Social Information Leakage: Effects of Awareness and Peer Pressure on User Behavior

Mariam Nouh[1], Abdullah Almaatouq[1], Ahmad Alabdulkareem[1], Vivek K. Singh[2], Erez Shmueli[2], Mansour Alsaleh[1], Abdulrahman Alarifi[1], Anas Alfaris[1,2], and Alex 'Sandy' Pentland[2]

[1] Center for Complex Engineering Systems (CCES), King Abdulaziz City for Science and Technology (KACST), Riyadh, Saudi Arabia
{mnouh,aalmaatouq,aabdulkareem,maalsaleh,aarifi}@kacst.edu.sa
[2] Massachusetts Institute of Technology (MIT), Cambridge, MA, USA
{singhv,shmueli,anas,pentland}@mit.edu

Abstract. Today, users share large amounts of information about themselves on their online social networks. Besides the intended information, this sharing process often also "leaks" sensitive information about the users - and by proxy - about their peers. This study investigates the effect of awareness about such leakage of information on user behavior. In particular, taking inspiration from "second-hand smoke" campaigns, this study creates "social awareness" campaign where users are reminded of the information they are leaking about themselves and their friends. The results indicate that the number of users disallowing the access permissions doubles with the social awareness campaign as compared to a baseline method. The findings are useful for system designers considering privacy as a holistic social challenge rather than a purely technical issue.

Keywords: Social information leakage, Online social networks, Privacy, Peer pressure.

1 Introduction

The growing popularity of Online Social Networks (OSNs), such as Facebook, Twitter and LinkedIn, has made them integral parts of contemporary online activities. Although OSNs are widely used and represent a rich source of information, much of their data is also sensitive and personal (e.g., demographic, interests, etc.) [1]. OSN users usually disclose such personal information in order to participate in social communities or in return of services [2,3]. However, disclosing personal information in this case can be a double-edged sword. For example, such exposure might make the user vulnerable to personalized attacks such as stalking, identity theft, reputation slander, personalized spamming and phishing. While most of the OSN services offer various levels of privacy protection (e.g., allowing only authorized list of other OSN users, applications and

T. Tryfonas and I. Askoxylakis (Eds.): HAS 2014, LNCS 8533, pp. 352–360, 2014.

third parties, etc.), users' information may extend beyond the defined bounds, which in a privacy context is referred to as information leakage [4].

Information leakage is the phenomena where explicit information provided to a third party can be used to derive implicit and previously hidden information about an entity. Many of the literature suggest some reactive measures to be taken to minimize the effects of information leakage. Such measures include suggesting some friends to un-friend in order to minimize the amount of leakage [5]. In this study, we would like to address this issue in a preemptive way rather than in a reactive way. Since the user is the entity in charge of such decision, we would like to test whether informing the user before sharing personal information can minimize this effect preemptively. Additionally, we examine the extent to which peer pressure influences user's behavior. We present the design of our experiment as well as the results and drawn conclusions. The experiment goal is to investigate how different users behave when they know that they are leaking sensitive information about themselves or their peers and how this affects their decisions. Our study was conducted on an online social network platform with around 200 participants. The results show that users were more responsive to the peer pressure variable, and thus, gives an indication that if users consider their peers when making their online privacy decisions they will most probably leak less information and increase their privacy level.

2 Related Work

Social information leakage (SIL) in OSN has different types and forms. Information leakage may occur from a user's network (friends) to the user, from a user to himself, or from a user to his network. The first has been heavily studied in the literature [6,7]; mainly it occurs by correlating and aggregating information from a user's network to reveal sensitive information about the user. The second is when a user shares different pieces of information (attributes) in one or more OSNs. By correlating these pieces of information one can construct a user's social print, which may pose several threats to the user's privacy and aids in launching several targeted attacks (e.g., social engineering, password recovery) [8]. The last type, as yet understudied scenario, is when a user leaks information about his network (friends) by explicitly sharing his own information.

As individual decisions are known to be sub-optimal in social settings [9], we need to provide incentivization as seen in [10] or peer-pressure mechanisms as seen in [11,12] in order to nudge users towards better social outcomes.

Moreover, previous research showed that relying on traditional ways for managing privacy, such as textual privacy management settings, proven to be inconvenient. Either due to the user's inexperience in dealing with these settings, or due to the high complexity of the privacy settings. Lipford et al. [13] performed an experiment to study this issue. They designed a privacy management interface focused on showing audience point of view. Their results showed that providing visual feedback of the outcome of privacy settings can improve users' understanding of their privacy and help them make more accurate decisions.

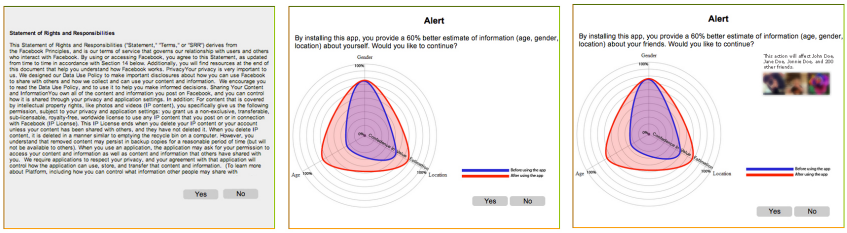

Fig. 1. Types of SIL Messages

In designing our experiment, we incorporate visual elements as well as peer-pressure mechanisms in order to study their effects on users' privacy decisions.

3 SIL Experiment Design

The main question that this study aims to answer is whether users knowing the implication of sharing pieces of information may make them change their behavior. We design our experiment by defining the hypothesis that we want to test, how participants are allocated to different groups, who are our targeted population, and what are the evaluation metrics that we want to measure. Specifically, in this study we try to answer the following questions:

1. If the user is presented with a numerical quantification of the amount of information leaked, does this affect his/her behavior?
2. What are the differences in user behavior when informed about leaking information about themselves, as opposed to leaking information about their peers.

This is especially interesting because: a) it grants the user a sense of 'agency' and b) it brings out the effect of direct and indirect peer pressure on user behavior. Previous results in smoking campaigns (e.g. second hand smoke affects your dear ones) as well as healthy behavior adoption [14] have suggested an impact of social peer pressure on user behavior. This affect is as yet not studied or quantified from a behavioral privacy aspect.

3.1 User Groups

In order to test our hypothesis we design the experiment such that users are randomly assigned to one of three groups. The social information leakage message consists of three main components: text message, visual message, and social message. Each group is presented with a different combination of these components (see Figure 1).

Control Group: This group is presented with only a text-based message shown as a typical terms and conditions page. This group acts as our control group (baseline). Users have two options either to accept and proceed to the app page, or decline and exit the app.

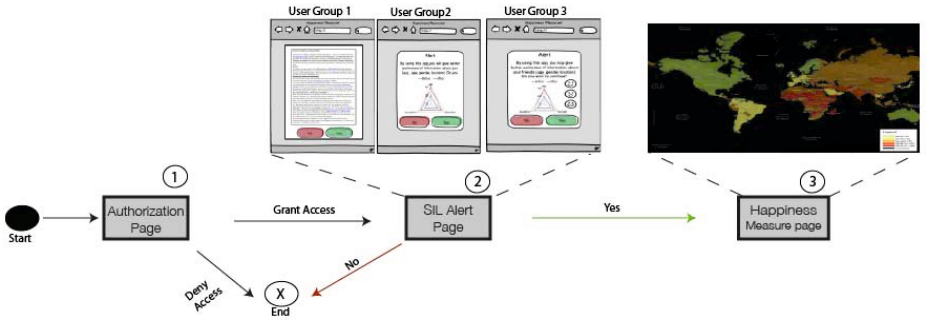

Fig. 2. The Sequence of Steps

User to User Group: This group is presented with both a text and a visual message. The text message states that using the app will result in leakage of some of the user's sensitive information. A spider graph is shown to visually emphasize the before/after effect on leaking the user's information. Users have two options either to accept and proceed to the app page, or decline and exit the app.

User to Network Group: This group is presented with all three components of the SIL alert message (i.e., text, visual, and social messages). The text message states that using the app will result in leakage of sensitive information about the user's friends. Similar to the previous group a spider graph is shown to visualize the amount of leakage. Additionally, profile pictures of user's friends are pulled and displayed as part of the social message to add the social pressure and test the peer effect on user's behavior. Users have two options either to accept and proceed to the app page, or decline and exit the app.

3.2 Population and Metrics

Our goal is to have population representation from different age groups, gender, and ethnic backgrounds. For each participants, we measure two variables: (1) user's action when presented with the information leakage alert message. (2) How much time it takes the user to respond to the message.

4 Experiment Setup

Due to it's popularity, we chose Facebook as a platform to conduct our experiment. We built a Facebook app and assigned each user who participated in the experiment to one of the three user groups presented earlier. The Facebook app is called Happiness Measure, which shows how happy the user is according to his/her current location. The app presents the users with a heat map of the world's happiest countries according to the 2013 World Happiness Report [15].

The actual functionality of the app was selected for simplicity and mainly to attract users to participate.

When the user first connects to the app, he/she is presented with an authorization request to allow the application access to his/her basic profile information (e.g., name, age, location, etc.). If the user grants the application access to her information, she is counted as part of our experiment and proceeds to the SIL alert page, otherwise their information is discarded. At this stage we measure the time it takes the user to respond to the SIL message (i.e., decision time), as well as record his response to the message (i.e., decision). Figure 2 shows the sequence of steps the users follow in the experiment.

For the sake of data collection, we utilized the Facebook advertisement services to promote the app and encourage people to participate in the study. The data collection started from November 23rd to December 5th, 2013. The advertisement campaign targeted 30 countries from different continents to allow for a diverse set of participants. A list of the targeted countries is shown in Table 1.

5 Experiment Results

We divided our analysis to two main parts. First is the decision analysis, where we look at the different decisions users made in each group. Second, we study the time factor to know how much time elapsed before users reached their decisions. The application received around 300 users' clicks, of those around 200 users completed the experiment while the rest decided to close the application before answering the SIL message. In this section, we analyze and discuss the results of those users, and we aim to answer the two questions stated previously in Section 3.

5.1 Decision Analysis

Table 2 summarizes the number of users per group together with their decisions. Decision analysis focuses on studying the differences between each group in terms of user's response to the SIL message. In Group 1, 91% of the users agreed to the terms and conditions page and thus responded with Yes to grant the app access, while 9% responded with No. Similarly, in Group 2 90% of the users responded with Yes knowing that the app will leak some of their personal information and only 10% responded with No. However, in Group 3 when presented with their friends' pictures 20% of the users decided to deny the application access and

Table 1. List of Targeted Countries

Tunisia	Algeria	Brazil	Canada	Chile	China	Egypt
India	Iraq	Mexico	Morocco	Pakistan	Qatar	Switzerland
Turkey	Colombia	UK	Jordan	Russia	Italy	France
UAE	USA	Saudi Arabia	Greece	Germany	Ghana	Slovakia

Table 2. Data Description

Groups	Num. of Users	Yes	No
Group1: Control Group	82	91%	9%
Group2: User to User Group	63	90%	10%
Group3: User to Network Group	54	80%	20%

Table 3. Median Time in Seconds

Groups	Median Yes	Median No	Median Total
Group1: Control Group	6s	8s	6s
Group2: User to User Group	9s	6s	8s
Group3: User to Network Group	9s	13s	10.5s

responded with No. The results show that the third group behaves differently than the other two groups. Thus, this indicate that peer pressure have an effect on user's decisions. Figure 3 shows the percentages of each decision per user group.

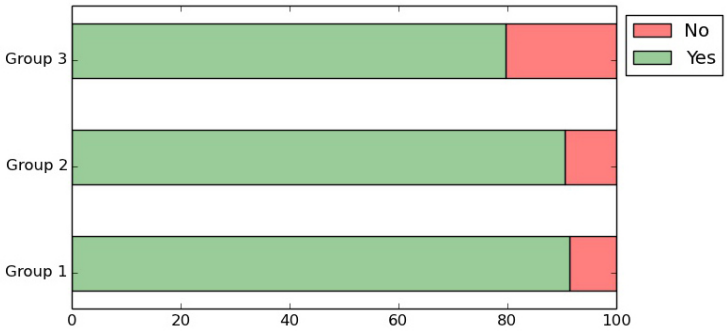

Fig. 3. Decisions per User Group

5.2 Decision Time Analysis

Decision time analysis focuses on studying the differences between user groups in terms of the time spent to make a decision on the SIL message. Again the results show that users in Group 3 behaved differently than users in the other two groups. The median total time elapsed to make a decision for users in Group 3 was 10.5 seconds (either yes or no), 8 seconds for users in Group 2, and 6 seconds for users in Group 1 (See Table 3). Moreover, we studied the behavior of each group based on the type of decision they took (i.e., Yes or No decision). Again, the analysis showed that Group 3 spent more time to make a decision than the other two groups. In the case of No decision, the median time of users in Group 2 was significantly less than the other two groups. Figures 4 and 5 represent a

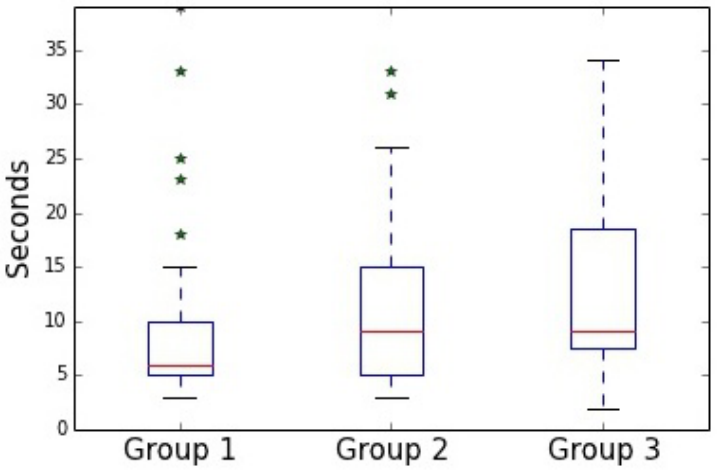

Fig. 4. Time to Take the "Yes" Decision

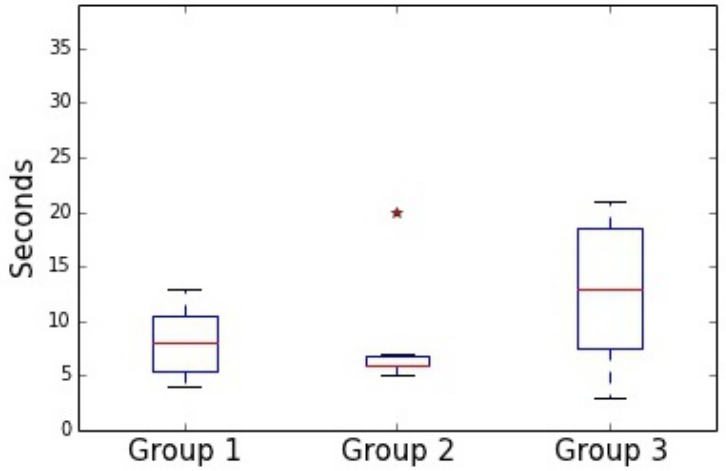

Fig. 5. Time to Take the "No" Decision

box plot of the time spent by each group to make a Yes and No decisions. The box plot representation gives a good overview of the distribution of the data, the median, minimum, maximum, and outliers. The box contains 50% of the data points, with the red line inside the box represent the median time for each experiment group.

6 Discussion

We can see that in both kinds of analysis Group 3 behaved differently than the other two groups. Additionally, unlike our initial expectation Group 1 and Group 2 were fairly similar in making their decisions. Thus, the quantification of information leaked seems not to have the major effect on user's decisions and behavior which answers the first question in our hypothesis. Moreover, when users are informed that they are leaking sensitive information about themselves they were more willing to proceed with the app as opposed to when their friends or peers were affected. This behavior was consistent in both the decisions they took and in spending more time to think when they were aware that their decision will affect their friends. This observation was also consistent among subpopulations. We looked at two demographic features, age and gender. We analyzed how each subpopulation responded to the experiment and the results were fairly consistent. Thus, this suggests that users intend to leak less information when they consider the effects on their peers.

7 Conclusion

In this paper we presented a user study on how knowing that you are compromising your friends' privacy may affect your decision-making. We discussed the experiment design, procedure, and results. Our analysis showed that peer pressure has influence on users' behavior. As a future work, we intend to study further what made Group 3 behave differently by conducting a second experiment with variations of Group 3 SIL message components. To the best of our knowledge, this is the first attempt to study peer pressure and its effects on users' behavior with regard to privacy. We hope the presented results will bring useful insights for policy and application design of future social applications with respect to privacy.

Acknowledgments. The authors would like to thank King Abdulaziz City for Science and Technology (KACST) for funding this work. In addition, the authors thank the Center for Complex Engineering Systems (CCES) at KACST and MIT for their support.

References

1. Chaabane, A., Kaafar, M.A., Boreli, R.: Big friend is watching you: Analyzing online social networks tracking capabilities. In: Proceedings of the 2012 ACM Workshop on Workshop on Online Social Networks, WOSN 2012, pp. 7–12. ACM, New York (2012)
2. Lampe, C.A., Ellison, N., Steinfield, C.: A familiar face(book): Profile elements as signals in an online social network. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2007, pp. 435–444. ACM, New York (2007)

3. Strater, K., Lipford, H.R.: Strategies and struggles with privacy in an online social networking community. In: Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction, BCS-HCI 2008, vol. 1, pp. 111–119. British Computer Society, Swinton (2008)
4. Becker, J., Chen, H.: Measuring Privacy Risk in Online Social Networks. In: Proceedings of W2SP 2009: Web 2.0 Security and Privacy (May 2009)
5. Heatherly, R., Kantarcioglu, M., Thuraisingham, B.: Preventing Private Information Inference Attacks on Social Networks. IEEE Transactions on Knowledge and Data Engineering 25(8), 1849–1862 (2013), doi:10.1109/TKDE.2012.120
6. Blenn, N., Doerr, C., Shadravan, N., Van Mieghem, P.: How much do your friends know about you?: reconstructing private information from the friendship graph. In: Proceedings of the Fifth Workshop on Social Network Systems (SNS 2012), Article 2, 6 p. ACM, New York (2012), http://doi.acm.org/10.1145/2181176.2181178, doi:10.1145/2181176.2181178
7. Singla, P., Richardson, M.: Yes, there is a correlation: - from social networks to personal behavior on the web. In: WWW 2008: Proceeding of the 17th International Conference on World Wide Web, pp. 655–664. ACM, New York (2008)
8. Irani, D., Webb, S., Pu, C., Li, K.: Modeling unintended personal-information leakage from multiple online social networks. IEEE Internet Computing 15(3), 13–19 (2011)
9. Hardin, G.: The Tragedy of the Commons. Science 162(3859), 1243–1248 (1968)
10. Singh, V.K., Jain, R., Kankanhalli, M.S.: Motivating contributors in social media networks. In: Proceedings of the First SIGMM Workshop on Social Media, WSM 2009, pp. 11–18. ACM, New York (2009)
11. Mani, A., Rahwan, I., Pentland, A.: Inducing Peer Pressure to Promote Cooperation. Sci. Rep. 3, 1735 (2013)
12. Bond, R.M., Fariss, C.J., Jones, J.J., Kramer, A.D., Marlow, C., Settle, J.E., Fowler, J.H.: A 61-million-person experiment in social influence and political mobilization. Nature 489(7415) (September 2012)
13. Lipford, H.R., Besmer, A., Watson, J.: Understanding privacy settings in facebook with an audience view. In: UPSEC. USENIX Association (2008)
14. Aharony, N., Pan, W., Ip, C., Khayal, I., Pentland, A.: Social fmri: Investigating and shaping social mechanisms in the real world. Pervasive Mob. Comput. 7(6), 643–659 (2011)
15. John Helliwell, R.L., Sachs, J.: World Happiness Report (2013), http://unsdsn.org/files/2013/09/WorldHappinessReport2013_online.pdf