iTrust: Interpersonal Trust Measurements from Social Interactions

Xiaoming Li, Qing Yang, Xiaodong Lin, Shaoen Wu and Mike Wittie

Abstract

Interpersonal trust is widely cited as an important component in several network systems such as peer-to-peer (P2P) networks, e-commerce and semantic web. However, there has been less research on measuring interpersonal trust due to the difficulty of collecting data that accurately reflect interpersonal trust. Currently, friends of a user in almost all online social networks (OSN) are indistinguishable, i.e., there is no explicit indication of the strength of trust between a user and his/her close friends, as opposed to acquaintances. To address this issue, we quantify interpersonal trust by analyzing the social interacting frequencies between users and their friends on Facebook. We consider bidirectional interacting data in OSN to deconstruct a user's social behavior and apply Principal Component Analysis (PCA) to estimate the interpersonal trust. A Facebook app, *itrust*, is developed to collect interaction data and calculate interpersonal trust. Results show that *itrust* achieves more accurate interpersonal trust measurements than existing methods.

Introduction

Recently, interpersonal trust has been applied in various systems as a key factor in decision-making processes. Taking e-commerce as an example, the opinions from trustworthy friends strongly influence a user's purchasing decisions. By leveraging a buyer's trust of his/her friends within OSN, it is possible to provide him/her with online reviews she can entirely trust. Another common example is wireless vehicular networks, in which a vehicle could determine to which neighbor to forward data by evaluating other vehicles' trustworthiness [1].

Although interpersonal trust is a very important concept in a human's life, there is no formal definition of interpersonal trust. However, most researchers agree that interpersonal trust is the willingness of accepting vulnerability or risk based on expectations regarding another person's behavior. Therefore, we define the interpersonal trust as "the probability that a trustee will behave as expected by a trustor." [2]

Quantifying interpersonal trust is a challenging problem because it is difficult to find an appropriate dataset that accurately reflects interpersonal trust. Even if such a dataset were available, accurate estimation of interpersonal trust from the dataset is non-trivial [3]. The first

attempt to quantify interpersonal trust is proposed in [4], which tried to estimate trust by analyzing email exchanges between different users. However, email communications are often the reflection of business-related activities, so they are inadequate for analyzing interpersonal trust in more general settings.

Thanks to the developments of online social networks (OSN), the richness of data generated within OSN provides unprecedented opportunities for analyzing interpersonal trust. Take Facebook as an example. In October 2013 the total number of Facebook users reached 1.26 billion and 1.23 billion of them are monthly active users. In United States, there are 128 million daily active users, i.e., about 40% of Americans use Facebook every day. Unfortunately, most OSNs do not incorporate interpersonal trust in the creation and management of relationships. The social role of a friendship was first considered in Google+ by introducing the concept of "circles." Users use circles as a way to distinguish their close friends, family and acquaintances. However, this "circles" concept does not quantify the interpersonal trust, only the nature of relationships between users.

If social networks consist of users interconnected via relationships [5], *is it possible to measure interpersonal trust from online social interactions*? We pose this question because Singh [6] has proved that social interactions have strong effects on interpersonal trust, while interpersonal trust also influences online interactions [7]. On the other hand, Onnela et al. [8] have discovered that there was a connection between tie strength and the duration of calls in mobile social networks. Because of these reasons, we pose the hypothesis that *interpersonal trust can be inferred from the frequency of social interactions in OSN*. For example, we might be able to use online interaction data in Facebook, e.g., inbox messages, photo tags and comments, to measure the interpersonal trust between users.

Friendships in most current OSNs are labeled as binary numbers, i.e., a user's close friends and acquaintances show no difference. To address this issue, we propose an innovative approach to quantify interpersonal trust based on online interactions in Facebook. We develop an app, *itrust*, to collect interaction data between a user and his/her friends. Then, we apply Principal Component Analysis (PCA) to estimate the interpersonal trustworthiness of his/her friends. Finally, a ranking list of his/her friends based on their trust values is returned.

We compare *itrust* to Vedran's weighting method [5] and the regression method proposed by Gilbert [9]. Experimental results show that *itrust* achieves a higher accuracy in estimating interpersonal trust than [5] and [9]. Besides, using the Kendall's tau and generalized Kendall's tau methods, we evaluate friends ranking (based on trust values) and find that *itrust* outperforms the other approaches.

The contributions of this paper are twofold. First, we identify a reasonable approach to estimate interpersonal trust in OSN. Second, the approach is implemented as a Facebook app, *itrust*, which can accurately rank a user's friends based on their interpersonal trust values.

Data Collection and Analysis

How many types of social interactions are there in OSN? What type of interactions reflect trust? To answer these two questions, we first collect all available interaction data in Facebook, the most popular OSN, for a set of 1.26 billion users in *itrust*. Then, we analyze the features of different types of interactions and their impacts on interpersonal trust.

itrust

Facebook provides an API for developers to collect data from any user (if permitted) in the network. On this basis, we developed an application, *itrust* (http://www.cs.montana.edu/itrust), to collect data on users' interactions and generate a ranking of the trustworthiness of a user's friends. When a user logs in to *itrust*, the app asks the user to authorize permissions to access his/her public profile, friend list, messages, news feed, relationships, status updates, and photos. After authorization, *itrust* begins to collect social interaction data. Note that collecting user interaction data may cause privacy issues [10], so *itrust* does not save any contents, but only counts the number of interactions. Based on such interaction data, *itrust* generates friends ranking list and shows it to the user.

The system architecture of *itrust* is shown in Fig. 1. The data tier stores interaction counts obtained from Facebook and friends trustworthiness results computed by the ranking calculation module. Such trustworthiness information could be used by external applications, e.g., a P2P program, to determine from which peer to download files [1].

The logic tier performs data normalization and interpersonal trustworthiness calculations. The presentation tier interacts with users, i.e., asks users to authorize permissions and display ranking list to users. In Facebook, there is no direct way of evaluating the accuracy of the interpersonal trustworthiness computed by *itrust*. Therefore, we develop the ranking evaluation module to allow a user to input his/her opinion about his/her friends' trustworthiness, which is considered the ground truth.

If users do not revoke their permission authorizations, *itrust* continuously monitors the users' interaction data, and thus addresses the temporal dynamics on interpersonal trust relationships. In other words, *itrust* provides a real-time measurement of interpersonal trust between users in Facebook.

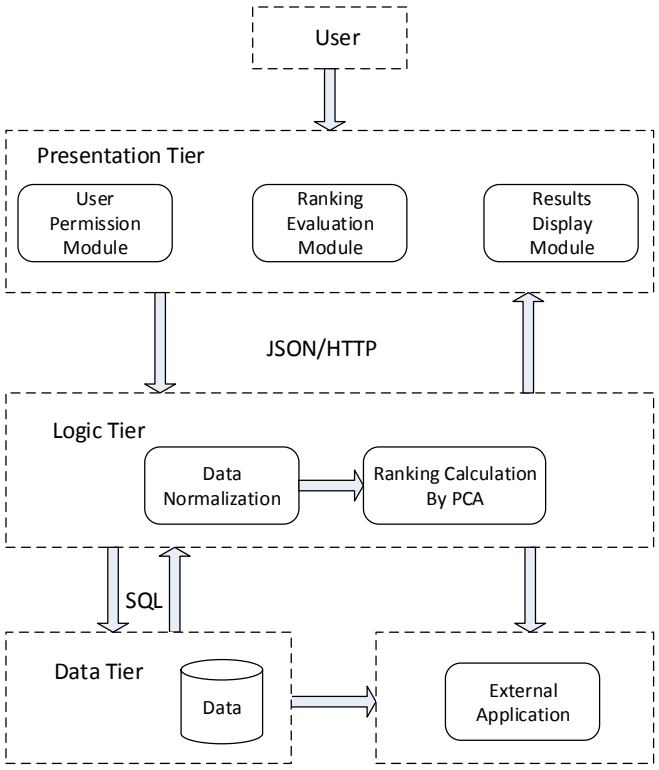

Fig.1 System architecture of *itrust*

Social Interaction Data

By 03/25/14, there is a total of 59 participants who use *itrust*. To collect trust-related interaction data, traditional methods that crawl Facebook users' webpages are not applicable because we require each participant to sign a consensus form to evaluate the trustworthiness of his/her friends. Despite the difficulty of recruiting more participants, we believe the data obtained from 59 users is generic enough to support our conclusions. First, compared to the linear regression model in [9], which recruited 35 Facebook users, we almost double the number of participants. Second, the 59 users in our study are diverse in race, age, educational background and work experience, i.e., our findings are applicable in a more general setting.

For each user, we collected twelve different types of interaction data and obtained a total of 15,158 records. The collected interaction data include: inbox messages, photo comments, photo likes, album comments, album likes, tag photos, tagged photos, tagged photo comments, tagged photo likes, tag-together photos, status comments, and status likes. After analyzing each type of data, we discover five characteristics of the interaction data in Facebook. First, large variance exists in each type of interaction data. For example, the average number of messages sent by a user is 9.54, while the maximum is 4214. The average number of status likes is 0.35 but the maximum is 53. Second, different interactions reflect interpersonal trust in different ways. For example, a user could be tagged 5 times by his/her friend A in photos, and she might also

receive 5 status likes from another friend B. Although the numbers of interactions with A and B are the same, the user may trust A more than B. Third, several interactions show a high level of correlation between each other. For instance, the interactions 'tagged photo comments' and 'tagged photo likes' are highly correlated with the number of the tagged photos. Fourth, social interactions in Facebook are directional, as is interpersonal trust. We define the data sent by a user in Facebook as his/her *outgoing* interactions and the data she receives as the *incoming* interactions. Fifth, the amounts of interaction data generated by different users are different. To support this claim, we randomly select 32 users and display the proportions of their incoming and outgoing interaction data in Fig. 2. From this figure, we see that some users (e.g., A) often publish contents but seldom interact with others, so they tend to have more incoming data (e.g., receiving comments) but less outgoing data. On the other hand, some users (e.g., B) may have less incoming but more outgoing data. Moreover, although some users (e.g., C and D) have similar proportions of incoming and outgoing data, the total amounts of their (incoming and outgoing) interactions could be very different.

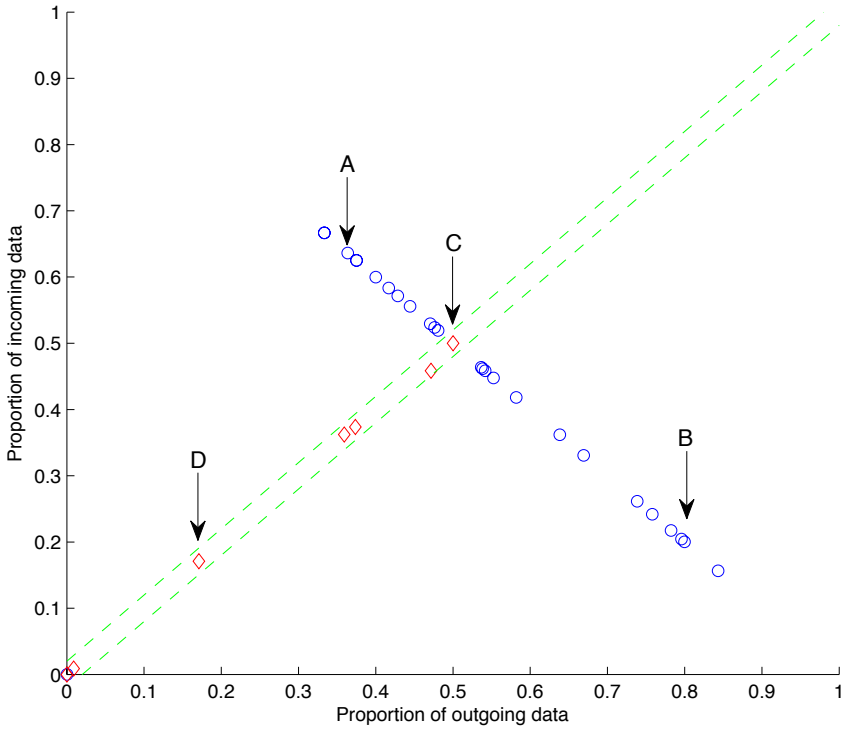

Fig.2 Incoming interaction data vs outgoing interaction data

In summary, interaction data in Facebook are disperse, diverse, correlated, directional and user-dependent; therefore, they must be processed before being used to infer the interpersonal trust information.

Trustworthiness Computation

People tend to interact frequently with a small group of people, e.g., with higher interpersonal trust, to maintain and nurture strong social ties. Inspired by the idea of measuring tie strength in social networks based on the duration of calls in the mobile phone context [8], in this paper, we investigate *whether users trust their friends proportionally to the frequency of interactions*.

Data Normalization

Interaction frequency is a sociological concept, defined as the total number of interactions per unit time. Compared to the number of interactions, interaction frequency is more accurate in measuring interpersonal trust considering some users may be newly created. To obtain the interaction frequency of a Facebook user, we first divide his/her interaction data by the lifetime of his/her account, measured in the number of months. In the following, we use incoming/outgoing data to refer the incoming/outgoing interaction frequency of a user, without causing any confusion.

Because we are interested in the trustworthiness of a user's friends, *itrust* uses a user's outgoing data to infer his/her friends' trustworthiness. Outgoing data, however, need to be normalized because they are not only dependent upon friends' trustworthiness, but also influenced by friends' activity levels. Due to social grooming, a user tends to interact more with active friends compared to inactive ones.

To normalize a user's outgoing data, we first need to measure his/her friends' activity levels. We illustrate the normalization process by an example. As shown in Fig. 3, Alice has two friends Bob and David, and we are interested to know the trustworthiness of Bob and David from Alice's perspective. The solid lines represent different types of interactions between Alice and Bob (and David), and the numbers on them denote the interaction frequencies. As shown in the figure, Alice has more interactions with Bob for any type of interaction compared to David. Intuitively, Alice should trust Bob more than David. The statement may be wrong, however, if Bob is an active user and David is an inactive one. In this case, even though Alice trusts David more than Bob, she has less chance to interact with David, due to the fact that David is inactive.

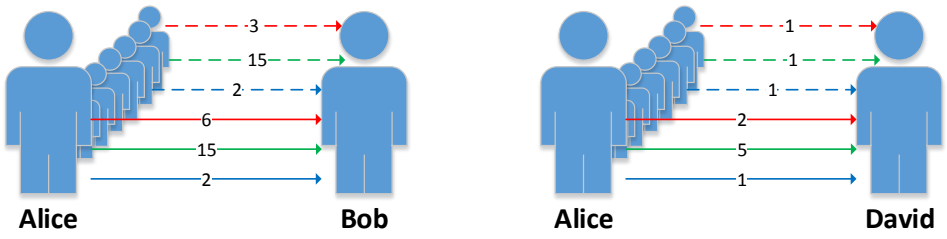

Fig.3 Data normalization based on friends' activity levels

We define the activity level of a user as the average incoming interactions from all his/her friends. In Fig. 3, the dash lines represent different types of interactions from Bob's (or David's) friends to Bob (or David), and the numbers on them denote the average interaction frequencies. We see the activity levels of Bob and David are <3, 15, 2> and <1, 1, 1>, respectively. Compared to David, Bob is more active in using Facebook, e.g., he publishes more status updates or uploads more photos, so his friends are more likely to interact with him.

To account for the activity level of Bob, the outgoing data from Alice to Bob will be normalized as follows. For each type of interaction between Alice and Bob, it will be normalized against the average incoming data of Bob (for that particular type of interaction). Therefore, we can calculate the normalized interaction between Alice and Bob as <6/3, 15/15, 2/2>=<2, 1, 1>. Similarly, the interaction between Alice and David is normalized to <2, 5, 1>.

With the above-mentioned method, for a certain user, we could obtain a normalized interaction vector for each of his/her friends. This vector includes twelve elements - normalized outgoing data for twelve types of interactions. Considering the interaction vectors from different users, an interaction matrix could be constructed, which will be used by *itrust* to compute users' trustworthiness.

Trustworthiness Ranking

Two principles need to be stated before we compute interpersonal trust in Facebook. First, the comparison of friends' trustworthiness is only valid from the perspective of a specific user. The reason is that different users perceive trust in different ways, and thus a user might be considered a close friend of one user, but an enemy of another. Second, trustworthiness is relative, so we only need to rank a user's friends based on their trustworthiness instead of computing the absolute trust values.

Although we have normalized interaction data, we cannot directly make use of them as correlations exist between different types of interactions. In other words, dependency and duplication in interactions must be removed. For example, 'tag photo' is the precondition of existing 'tag photo comments or likes', so the number of 'tag photo comments or likes' is highly dependent on the number of 'tag photo.' To address this issue, we introduce the PCA method, which not only removes correlation, but also objectively assigns weights/importance to different types of interactions.

PCA is a statistical procedure that uses orthogonal transformation to convert a set of observations of possibly correlated variables into a set of values of uncorrelated ones. One objective of PCA is to find a small set of linear combinations of the variables (interaction data), so that the compounded variables (compounded interaction data) are not correlated and thus avoid the multicollinearity problem. In *itrust*, PCA is applied to extract the internal structure or

feature of normalized 'interaction matrix.' These features are independent and represent a user's social interactions as well. Extracted features are actually the combination of different types of interactions with different weights. Certainly, the features extracted from a user's 'interaction matrix' should be the same for all of his/her friends. By analyzing the features of all users in our dataset, we discover that six compounded interactions could represent 95% of the original 'interaction vector.' Therefore, we use those six compounded interactions with corresponding weights to adjust the number of a user's outgoing data towards each of his/her friends. Finally, a user's friends are ranked based on the amount of his/her outgoing data transformed by PCA, i.e., the more outgoing data, the higher the trust levels.

Evaluation Results

After *itrust* finishes data normalization and compute trustworthiness, a separate page is displayed to allow a user to evaluate the trustworthiness of his/her friends by dragging a sliding bar ranging from 0 to 100. We notice that users are often uncertain about how to translate subjective and multidimensional feelings about interpersonal trusts to a pre-labeled and linear scale. In addition, individual interpretations of interpersonal trust vary, so users are aware that accurate trustworthiness values are not required. However, the aggregate values consistently indicate the relative differences of interpersonal trust between his/her friends. Through this, we obtain the ground truth of a user's friends ranking based on their trustworthiness.

Ranking Accuracy

The trustworthiness of friends on the top and bottom of the ranking list are usually more important than those in the middle. This is because most applications tend to make use of trustworthy friends (e.g., to download files in P2P network) and avoid untrustworthy friends (to buy a product they recommended in e-commerce). Due to the above-mentioned reason, we first examine how accurate the trustworthiness is for those on the top and bottom 20% percentiles of the ranking list.

$$\mathrm{S} = \frac{x + y}{0.4|F|}$$

With a friends ranking list generated by *itrust*, we use *x* (and y) to denote the number of friends appearing on the top 20% (and bottom 20%) on the list. $|F|$ is the total number of the user's friends. Based on the above equation, ranking accuracies of different methods are shown in Fig. 4. The figure indicates that *itrust* provides more accurate ranking results on both top and bottom of the ranking list.

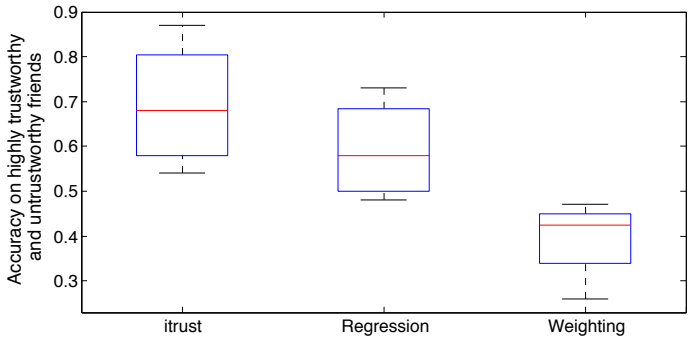

Fig.4 *itrust* accurately discovers highly trustworthy and highly untrustworthy friends

Ranking Evaluation Methods

The above-mentioned method only considers trustworthiness of friends on the top and bottom percentiles of the ranking list; it is necessary to evaluate the ranking accuracy for every friend on the list. To achieve this goal, we introduce the Kendall's tau method to quantify the difference of the ranking generated by users and that computed by *itrust*. Kendall's tau method is a well-recognized approach to compare two rankings. It uses the number of pair-wise disagreements to indicate the difference between two rankings. The smaller the difference, the more similar the rankings are. Fig. 5 shows an example where the ground-truth ranking is 'ABCD', and two special cases are 'ACDB' and 'BADC', respectively. According to the Kendall's tau method, both case 1 and 2 have the same number of pair-wise disagreements, i.e., two disagreements (BC and BD) in case 1, and two (AB and CD) in case 2. Fig. 6(a) shows the Kendall's tau coefficients between the ground-truth and the rankings generated by *itrust*, regression and weighting approaches, respectively. As shown in Fig. 6(a), *itrust* obtains the most accurate ranking results among these three methods, i.e., the average tau of *itrust* is 83%.

Fig.5 Illustrations of the Kendall's tau and generalized Kendall's tau methods

Kendall's tau fails, however, to take into account the importance of different pair-wise disagreements, which is critical for evaluating the accuracy of a ranking list. To model the

importance of each pair-wise disagreement, we adopt the Generalized Kendall's tau method. Generalized Kendall's tau (gtau) considers elements' weight, position weight, and trustworthiness similarities when it evaluates the difference between two rankings. Unlike the Kendall's tau, which counts the proportion of pair-wise agreements, gtau computes the sum of weighted pair-wise disagreements. Therefore, the Kendall's tau method gives results ranging from 0 to 1 while gtau returns results within various ranges, which highly depends on element and position weight, similarity values, and the size of the ranking list. In summary, the larger the results computed by gtau (or the smaller the results generated by Kendall's tau), the more similar the two ranking lists are.

At a glance of Fig. 5, the ranking error caused by swapping A and B should be bigger than that between B and C because many applications are only interested in trustworthy (or untrustworthy) information. Therefore, we add higher weights to the elements on top and bottom of a ranking list but lower weights to those in the middle. Moreover, ranking error caused by swapping B and D should be larger than that of C and D because B and D are farther apart than C and D. Finally, the ranking error caused by swapping A and B should be bigger than that of B and C because the trust values of A and B are more similar than B and C. In other words, both A and B are very close friends of the user but C is only an acquaintance.

Based on the generalized Kendall's tau method, we assign the element and position weights based on the standard normal distribution and compute trust similarities based on the trustworthiness values provided by the user. Evaluation results are shown in Fig. 6(b), which indicates that *itrust* offers much better ranking results compared to the other two methods.

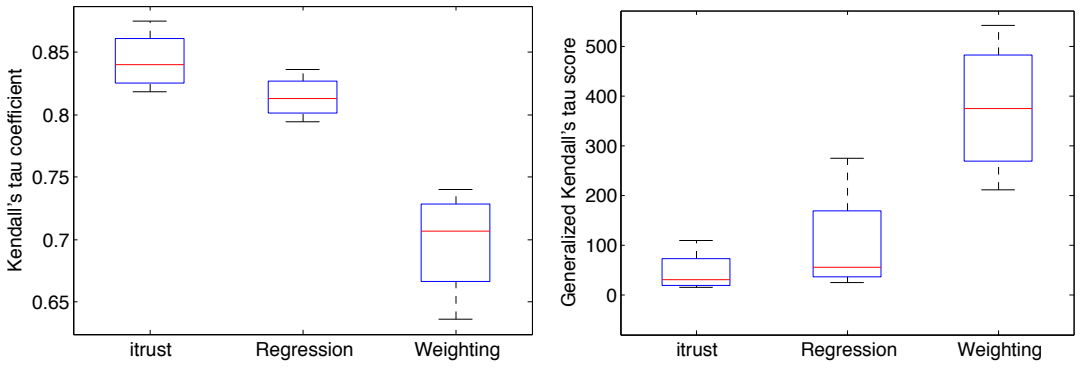

(a) (b)

Fig.6 Evaluations based on Kendall's tau and generalized Kendall's tau

Conclusion

In this article, interpersonal trust relationships between Facebook users are measured by analyzing users' online social interactions. We use normalized outgoing interaction frequency of a user to model the trustworthiness of his/her friends. With the PCA method, features of this

user's outgoing data are then extracted. Finally, a ranking list of the user's friends is generated. Evaluations show that *itrust* provides more accurate trust ranking lists than existing methods. We believe this work contributes to the understanding of interpersonal trust in OSNs. Although our work provides promising results on interpersonal trust measurements in OSN, there is still much work ahead of the research community. Specifically, larger datasets with more Facebook users need to be collected to further evaluate the performance of *itrust*. Moreover, whether *itrust* is applicable to other types of OSN, e.g., Twitter or LinkedIn, is still an open research issue.

References

1. Qing Yang, Honggang Wang. "Towards trustworthy vehicular social networks," IEEE Communication Magazine, 2015, Accepted.

2. Guangchi Liu, Qing Yang, Honggang Wang, Xiaodong Lin and Mike P. Wittie. "Assessment of multi-hop interpersonal trust in social networks by Three-Valued Subjective Logic", IEEE INFOCOM, pp. 1698-1706, 2014

3. Guangchi Liu, Qing Yang, Honggang Wang, Shaoen Wu and Mike P. Wittie. "Uncovering the Mystery of Trust in A Online Social Networks", IEEE CNS, 2015, Accepted.

4. Dijiang Huang; Arasan, V., "On Measuring Email-Based Social Network Trust," IEEE GLOBECOM, pp.1-5, 2010.

5. Podobnik, Vedran, Striga, D., Jandras, A and Lovrek, I. "How to calculate trust between social network users?" IEEE SoftCOM, pp. 1-6, 2012.

6. Singh, Thomas B. "A social interactions perspective on trust and its determinants." Journal of Trust Research 2.2 (2012): 107-135.

7. Vishwanath, Arun. "Manifestations of interpersonal trust in online interaction: A cross-cultural study comparing the differential utilization of seller ratings by eBay participants in Canada, France, and Germany", New Media & Society 6.2 (2004): 219-234.

8. J. P. Onnela, J. Saram¨aki, J. Hyv¨onen, G. Szab´o, D. Lazer, K. Kaski, J. Kert´esz, A. L. Barab´asi. "Structure and tie strengths in mobile communication networks", Proceedings of the National Academy of Sciences, 104.18 (2007): 7332-7336.

9. Gilbert, Eric, and Karrie Karahalios. "Predicting tie strength with social media", ACM SIGCHI, pp. 211-220, 2009.

10. Tinghuai Ma, Jinjuan Zhou, Meili Tang, Yuan Tian, Abdullah AL-DHELAAN, MZNAH AL-RODHAAN, and SUNGYOUNG LEE, "Social network and tag sources based augmenting collaborative recommender system," IEICE transactions on Information and Systems, vol. E98-D, no.4, pp. 902-910, Apr. 2015.

Biography

Xiaoming Li received his M.S. degree in Computer Science from Montana State University in 2014. He is currently a research associate in Nanyang Technological University. His research interests include online social networks and mobile computing.

Qing Yang, Ph.D, is a RightNow Technologies Assistant Professor in the Department of Computer Science, Montana State University. He received B.S. and M.S. degrees in Computer Science from Nankai University and Harbin Institute of Technology, China, in 2003 and 2005, respectively. He received his Ph.D degree in Computer Science from Auburn University in 2011. His research interests lie in the areas of wireless vehicular networks, network security, and trust in online social networks.

Xiaodong Lin received the Ph.D. degree in information engineering from Beijing University of Posts and Telecommunications, Beijing, China, in 1998 and the Ph.D. degree in electrical and computer engineering from the University of Waterloo, Waterloo, ON, Canada, in 2008. He is currently an Assistant Professor of information security with the Faculty of Business and Information Technology, University of Ontario Institute of Technology, Oshawa, ON, Canada. His research interests include wireless network security, applied cryptography, computer forensics, and anomaly-based intrusion detection.

Shaoen Wu received a Ph.D. in computer science in 2008 from Auburn University. He is presently an assistant professor with the Department of Computer Science at Ball State University. He has been an assistant professor in the School of Computing at the University of Southern Mississippi, a researcher scientist at ADTRAN Inc., and a senior software engineer at Bell Laboratories. His current research is in the areas of cyber security, wireless networking, cloud computing, and mobile computing.

Mike P. Wittie is a RightNow Technologies Assistant Professor of Computer Science at Montana State University. He received his PhD in Computer Science from the Computer Science Department at the University of California, Santa Barbara. He graduated from the University of Pennsylvania with an MSE in Computer Science and a BA in Cognitive Science. His research interests lie in networking problems related to dynamic content applications and soft real-time services.