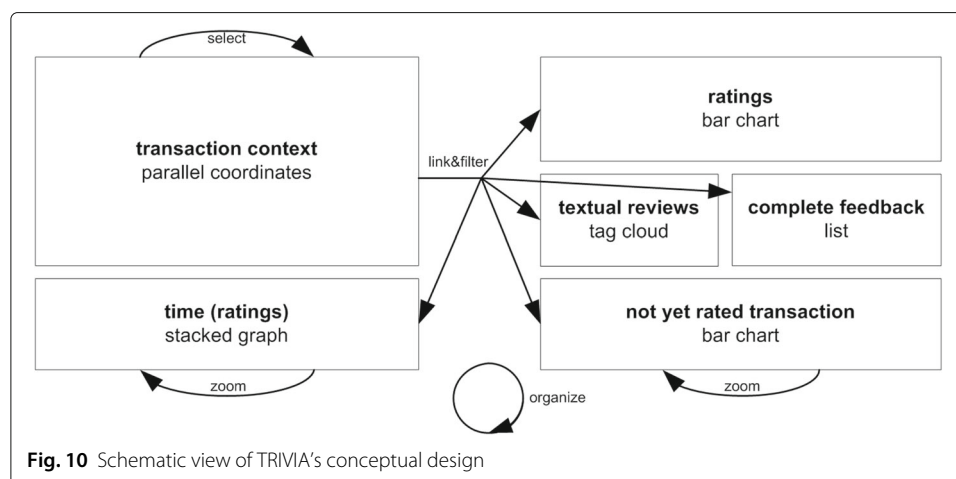


Fig. 10 Schematic view of TRIVIA's conceptual design

Thereby, the grey boxes illustrate the single visualization windows (as arranged in our software), whereas the arrows describe the interaction techniques applied. To connect all views to the parallel coordinates, we selected the interaction technique *link & filter*. We, moreover, allow to *select* and *zoom* into the visualizations as well as to *organize* the single windows.

Since our aim is to demonstrate the value added through a visual exploration of reputation data, the selected techniques can be considered to be an exemplary case. The choice of different and further techniques would also have been possible.

Implementation

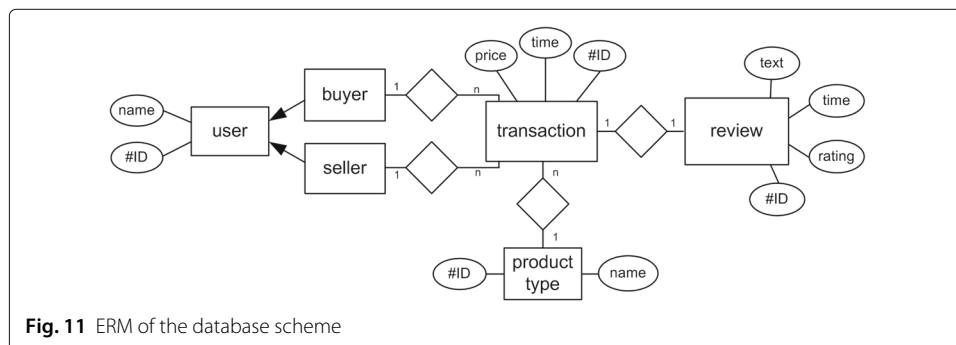
Based on the conceptual design, we implemented TRIVIA as a web-based application in a three-tier client-server-architecture¹. To simulate an electronic marketplace, we created an eBay-like scenario involving “users” who can take the role of a buyer or a seller, transactions that involve a product type, price and timestamp as well as reviews including a textual review, a rating and a timestamp. Figure 11 depicts the resulting database structure.

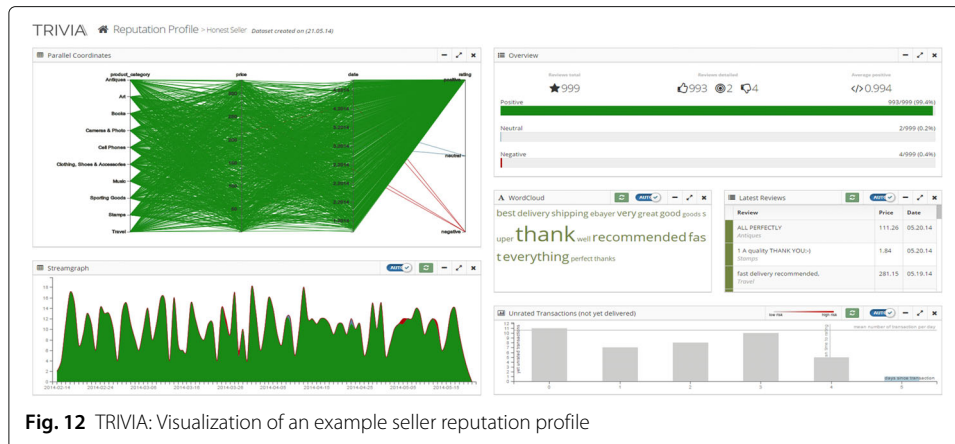
Implementing this scheme, we set up a MySQL-Database on server-side (persistent layer). The software logic was implemented in PHP on an Apache server (logic layer). On client-side (presentation layer) we employed the current web standards HTML5, JavaScript and CSS (Bootstrap). The visualizations were implemented using SVG and JavaScript with the d3.js-package² and the c3.js-package³. D3 (data-driven-documents) is a JavaScript library that allows to manipulate documents based on data. Applying d3, data are bound to the browser’s DOM enabling the user to instantly interact and manipulate the visualization. C3 builds on top of d3 providing various APIs and callbacks.

Figure 12 shows the GUI of TRIVIA visualizing an example seller profile. The arrangement and size of the single windows can be freely chosen to adapt the presentation of the reputation profile to personal preferences (organization). More detailed views of each window can be found in Section “Case studies”. Various interaction techniques furthermore allow to “dive” into the data. Figure 13, for instance, depicts TRIVIA with all negative ratings of the example seller profile being highlighted in the parallel coordinates visualization. The changes are also reflected in all other visualization to reveal coherence (link & filter).

Case studies

To evaluate the proper functioning of TRIVIA, we set up a testbed that allows to simulate transaction and review activities of an eBay-like electronic marketplace. In order



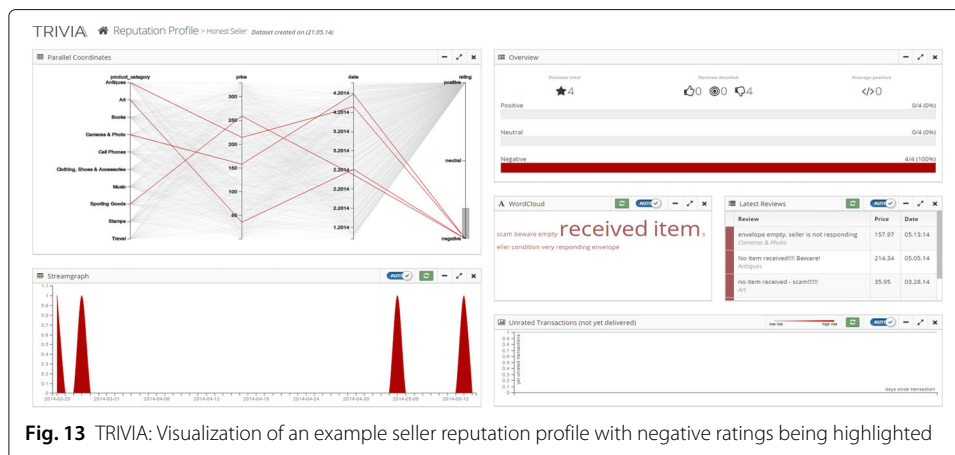


to demonstrate how value is added through a visual representation of seller reputation profiles in this environment, we simulated three seller attacks. The resulting reputation values are on the one hand calculated by two common reputation systems namely the Beta Reputation System [30] and a system computing a simple average of positive and negative ratings. On the other hand, we interactively visualize the malicious sellers' reputation profiles and point out its merit compared to presenting a numerical value only. Note that the primary objective of this evaluation is not to compare the reputation values of different systems and measure their effectiveness, but to demonstrate an alternative approach of how the problem of robustness and malicious seller detection could be addressed through user interaction while maintaining transparency. Thus, we did not use any of the testbeds introduced in literature such as the Agent Reputation and Trust (ART) [39] testbed or the Trust and Reputation Experimentation and Evaluation Testbed (TREET) [40] that were designed to evaluating the performance of agents or measuring payoffs in a game-based manner. Instead, we created a simple simulation data set that is visualized at one specific point in time.

All visualized reputation profiles illustrated here, can also be inspected online⁴.

The electronic market testbed

To carry out experiments in an eBay-like environment as described in Section “Implementation”, we set up an electronic market testbed. We simulated transaction activities for one honest seller (see Figs. 12 and 13), three malicious sellers and $n = 100.000$



buyers (n is large enough that buyers do not learn from their experience). Each malicious seller implements one of the three attacks *playbook*, *value imbalance attack* and *reputation lag attack*. The proliferation attack and the re-entry attack were not considered as these problems need to be solved through registration policies which are beyond the scope of reputation systems. Collusion and sybil attacks are usually implemented as advisor attacks. Since we focus on seller attacks and assume every review to be honest, we did not regard these types of attacks either. To simulate the reputation lag, the time to delivery (ttd) for usual transactions was set to an interval of $ttd = \{4, 5\}$ days. We assume that transactions are directly rated on delivery.

This setting was simulated for 100 days. After this period, we compared the output of two common reputation systems - the Beta Reputation System and a reputation system computing a simple average - to the visual representation of our software. The Beta Reputation System (BRS) uses the beta probability density function (beta PDF) to derive a reputation value. Result is a reputation score that is described in a beta PDF parameter tuple (α, β) , where α represents positive and β represents negative ratings [30]. In contrast to the simple average, the BRS includes a forgetting factor that discounts referrals based on their age. This difference makes the selection of those two systems particularly interesting, since both computation models have their weaknesses which can be mitigated by a visual representation. For the forgetting factor we set $\lambda = 0.95$ and chose a weight of $w = 1$ for each referral. The resulting reputation values were normalized to the interval of $[-1; 1]$ to make them comparable.

Case 1: The Playbook

A malicious seller implementing a Playbook tries to maximize his profit by employing a cheating strategy that bases on a sequence of actions (see Table 1). Attackers would have a “book” of different plays which allows them to dynamically decide which actions to apply [3]. Though playbooks are an attack toward reputation systems, they are not necessarily unethical as pointed out in [5]. Generating oscillation in a company’s reputation is often also used by commercial players.

To demonstrate a Playbook, we simulated a simple attack where the malicious seller first acts honestly to build a high reputation. After a specific time, he “milks” his high reputation by providing low quality products or services under low production costs. Exceeding a defined threshold value or time-frame, he changes his behavior and acts honest again. This sequence of actions could be implemented in an infinite loop [5]. In our simulation, the malicious seller acts honest for 20 days, then he starts cheating for four days. Due to a reputation lag of four to five days, he will act honest again before buyers recognize him to cheat. This sequence is repeated over the entire time. In each period, the transactions were randomly spread with an average of 10 transactions per day. Table 3 shows the reputation values calculated for this profile after 100 days.

The main difference between both metrics is obvious. Due to the forgetting factor, the BRS quickly adapts its output to the current situation while the simple average allows to

Table 3 Reputation values calculated by the BRS and the simple average

Reputation system/model	Reputation value (after 100 days)
BRS	0.91
Simple average	0.6875 (84 % positive ratings)




Fig. 14 TRIVIA: Visualization of ratings over time in a stacked graph

profit from a high reputation for a longer time. Because of the simple average’s persistence, its reputation value is significantly lower after 100 days. However, both systems compute only one reputation value (at every point in time) that serves as basis for a decision. Providing numerical values only, the reputation system completely neglects the fact that the ratings have been changing periodically. Although the quality of sold products might currently be good it can rapidly change. Additionally visualizing the progression in a stacked graph, users will get an idea of this coherence. Through allowing the user to focus on single parts of the entire timeline, they can become aware of the current situation and understand the malicious seller behavior as depicted on Fig. 14. Buyers that set a great value on high quality products or services should avoid to buy from this seller. Particularly, if a service is subscribed or leased for a longer time-frame oscillation might be undesirable.

Case 2: value imbalance attack

The value imbalance attack is currently a quite powerful exploit, since most common reputation systems used in practice do not involve transaction context in their computation mechanism. Sellers benefiting from the value imbalance problem build a high reputation selling cheap products while cheating on the expensive ones (see Table 1). In this way, A Californian seller managed to deceive victims for over \$300.000 [43].

To demonstrate the value imbalance attack, we again simulated a malicious seller profile over a time-frame of 100 days. In this time he randomly accomplishes 1000 transactions - 950 selling low price items (honest) and 50 selling high price items (dishonest). The idea is to maintain a reasonably high level reputation over the whole time. Table 4 shows the resulting reputation values calculated. After 100 days, the BRS provided a reputation value of 0.86 (ranges between 0.7 and 0.99 after bootstrapping). The simple average, in contrast, converges to 0.90 after 100 days.

Since the analyzed metrics do not consider transaction context as a computation parameter, the attack remains mostly undetected. Both provide a quite high reputation

Table 4 Numerical reputation values over time for the value imbalance attack

Reputation system/model	Reputation value (after 100 days)
BRS	0.86
Simple average	0.90 (84 % positive ratings)

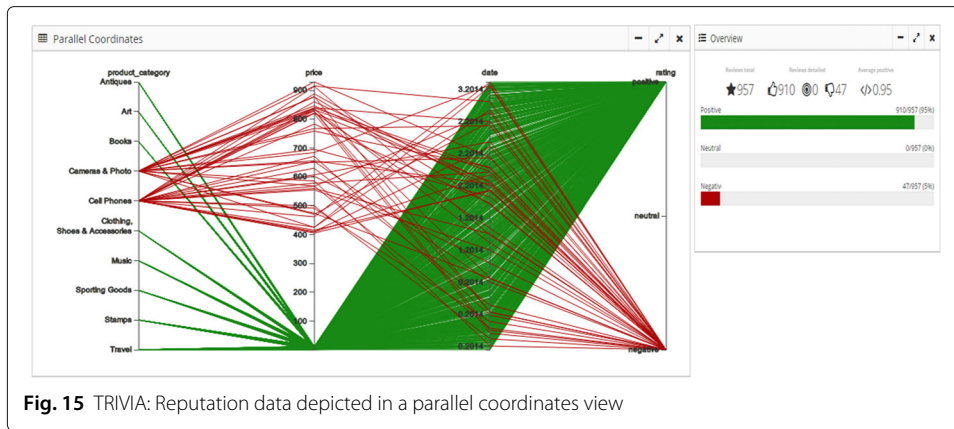


Fig. 15 TRIVIA: Reputation data depicted in a parallel coordinates view

value throughout the entire experiment. Depicting all context attributes in a parallel coordinates visualization, instead, the relation between negative ratings and the price can easily be revealed by the human analyst. While Fig. 15 gives an overview of all reviews in one picture, the correlation becomes obvious when highlighting the price range from 300 to 900 (see Fig. 16). For this price range the malicious actor received 47 negative ratings. The merit of the visualization compared to metrics that include transaction context such as recently proposed by Zhang et al. [43] is that all relevant data is transparently depicted and absolutely replicable to the user.

Besides the price, a different manifestation of this attack could be a discrimination based on different product categories. Further examples of using a parallel coordinates visualization to detect context-based attacks can be found in a recent work [16]. Thereby, they carry out case studies by means of real-life data.

Case study 3: reputation lag attack

Exploiting the reputation lag problem, a malicious seller behaves honestly to build a high reputation. After some time, he changes his behavior to maximize his profit by cheating on as many customers as possible before he is rated badly (depending on the market that might take quite a few days; here we assume 15 days). This weakness theoretically facilitates unlimited cheating [21]. Though the reputation lag attack is broadly known, there is nearly no approach to cope with this challenge.




Fig. 16 TRIVIA: Reputation data depicted in a parallel coordinates view with high price transactions selected



Fig. 17 TRIVIA: Visualization of the “number yet unrated transactions” on “days since transaction” after 95 days

To demonstrate how this attack can be detected visually, we simulated a malicious seller accomplishing honest transactions for 90 days. In this time-frame he randomly performs an average of 5 transactions per day (mean number of transactions per transaction day, $\overline{ntd} = 5$). The mean time to delivery (rating) is $\overline{tt\bar{d}} = 4.5$ days. After this honest period, the malicious seller starts cheating for the remaining 10 days by selling very low quality products. Since he wants to profit from this attack for a longer time-frame, he delays the delivery. He, moreover, increases his number of sales (through e.g. lowering the price). The beta reputation system as well as the simple average completely lack in providing information about transactions not yet rated.

Visualizing yet unrated transactions, instead, a buyer gets an overview of how many customers are currently waiting for their delivery. Figure 17 depicts the dishonest period after 5 days (on day 95 of the whole experiment). Obviously, the seller started notably more transactions as usual since the mean number of transactions per transaction day is widely exceeded. Though such a picture does not always depict an attack, it calls a buyer to think about the reliability of the seller and might therefore bound the number of fraud victims from above.

Analyzing the same profile five days later (Fig. 18), it becomes even more obvious that something unusual is happening as the mean time to rating (delivery) is now widely exceeded for many transactions as well.

Discussion and conclusion

The objective of this research is to increase the robustness of reputation systems while maintaining their transparency by means of an interactive visual representation of reputation profiles. To take a step into this direction, we adapted a range of interactive visualizations to depict reputation data. Based on these, we introduced TRIVIA a visual analytics tool that focuses on the detection of malicious sellers in electronic marketplaces as a large application area of online reputation systems.



Fig. 18 TRIVIA: Visualization of the “number yet unrated transactions” on “days since transaction” after 100 days

As demonstrated above, we could show that an interactive visual representation can notably contribute to a transparent presentation of reputation data, as all input data is depicted in one integrated view. The interaction techniques implemented in TRIVIA allow the users to gain insights by themselves and “dive” into the data. As a consequence, a higher level of trust toward the reputation system could be achieved, because it does not only provide a pre-calculated reputation value but rather allows to evaluate the referrals and derive one’s reputation. The visualization techniques we adapted have proven to be adequate for the specific cases.

We furthermore found that an interactive visualization can contribute to the robustness of current reputation systems against three major attacks. In contrast to most advanced computation models that have become quite complex, presenting the share of positive, neutral and negative ratings combined with interactive visualizations has shown to be sufficient to detect malicious sellers and even obtain better results. While most systems compute a static reputation value, we found that involving the user in reputation assessment and letting him dynamically change the parameter used for the computation, malicious behavior can reliably be identified.

One counter-argument a reader could forward is that although we criticized the high complexity of computation models, the models our visualizations are based on are at least as complex. Here, we are on the same page with Marsh et al. [15] who argue that “*Most especially, if complex models are used in the background for whatever reason, the user interface has a role to make the model understandable without losing any of its predictive power*”. Obviously, a proper visualization of a model can provide much more information than the numerical outcome of a metric. Nevertheless, as this is one of few proposals trying to involve the user in reputation assessment through interactive visualizations, there is vital necessity for further research on finding and evaluating the ideal presentation. Also, please note that the objective of this work is not to propose a visual analytics software that can directly be used in practice, but to point out the benefits of transparently and interactively presenting different facets of reputation data using visual analytics.

Reviewing the evaluation of TRIVIA in this paper, there are some limitations that should be addressed in future research. It remains to show (1) whether an average end-user is able to use the single visualizations and the visualization cockpit for detecting attacks, (2) whether this approach indeed increases the understandability of reputation data for end-users, and (3) whether such visualizations will be adopted by end-users quickly (with regard to perceived usefulness and usability). To this end, we will conduct a user study that analyses and compares the participants’ usage of the single visualization and interaction techniques as well as their combination. An exemplary international user study, that compared the usage of an eBay-like feedback profile versus a new interface involving an interactive parallel coordinates visualization was already conducted in a recent work [41]. During the study, participants had to solve four cases. In each case, they were asked to buy a specific item, compare two seller profiles offering the respective item and give a preference for one of both sellers. In each case, one seller was malicious whereas the other seller was honest. Results of the study show that the detection ability and robustness against context-based attacks could notably be increased through using the new interface including a parallel coordinates visualization compared to the eBay-like interface. Also, the understanding of malicious behavior could be significantly enhanced. During the study participants had to describe why they made a decision. Measuring the

perceived usability, no trade-offs between using the eBay-like interface compared to the interactive interface were detected. Furthermore, users of the novel interface denoted that they would use such an interactive visualization if it was available. Based on these outcomes, we feel confident that a user study analyzing TRIVIA may also provide some valuable insights.

Overall, we believe that our results provide a promising perspective on an alternative approach for reputation assessment and detecting malicious actors. We thus motivate to carry out further research in enhancing presentation and making use of the user's cognitive capabilities in online reputation systems.

Endnotes

¹ <http://trust.bayforsec.de/visualrepcockpit/>

² <http://d3js.org/>

³ <http://c3js.org/>

⁴ <http://trust.bayforsec.de/visualrepcockpit/>

Authors' contributions

JS proposed the initial idea of this paper. He developed TRIVIA and conducted the evaluation of the proposed ideas. GP supervised the research, contributed to the paper writing and made suggestions. Both authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Received: 9 October 2015 Accepted: 8 October 2016

Published online: 03 November 2016

References

1. Yao Y, Ruohomaa S, Xu F (2012) Addressing common vulnerabilities of reputation systems for electronic commerce. *J Theor Appl Electron Commer Res* 7(1):1–20
2. Diekmann A, Jann B, Przepiorka W, Wehrli S (2014) Reputation formation and the evolution of cooperation in anonymous online markets. *Am Sociol Rev* 79(1):65–85
3. Kerr R, Cohen R (2009) Smart cheaters do prosper: Defeating trust and reputation systems. In: *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 2. AAMAS '09*. International Foundation for Autonomous Agents and Multiagent Systems, Richland. pp 993–1000
4. Dellarocas C (2000) Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In: *Proceedings of the 2nd ACM Conference on Electronic Commerce*. ACM, New York. pp 150–157
5. Jøsang A, Golbeck J (2009) Challenges for robust of trust and reputation systems. Saint Malo
6. Zhang L, Jiang S, Zhang J, Ng WK (2012) Robustness of trust models and combinations for handling unfair ratings. In: *Dimitrakos T, Moona R, Patel D, McKnight DH (eds). Trust Management VI. IFIP advances in information and communication technology*. Springer, Berlin and Heidelberg Vol. 374. pp 36–51
7. Tavakolifard M, Almeroth KC (2012) A taxonomy to express open challenges in trust and reputation systems. *J Commun* 7(7):538–551
8. van Valen L (1973) A new evolutionary law. *Evol Theory* 1:1–30
9. Carroll L, Tenniel J (1897) *Through the Looking Glass: And What Alice Found There*. Altemus' illustrated young people's library. Henry Altemus, Philadelphia
10. Hammer S, Kieffhaber R, Redlin M, Andre E, Ungerer T (2013) A user-centric study of reputation metrics in online communities. *CEUR Workshop Proceedings 997*, CEUR-WS.org, Rome
11. Keim DA, Kohlhammer J, Ellis G, Mansmann F (2010) *Mastering the Information Age - Solving Problems with Visual Analytics*. Eurographics Association, Geneva. <http://books.google.de/books?id=vdv5wZM8i0C>. Accessed 19 Oct 2016
12. Marsh SP (1994) *Formalising trust as a computational concept*. PhD thesis. <http://www.cs.stir.ac.uk/research/publications/techreps/pdf/TR133.pdf>. Accessed 19 Oct 2016
13. Resnick P, Kuwabara K, Zeckhauser R, Friedman E (2000) Reputation systems. *Commun ACM* 43(12):45–48
14. Sänger J, Pernul G (2014) Reusability for trust and reputation systems. In: Zhou J, Gal-Oz N, Zhang J, Gudes E (eds). *Trust Management VIII. IFIP advances in information and communication technology*. Springer, Berlin Heidelberg Vol. 430. pp 28–43
15. Marsh S, Basu A, Dwyer N (2012) Rendering unto caesar the things that are caesar's: Complex trust models and human understanding. In: *Dimitrakos T, Moona R, Patel D, McKnight DH (eds). Trust Management VI. IFIP Advances in Information and Communication Technology*. Springer, Berlin and Heidelberg Vol. 374. pp 191–200
16. Sänger J, Pernul G (2014) Visualizing transaction context in trust and reputation systems. In: *Proceedings of the 9th International Conference on Availability, Reliability and Security (ARES)*

17. Sänger J, Richthammer C, Kunz M, Meier S, Pernul G (2015) Visualizing unfair ratings in online reputation systems. In: ECIS 2015 Completed Research Papers. Paper 159. IFIP Advances in Information and Communication Technology. AISel, AIS Electronic Library. pp 1–15
18. O'Donovan J, Smyth B, Evrim V, McLeod D (2007) Extracting and visualizing trust relationships from online auction feedback comments. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence. IJCAI'07. Morgan Kaufmann Publishers Inc, San Francisco. pp 2826–2831
19. Hansson K, Karlström P, Larsson A, Verhagen H (2011) Actory?: A tool for visualizing reputation as a means to formalize informal social behavior., Montpellier
20. Nurse JRC, Agrafiotis I, Goldsmith M, Creese S, Lamberts K (2014) Two sides of the coin: measuring and communicating the trustworthiness of online information. *J Trust Manag* 1(5)
21. Kerr R, Cohen R (2006) Modeling trust using transactional, numerical units. In: Proceedings of the 2006 International Conference on Privacy, Security and Trust: Bridge the Gap Between PST Technologies and Business Services. PST '06. ACM, New York. pp 21–12111
22. Hoffman K, Zage D, Nita-Rotaru C (2009) A survey of attack and defense techniques for reputation systems. *ACM Comput Surv* 42(1):1–31
23. Fraga D, Bankovic Z, Moya JM (2012) A taxonomy of trust and reputation system attacks. In: Proceedings of the 2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications. TRUSTCOM '12. IEEE Computer Society, Washington DC. pp 41–50
24. Shneiderman B (1996) The eyes have it: A task by data type taxonomy for information visualizations. In: Proceedings of the 1996 IEEE Symposium on Visual Languages. VL '96. IEEE Computer Society, Washington DC. pp 336–343. <http://dl.acm.org/citation.cfm?id=832277.834354>. Accessed 19 Oct 2016
25. Keim DA (2001) Visual exploration of large data sets. *Commun ACM* 44(8):38–44
26. Mazza R (2009) Introduction to Information Visualization. 1st. Springer
27. Seifert C, Kump B, Kienreich W, Granitzer G, Granitzer M (2008) On the beauty and usability of tag clouds. In: Proceedings of the 2008 12th International Conference Information Visualisation. IV '08. IEEE Computer Society. pp 17–25
28. Hassan S, Sänger J, Pernul G (2014) Soda: Dynamic visual analytics of big social data. In: Proceedings of the International Conference on Big Data and Smart Computing (BIGCOMP). pp 183–188
29. Alencar AB, de Oliveira Maria Cristina F, Paulovich FV (2012) Seeing beyond reading: a survey on visual text analytics. *Wiley Interdisc Rev Data Mining Knowl Discov* 2(6):476–492
30. Jøsang A, Ismail R (2002) The beta reputation system, Bled Slovenia
31. Aigner W, Miksch S, Schumann H, Tominski C (2011) Survey of visualization techniques. In: Visualization of Time-Oriented Data. Human-Computer Interaction Series. Springer, London. pp 147–254
32. Lee B, Riche NH, Karlson AK, Carpendale S (2010) Sparkclouds: Visualizing trends in tag clouds. *IEEE Trans Vis Comput Graph* 16(6):1182–1189
33. Riegelsberger J, Sasse MA, McCarthy JD (2005) The mechanics of trust: A framework for research and design. *Int J Hum-Comput Stud* 62(3):381–422. doi:10.1016/j.ijhcs.2005.01.001
34. Buntain C, Golbeck J (2015) Trust transfer between contexts. *J Trust Manag* 2(1):1–16. doi:10.1186/s40493-015-0017-1
35. Rehak M, Gregor M, Pechoucek M, Bradshaw J (2006) Representing context for multiagent trust modeling, Hongkong
36. Wegman EJ (1987) Hyperdimensional Data Analysis Using Parallel Coordinates
37. Heer J, Shneiderman B (2012) Interactive dynamics for visual analysis. *Commun ACM* 55(4):45–54
38. Keim DA (2002) Information visualization and visual data mining. *IEEE Trans Vis Comput Graph* 8(1):1–8
39. Fullam K, Klos T, Muller G, Sabater-Mir J, Barber KS, Vercouter L (2006) The agent reputation and trust (art) testbed. In: Stølen K, Winsborough W, Martinelli F, Massacci F (eds). Trust Management. Lecture Notes in Computer Science. Springer, Berlin Heidelberg Vol. 3986. pp 439–442
40. Kerr R, Cohen R (2010) Treet: the trust and reputation experimentation and evaluation testbed. *Electron Commer Res* 10(3-4):271–290
41. Sänger J, Hänsch N, Glass B, Benenson Z, Landwirth R, Sasse MA (2016) Look before you leap: Improving the users' ability to detect fraud in electronic marketplaces. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. CHI '16. ACM, New York. pp 3870–3882. doi:10.1145/2858036.2858555, <http://doi.acm.org/10.1145/2858036.2858555>
42. Spitz S, Tüchelmann Y (2011) A survey of security issues in trust and reputation systems for e-commerce. In: Calero JA, Yang L, Mármol F, García Villalba L, Li A, Wang Y (eds). Autonomic and Trusted Computing. Lecture Notes in Computer Science. Springer, Berlin Heidelberg Vol. 6906. pp 203–214
43. Zhang H, Wang Y, Zhang X (2012) A trust vector approach to transaction context-aware trust evaluation in e-commerce and e-service environments. In: Proceedings of the 5th IEEE International Conference on Service-Oriented Computing and Applications, Taipei Vol. SOCA. pp 1–8