

NetTrust - Recommendation System for Embedding Trust in a Virtual Realm

Preeti Hariharan
School of Informatics
Indiana University
Bloomington, Indiana, USA
pharihar@cs.indiana.edu

Farzaneh Asgharpour
School of Informatics
Indiana University
Bloomington, Indiana, USA
fasgharp@indiana.edu

L. Jean Camp
School of Informatics
Indiana University
Bloomington, Indiana, USA
ljcamp@indiana.edu

ABSTRACT

NetTrust is a system that embeds social context in browsing by combining individual histories, social networks, and explicit ratings. NetTrust embeds two types of recommendation systems to communicate recommendations about the trustworthiness of a website. The first recommendation system uses shared browsing histories from a user's self-selected social networks to create both explicit and implicit data collection. The second recommendation system is based on listings by third parties. NetTrust targets the human element of trust and proposes reputation based solutions to issues of trust on the web. The reputation systems used in NetTrust allow an individual to select their own trusted sources of information and rate particular sites as trustworthy (or not). NetTrust allows an individual to select their own trusted sources of information from a market of ratings agencies and combine these ratings with the reputation information from their individual social network.

Categories and Subject Descriptors

H.5.3 [Group and Organization Interfaces]: Collaborative computing, Web-based interaction; K.4.4 [Electronic Commerce]: Security; H.1.2 [User/ Machine Systems]: Human factors, Software psychology; D.2.2 [Design Tools and Techniques]: User interfaces

General Terms

Usability, Toolbar, Experiment, Security, Trust, Phishing

Keywords

Recommendation Systems, NetTrust, Security, Privacy, rating, Phishing, Peer Production

1. INTRODUCTION

A reputation system collects, distributes, and aggregates feedback about participants' past behavior. Recommender

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM Recommender Systems 2007 Minneapolis, Minnesota, USA
Copyright 2007 ACM 0-12345-67-8/90/01 ...\$5.00.

systems use the opinions of a community of users to help individuals in that community more effectively identify content of interest from a potentially overwhelming set of choices [10, 1].

Recommendation Systems must be robust and reliable [5]. To have a viable system, a reputation system must present ratings that are usable and appropriate for the context. For many systems, the context is defined by the scope of the reputation system: movie rentals, eBay transactions or books. For NetTrust the context is highly variable. The goal of the system is to provide reputation that is useful and usable in a wide range of contexts. The designers of NetTrust present reputation information which informs user choice. The second goal of NetTrust is to make security technology enjoyable to use. It thereby removes the perceived trade-off between usability and security.

A reputation system can encourage more informed user trust behavior, and deter participation by those merchants who are unskilled or dishonest. As a solution to the ubiquitous problem of trust in new short-term commercial relationships on the Internet, reputation systems have an immediate appeal. The participants of a reputation system can themselves create a safe community [13]. In terms of phishing and masquerade fraud, the absence of reputations has been utilized by third parties who leverage the de-contextualization of browsing to fool the human in the loop [16].

The research question we address in this work is: Can a recommendation system help in improving the state of web security and reduce the efficacy of phishing and pharming?

To be specific, we describe a reputation-based application called NetTrust. NetTrust is an application that is presented to the user as a toolbar plug-in for Firefox Web Browser. This toolbar is constructed on sound technical and social foundations to prevent masquerades. NetTrust embeds social context in web-based trust decisions by combining individual histories, social networks, and explicit ratings. This paper is a part of the ongoing research on NetTrust.

After describing NetTrust, we describe two experiments. We extend previous work by implementing a user study about the usability and utility of NetTrust toolbar. We study both classic usability and the more specific issue of trust behavior.

In Section 2, we present the research motivating this work. In Section 3, we review the NetTrust system, addressing its principles and architecture. In Section 4, we discuss the design and analysis of the usability test on NetTrust and describe our findings. Section 5 concludes the paper.

2. BACKGROUND

In an online virtual realm, anonymity of users allows for misuse of information [8]. Thus there is a need for trust. Informal trust decisions require trust signals that are meaningful, useful and timely. Trust involves a combination of many factors that work together to “facilitate cooperative behavior” [14, 9] by reducing complexity and enabling people to make decisions in high levels of uncertainty. Dingleline argues that provision of a trust based reputation mechanism will maintain accountability in online peer-to-peer reputation systems [5]. Recent research has shown that incorporating trust and reputation models into the recommendation process can have a positive impact on the accuracy and robustness of recommendations [12].

2.1 Models of Trust

A potentially useful approach in dealing with recommender attacks by malicious or untrustworthy users is to model user trust explicitly. Massa et al. in [11] has built a trust model from explicit trust ratings given by users in the popular Epinions.com service.

The Xtrust System deploys trust in a recommender system. Xtrust does not use contextual information [3]. Xtrust calculations discount individuals’ assertion of trustworthiness uniformly in all situations.

In Pretty Good Privacy (PGP), individuals publicize unique public key identifiers with claims of identity associations. Each individual is responsible for disseminating and gaining information. Identity verification is based purely on social network information [7]. There is no centralized agency of trust. PGP is designed only to verify identity claims in email.

Blaze, Feigenbaum, and Strauss introduced PolicyMaker Trust Management System [2], a decentralized trust management system. However, Policymaker requires users to be aware of all possible contexts and explicitly rate on the basis of this knowledge.

Consider for example, a student in Computer Science deciding which courses to take for a semester. The student would enquire with his colleagues studying Computer Science rather than a friend who is a music major. Note that when it comes to deciding the best place to buy a CD, the student may prefer asking the friend who was a music major.

The above scenario shows that people place different levels of trust on different people in different contexts. Thus it is very important to keep track of this contextual information while designing a trust based management system.

Moreover, in the e-commerce domain, trust cannot be completely defined either by the security of the information system or by the individual notion of trust through socialization. It must use a combination of the two factors, thereby enabling people to use a simple interface to make educated decisions in a situation characterized by high levels of uncertainty.

3. NETTRUST

3.1 Overview

The purpose of NetTrust is to communicate meaningful information with respect to trust by supplying a recommender system that embeds both social networks and centralized authorities. With the help of user’s social network as well

as global authorities, the application aggregates reputation ratings into a simple display. Thus the user will be able to make a contextual decision concerning the trustworthiness and reliability of a website.

Phishing has become the latest threat to internet users. A phishing site impersonates as a trusted website, for example a well-known bank, in order to attain personally identifiable information. Such information includes passwords and account numbers, credit card numbers, and social security numbers.

There are some toolbars in market, which target only phishing threats. For example, Spoofguard [4] is one such toolbar which uses real-time information like links, images in the website, lack of SSL certificate and other such information to calculate a ‘spoof score’ to signal trust. However, these features are under the control of the malicious phish guard.

NetTrust uses real-time social network information in addition to the above mentioned features, which makes it so unique.

NetTrust uses features that are not under the control of the malicious agent: user social network, user history and the history of the user’s social network. In addition, the NetTrust toolbar takes advantage of a characteristic of phishing sites to prevent one phishing victim from misdirecting others - the temporal patterns of phishing sites. Phishing sites go up, are identified, and are taken down [15]. Phishing sites do not stay up over long periods of time. The impermanence of phishing sites is integrated into the reputation system as described below.

The NetTrust toolbar is the mechanism for interaction with the user. It integrates social network information. Individuals may have multiple social networks: home, family, hobby, political or religious.

Each social network is associated with a pseudonym, which we call a “nym”. Each nym is a member of disparate social networks. Social relationships are often contextual, so we find ourselves sharing varied information with different sets of people [6]. Each person can construct as many nyms or persona that they desire, where each persona interacts with a different social network.

The use of pseudonyms in NetTrust enables the construction of boundaries between the various roles assumed by one person. The recommendation system depends to a great extent on the role (nym) selected by the user and the corresponding social network.

The aggregated ratings shown on the NetTrust toolbar are different for different users, and are dependent on the trust that his social network has for that website.

In addition to the user’s chosen social network, the user can subscribe to receive information from some approved trusted third parties. We call these centralized trusted third parties “Broadcasters”. They are called broadcasters to emphasize that they distribute trust information but have no privileges or authorization on the user’s machine after they have been selected. Broadcasters distribute “red” or “green” lists which can be stored and searched locally. These broadcaster ratings act as signals of trust for a user, for a website.

Sybil attacks are said to take place whenever the user adds himself more than once in his own social network. These attacks are mitigated in NetTrust using a technique as explained below: Consider Alice and Bob as virtual friends. Alice first creates a pseudonym for herself. She then sends



Figure 1: NetTrust toolbar (Net Trust, Login, Rate Site, Friends, Rating Lights) in the usability test



Figure 2: NetTrust toolbar (Comments and Broadcasters panels) in the usability test

an invitation nonce to Bob. Once Bob joins the system, Alice’s history is shared with Bob’s chosen pseudonym. To mitigate this we limit the number of people joining any social network (twenty per pseudonym) to decrease the likelihood of a stranger joining the network. In economic terms, a participatory slot network in a social network is a scarce (and thus valuable) resource.

3.2 Principles

This section explains in brief the architecture of the NetTrust reputation system.

Explicit Data Collection: When a user visits any website, NetTrust allows the user to explicitly rate the site. The user with his chosen nym makes a choice to rate the website using a positive scale of 1-5 or a negative scale of -1 to -5. Along with rating a website, the user can also enter a comment about the website. Thus NetTrust allows for a two way interaction. Note here that the aggregated ratings seen on the toolbar are cumulative, including the users’ rating and the ratings given by his friends for the same website.

Implicit Data Collection: If a user, Alice, repeatedly visits a website X, we can assume that Alice has found X trustworthy even though she may not explicitly claim so. Therefore, NetTrust keeps track of the visiting frequency of each website and use it as one of the involving factors in rating web sites. We call this kind of rating *implicit rating*. A user can turn off this feature of NetTrust by switching the NetTrust toolbar to persona=“None” state.

Third Party Broadcasters As described above broadcasters such as the Better Business Bureau or Pay Pal can be added. If the user has added a broadcaster, the toolbar will extract the ratings of the website the user is visiting from the broadcaster ratings. Broadcaster’s ratings are shown as positive with a happy face and negative as yuck face. Figure 2 illustrates the NetTrust toolbar broadcaster ratings.

Recommendation System: On the basis of the user’s implicit rating and explicit rating, the user is given aggregated ratings about the website. These ratings appear as rating lights on the NetTrust toolbar (refer to Figure 1). Apart from numeric ratings, the user is provided feedback about the website in the form of comments as described above. The user can view all the comments given by the buddies who have reviewed the website in the past.

These recommendations are dynamic and are dependent on the number of buddies one has, their ratings for the given website and the number of times the buddies have visited that website (assuming the one-week anti-phishing time delay has not expired).

4. EXPERIMENTS

We conducted two tests to measure the performance of NetTrust toolbar: 1. a usability test to measure the toolbar’s ease of use 2. a trust test to measure how much users’ decisions have been effected by NetTrust toolbar.

4.1 Usability test

Usability Test Setting: The usability test was done using a middle fidelity prototype of the toolbar. Figures 3 and 4 illustrate the NetTrust toolbar as it was in the usability test. Each participant first filled a pre-questionnaire regarding his/her demographics as well as his/her general behavior in online transactions. Then, we briefly introduced the NetTrust toolbar to the participant. We introduced the toolbar as an instrument which enables them to share their opinions about the trustworthiness of a visited web site with a selected group of their friends. After introduction to the toolbar the participants received no other guidance on use or details of the toolbar.

The website zappos.com, which offers general merchandise such as clothing and jewelries, was used to conduct the usability test. The participants could click on links and navigate through the page.

The toolbar was pre-loaded with a list of friends, their comments about zappos.com, and the implicit and explicit ratings. We asked the participants to assume the preregistered list of friends as their own registered friends. The participant were then asked to perform a series of tasks as follows. The usability test was concluded by having the participant filling a post-questionnaire regarding his/her general impression about the toolbar.

Tasks: Throughout the usability test, each participant was asked to think aloud while performing the following tasks: 1. add and delete a friend, 2. add and delete a broadcaster, 3. find and read the comments about zappos.com given by their friends, 4. add comments about the website zappos.com, 5. interpret the rating lights and broadcasters’ ratings, 6. rate the website zappos.com.

Participants: We had 7 male and 8 female participants all within the age range of 18-36. The participants were all frequent internet users. Except two participants, all of them shopped online using their credit cards.

Problems Discovered: All users completed all tasks. 13 out of 15 found the toolbar was easy to use and the interactions were simple. Two participants were able to perform the tasks but expressed some frustration during task completion. As evidence, all the participants could add/delete friends, add/read comments and add/delete a broadcaster. Since the toolbar is more for making decisions than perform-



Figure 3: NetTrust toolbar’s status for the Website remindingyou.name used in the Trust Experiment




Figure 4: NetTrust toolbar’s status for the Website memoryminders.us used in the Trust Experiment

ing complicated tasks, we discovered the following cognitive interaction problems:

1. 12 out of 15 participants did not understand the meaning of the word “Broadcasters”. However, all of them interpreted the happy and yuck faces correctly. 90% of participants knew the “BBB” as “Better Business Bureau”, but only 50% knew the meaning of FDIC. As a result multiple users stated “BBB’s happy face means BBB trust this site”, or the opposite with the yuck face.
2. All the participants interpreted the green and red lights as positive and negative rate. When there are no ratings of a type (positive or negative) the ratings bar is gray. Thus, 14 out of 15 participants were confused about having two sets of lights one in green and the other in gray next to one another. The color gray did not convey any meaning to them.
3. When asked to rate a site the participants first clicked on the the word “Ratings” next to the rating lights. Then, discovering that the rating lights panel is not an interactive panel, they looked for some other thing. However, with no exception, as soon as they noticed the “Rate Site” panel with thumbs up and down signs, they could easily rate the website. This observation confirms that, the word “Rating” in the lights panel was somehow confusing but the thumbs up and down.

Modifications: The following modifications were suggested to improve the usability of the NetTrust toolbar.

1. We are considering changing the word “broadcasters” to “Authorities”. We think this word is more transparent to the users.
2. The two rating lights, one referring to negative and the other to positive rating, match with the rating mechanism on the toolbar. We decided to have two rows of lights but each set be presented with a relevant sign of positive or negative.
3. The word “Rating” next to the rating lights is to be removed, and letting the lights stand framed by thumbs up (+) and down (-) alone.
4. The word “FDIC” is to be changed to “Bank”.

Table 1: Interesting features of NetTrust

like the most	% of particip.
Seeing friends’ comments	53%
Transparency	13%
Richness in identities involving trust decision	7%
Simple interaction	20%
Broadcasters	7%
Easy to interact with	20%

Likes and Dislikes: Table 1 shows participant’s opinions on what they like the most about the toolbar along with the percentage of participants who hold that opinion.

Based on this test, the most interesting feature of the NetTrust for the participants has been enabling the users to see their friends comments and to share their opinions with their peers.

73% of the users said, they like to use the NetTrust toolbar on a regular base. The rest believed “it is too much to check all of these factors about every single website”, or some mentioned “I can judge about the web sites myself, I don’t need to check with anyone else”.

The most uninteresting feature of the NetTrust in its first version has been the broadcasters. This may has been due to the lack of transparency of the word “Broadcasters” for the users.

4.2 Trust Experiment

Experimental Setup: In order to measure the impact of NetTrust on users’ decisions, we conducted an experiment having two groups of participants. One group had the NetTrust toolbar installed on their browser and the other did not have the toolbar. We refer to these groups as group 1 and group 2 respectively.

In this experiment we used three fabricated websites. The websites were offering different online products. To purchase the service offered by each website one needed to share certain information such as name, mother’s maiden name, and some financial information such as credit card number. For participants in group 1, who had the NetTrust toolbar installed on their browsers, the toolbar was preloaded with a certain rating state, friends’ comments and broadcaster’s ratings about the related website.

Group 1: The participants in group 1, were introduced to




Figure 5: NetTrust toolbar’s status for the Website elephantmine.net used in the Trust Experiment

the toolbar and were trained to conduct certain tasks on the toolbar. The tasks included rating a site, adding and removing buddies and broadcasters, as well as viewing buddies comments on the website they were visiting. After the training session they were asked to navigate through each of the three websites, read the information given by the NetTrust toolbar about the web site and finally answer a questionnaire in this regard. They were asked to assume the pre-registered buddies on the toolbar as their personal network of friends.

Group 2: Participants in the second group were asked to do the same thing, but they did not have the toolbar installed on their browser. This means, they were asked to navigate through each website and fill a questionnaire, but considering their own personal judgement about the trustworthiness of each website.

Comparing the results of the two groups we measured the impact of NetTrust on users’ trust. The two groups had the same questionnaires. Refer to the Appendix for the questionnaire.

Participants: We had 20 participants within the age of 18-35, all of them were frequent computer users. The participants were divided equally into two groups of 10. The two groups were conducting the experiment in two separate rooms, supervised by two observers.

Toolbar Status: The first two websites, the toolbar showed a large number of ”buddies” visiting the site, 4 for website 1 and 3 for website 2, respectively, as well as positive ratings for the broadcasters. The last website showed only 1 friend visiting the site and negative rating from the broadcasters. Figures 3, 4, 5 show the status of the toolbar for each of these websites.

Results: Comparing the final results from the two groups of participants we found: For elephantmine.net, 80% of the participants without the NetTrust toolbar said they would not disclose their credit card information. In contrast, the NetTrust toolbar was as shown in Figure 2. It showed a number of their friends who had seen the website before. After reading the rating lights and the comments given by the buddies, around 34% of the participants would not disclose credit card information to this website. This indicates that with the help of social network information, the toolbar increased the propensity to trust.

Similarly for remindingyou.name, 100% of the participants without the NetTrust toolbar said they would not disclose their credit card information. In contrast, the NetTrust toolbar was as shown in Figure 3. Though the rating lights were

still red, the website was visited by almost three of the user’s friends who had given it good comments. This increased the extent of trust placed on this website. Almost 77% of the participants with the NetTrust toolbar installed on their browser, said they would be willing to share their credit card information to buy the products listed on this website.

For the website memoryinders.us, 75% of the participants were not willing to share their credit card information online. With the NetTrust toolbar as shown in Figure 4, it was found that most of the friends in the users social network, disliked this site and considered it to be unsafe for credit transactions. Moreover, the rating lights indicated a high alert ”red” sign. The broadcasters conveyed conflicting information in this case.

From the study, it was found that 100% of the participants with the NetTrust toolbar installed, would refuse to sharing personal credit card information. This implies that a toolbar with meaningful contextual information involving immediate social network has a great effect on the users propensity to trust a website.

This initial test illustrates that NetTrust changes user trust behavior both positively and negatively.

5. CONCLUSION

NetTrust is a socially aware aggregation recommendation systems that gives meaningful signals. It integrates privacy-enhanced signaling into browsing. The user interface has been found to be very easy to use. We are currently working on revising the architecture of the system to make it as decentralized as possible. NetTrust embeds both implicit ratings (by shared history), explicit ratings, and third party broadcaster ratings. Not only does the application help users make educated decisions concerning websites, but it also brings social networks and groups of people together. NetTrust is a radically new way of proposing a social-technical solution to the problem of trust in digital world.

6. REFERENCES

- [1] Pew internet, online rating systems, 2004. Available online at: http://www.pewinternet.org/PPF/r/140/report_display.asp.
- [2] M. Blaze, J. Feigenbaum, and M. Strauss. Compliance checking in the policymaker trust-management system. In *Proceedings of the Financial Cryptography 98*, volume 1465 of *Lecture Notes in Computer Science*, page 254274, Berlin, 1998. Springer.
- [3] M. Branchaud and S. Flinn. Xtrust: A scalable trust management infrastructure. In *Second Annual Conference on Privacy, Security and Trust*, pages 207–218, Fredericton, New Brunswick, Canada, October 2004.

- [4] N. Chou, R. Ledesma, Y. Teraguchi, and J. C. Mitchell. Client-side defense against web-based identity theft. In *Proceeding of Network and Distributed System Security Symposium*, 2004.
- [5] R. Dingedine, M. J. Freedman, and D. Molnar. *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, chapter 16. O'Reilly and Associates, 2001.
- [6] J. Donath and D. Boyd. Public displays of connection. *BT Technology Journal*, 22(4), October 2004.
- [7] S. Garfinkel. *PGP: Pretty Good Privacy*, pages 235–236. O'Reilly & Associates Inc, Sebastopol, CA, 1994.
- [8] D. Gefen. E-commerce: the role of familiarity and trust. *The International Journal of Management Science*, 28:725–737, 2000.
- [9] A. Genkina and L. J. Camp. *Social Networks*, pages 523–550. John Wiley & Sons, Inc., Hoboken, New Jersey, 2007. Chapter 14 in: Phishing and Countermeasures, Understanding the Increasing Problem of Electronic Identity Theft.
- [10] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1):5–53, 2004.
- [11] P. Massa and B. Bhattacharjee. Using trust in recommender systems: an experimental analysis, 2004.
- [12] J. O'Donovan and B. Smyth. Is trust robust?: An analysis of trust-based recommendation. In *IUI 06*, pages 101–108, New York, NY, USA, 2006.
- [13] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems: Facilitating trust in internet interactions. *Communications of the ACM*, 43(12):45–48, 2000.
- [14] B. Schneiderman. Designing trust into online experiences. *Communications of the ACM*, 43(12):57–59, December 2000.
- [15] R. C. T. Moore. An empirical analysis of the current state of phishing attack and defense. *Workshop on the Economics of Information Security*, 2007. available online at: <http://weis2007.econinfosec.org/papers/51.pdf>.
- [16] M. Wu, R. C. Miller, and S. L. Garfinkel, editors. *Do Security Toolbars Actually Prevent Phishing Attacks?*, New York, NY, 2006. SIGCHI Conference on Human Factors in Computing Systems, CHI 06, ACM Press.

7. APPENDIX

Questions Given in the Trust Experiment

Question 1. From the following list of information check the ones that you be willing to give to memoryminder.us?

1. Your credit card number
2. Your Social Security number
3. Your year of birth
4. Your IM buddy list
5. Your list of email contacts
6. Your coworkers' names
7. Your friend's names

8. Your parents' names
9. Your family members' names
10. Your family members' birthdays
11. Your family's wedding anniversaries
12. Your family members' shopping preferences

Question 2. From the following list of information check the ones that you be willing to give to remindingyou.name?

1. Your credit card number
2. Your Social Security number
3. Your year of birth
4. Your IM buddy list
5. Your list of email contacts
6. Your coworkers' names
7. Your friend's names
8. Your parents' names
9. Your family members' names
10. Your family members' birthdays
11. Your family's wedding anniversaries
12. Your family members' shopping preferences

Question 3. From the following list of information check the ones that you be willing to give to elephantmine.net?

1. Your credit card number
2. Your Social Security number
3. Your year of birth
4. Your IM buddy list
5. Your list of email contacts
6. Your coworkers' names
7. Your friend's names
8. Your parents' names
9. Your family members' names
10. Your family members' birthdays
11. Your family's wedding anniversaries
12. Your family members' shopping preferences

Question 4. How do you decide whether or not to provide information to a web site?

Question 5. When you encounter a web site prompt for information how often do you:

1. Provide accurate information:
 - Never Sometimes Usually Often Always
2. Provide incorrect information:
 - Never Sometimes Usually Often Always
3. Just leave the site:
 - Never Sometimes Usually Often Always

Question 6. Your age:

Question 7. Your gender:

Question 8. Your years of computer experience: