What Does Your Profile Really Say About You?: Privacy Warning Systems and Self-disclosure in Online Social Network Spaces

Lia Emanuel

CREATE Lab

Department of Psychology

University of Bath

Bath, UK BA2 7AY

L.Emanuel@bath.ac.uk

Chris Bevan

CREATE Lab

Department of Psychology

University of Bath

Bath, UK BA2 7AY

crb23@bath.ac.uk

Duncan Hodges

Cyber Security Centre

Department of Computer Science

University of Oxford

Oxford, UK OX1 3QD

duncan.hodges@cs.ox.ac.uk

Abstract

This paper reports current progress on the design and initial evaluation of an innovative privacy feedback system aimed to provide social network users with tailor-made feedback about their identity exposure online. Preliminary results suggest our feedback system, based on a research driven model of identity, appears to reduce the amount of information individuals disclose about themselves in social network spaces. We discuss the impact of our feedback system on the way individuals share information online, as well as suggestions for a more fine-grained evaluation and future development of this feedback system.

Author Keywords

Online privacy; identity information; social network sites; feedback systems.

ACM Classification Keywords

H.5.m Miscellaneous

General Terms

Human Factors, Design, Security

Introduction

Nearly half of the top ten sites as ranked by Alexa are social networks or have social network-like elements within them. Many individuals maintain multiple accounts on different sites, as different sites satisfy a subtly different interaction experience. Whether it is for professional networking or finding romance, the underlying context of a social network can influence the type of information users disclose [1].

There has been significant academic research on the privacy and security concerns with data exposure on social network sites (SNS) e.g. [6,7]. More interestingly, the online community is taking steps to highlight the potential vulnerabilities introduced by overly disclosing on SNS; sites such as pleaserobme.com, weknowwhatyouredoing.com and the @needadebitcard twitter feed are proof of both the presence, and prevalence, of this over-disclosure.

In addition to the security implications, the use of large amounts of publicly disclosed 'private' data for marketing and advert targeting is often overlooked by the public, yet arguably is more likely to impact them. Although feedback systems exist to warn users of general information tracking of this type (e.g. Ghostery) and can reduce the amount of information a user provides [8], these systems primarily raise privacy awareness around marketing and e-commerce actions. SNS, on the other hand, are distinctly different. They are interactional in nature, whereby information is shared with an online community to exchange opinions, beliefs and activities [2]. For this reason, we believe privacy warning systems need to adapt to the way in which users interact and exchange information within SNS.

Social networking users are providing more content than ever before [4]. This content is, by definition, largely personal in nature and there is very little to provide users with privacy advice. Indeed, the business-models that enable most social networks rely on encouraging large-scale sharing of personal data. We propose introducing a higher-level privacy warning system. Rather than a notification warning users that they are being tracked or providing a generalized warning of 'high-risk' information disclosure, we aim to tailor the warning feedback to the individual and the specific information being disclosed.

This project focuses on user's disclosure behavior in two different online social network contexts, namely a dating network and a professional network. This paper reports progress so far on the development of our privacy warning system, and preliminary data on the effect of this warning system on user's disclosure behavior within the two SNS using experiment design methods. In the evaluation of the warning system we had two main research questions:

RQ1: Prior to any warning system intervention, do users tend to disclose information differently depending on the context of the social network?

RQ2: Following tailored feedback about one's potential online exposure are users more stringent with the information they disclose, as compared to those receiving no feedback? Does context play a role in the effect of the feedback system?

Our goal is that the implementation of this warning system will benefit the social network user by providing relevant, real-time, feedback about their identity
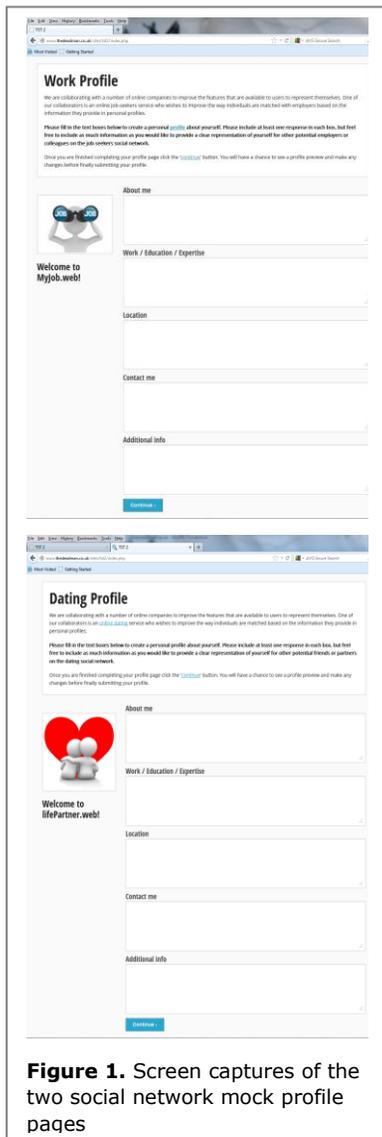
Figure 1. Screen captures of the two social network mock profile pages

exposure online and allow them to make an informed decision about the implications of that exposure with regards to their privacy.

Development of the privacy feedback system

In order to help model identity we considered the facts about an individual as nodes – these facts could be things like email addresses, job descriptions, names etc. It is often possible to infer new facts from the current set; these inferences or transforms can be modeled as new links between nodes. A more complete description of this model, known as the SuperID Model [5], can be found in e.g. [3].

For this study, we used 37 possible input facts, which given the state-of-the-art transforms at the time of writing [3] resulted in 52 possible new output facts. The model is held on a server, which provides an API allowing others to query what new facts can be inferred from the current set. This model can easily be updated as new inferences are discovered. A plugin installed in the user's browser searches each page for a form where a user submits personal data (e.g. when they were signing up for a new service). As the user completes the form, the plugin identifies certain types of fact, queries the server with the set of facts that are present and highlights to the user what extra information could be derived. It is important to note that simply the presence of a fact is sent to the server, **not** the value of the fact (i.e. the presence of a name is sent not that the name was Alice).

Testing the privacy feedback system

In order to explore the impact of this type of feedback system on user's disclosure behavior, mock social network profile pages were created that could

incorporate the privacy feedback plugin described. The mock profiles showed a template for a fictitious dating social network and a professional, job-seekers network profile (Figure 1). The profile had five themed text boxes: About me, Work/Education/Expertise, Location, Contact/Getting in Touch and Additional Information, for users to fill in.

Preliminary data from 41 participants (12 males, 29 females; age M = 23.02, SD = 6.20) has been collected. Participants were recruited from online University notice boards, and were randomly assigned to create either the dating social network profile (N=23) or the professional social network profile (N=18). Participants were asked to create a personal profile page using the five text boxes under the pretense that they were aiding online companies to improve their profile features. Once participants created and submitted their initial profile information, it was manually coded by classifying profile information into facts about the individual. These facts were submitted to the server to determine what information could be derived about the participant, given the information they had provided in their profile. While profiles were being manually coded participants were asked to complete a demographics questionnaire.

Participants then reviewed their profile before submitting it to go 'live' on the social networks. Approximately half of all participants, through random allocation, received feedback from the privacy warning system during the review of their profile. Feedback included each fact given by the participant, how it was identified by the server and what additional information could be derived by other users on the social networks from that fact if it were made available online.

Figure 2 shows an example of the feedback a participant could receive about their profile page. During the review of their profile, participants could edit information and were asked to set a privacy level for each of the five text boxes individually, such that the information in a text box could be viewed by: Only You, Approved Contacts, Groups/Networks You Join, Everyone.
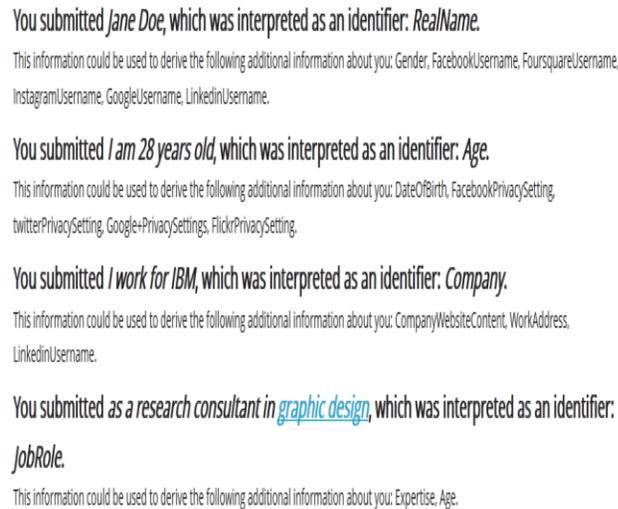
You submitted *Jane Doe*, which was interpreted as an identifier: *RealName.*
This information could be used to derive the following additional information about you: Gender, FacebookUsername, FoursquareUsername, InstagramUsername, GoogleUsername, LinkedinUsername.

You submitted *I am 28 years old*, which was interpreted as an identifier: *Age.*
This information could be used to derive the following additional information about you: DateOfBirth, FacebookPrivacySetting, twitterPrivacySetting, Google+PrivacySettings, FlickrPrivacySetting.

You submitted *I work for IBM*, which was interpreted as an identifier: *Company.*
This information could be used to derive the following additional information about you: CompanyWebsiteContent, WorkAddress, LinkedinUsername.

You submitted *as a research consultant in graphic design*, which was interpreted as an identifier: *JobRole.*
This information could be used to derive the following additional information about you: Expertise, Age.

Figure 2. Example of privacy warning feedback about a user's information exposure

Results

Initial profile submitted

We were interested in exploring whether individuals disclose information differently depending on the social network context, prior to any type of privacy feedback. Across both network contexts, 23 of the possible 37 input facts were present in the initial profiles. There

was no difference in the number of facts disclosed in participants dating profiles ($M = 7.22$, $SD = 2.13$) and participants professional profiles ($M = 8.39$, $SD = 2.70$), $t(39) = -1.55$, $p = .13$.[1] To explore the content of the information disclosed, Table 1 illustrates the proportion of each fact type present in the dating network profiles and the professional network profiles.

Fact Type	Dating	Professional
Personal Information	*63.3%*	*40.8%*
Name	6.1%	4.1%
Age	7.9%	4.1%
Gender	1.8%	0.7%
Location	15.9%	12.9%
Nationality/Birth place	2.4%	2.0%
Language	0.6%	0.7%
Interests	9.1%	6.8%
Personality*	19.5%	9.5%
Work Background	*27.3%*	*42.2%*
Job Role	12.2%	12.9%
Education	8.5%	10.2%
Resume	1.8%	11.6%
Expertise	3.0%	3.4%
Company	1.8%	3.4%
Work Address	0.0%	0.7%
Contact Information	*9.1%*	*17.0%*
Email address	8.5%	12.2%
Phone number	0.6%	4.8%

Table 1. Proportion of fact types submitted in initial profiles by dating and professional context. * Seven facts related to personality traits were collapsed.

---

[1] Exploratory analysis indicated there was no relationship between the number of facts disclosed and reported occupation status, $r = -.05$, $p = .75$, relationship status, $r = -.12$, $p = .46$, and gender, $r = .02$, $p = .91$.

Profile review and privacy system

To explore *RQ2* we assessed the impact of the feedback system on information disclosure in two ways, the amount of facts present in the reviewed profiles and the privacy settings selected by participants. Changes participants made to their profile during the review period were considered in terms of facts that were altered (i.e., either edited or deleted) or facts that were added to their initial profile information. Overall, only 7.7% of facts present in the initial profiles were altered. The greatest proportion of facts altered were in the work background (45.8%), followed by personal information (41.7%) and contact information (12.5%) fact types. Notably, only participants who received no feedback via our warning system provided additional information during the review of their profile (Table 2). This group added personal information (62.5% of added facts) and contact information (37.5% of added facts).

Feedback	No changes	Added Facts	Deleted Facts
No Feedback	84%	4%	12%
Feedback given	81%	0%	19%

Table 2. Percentage of participants who made no changes, added facts and deleted facts in the review of their profile by privacy system feedback.

In order to explore the effect of the feedback system on the number of facts present in participant's final profile, and the role of the SNS context, a 2 (privacy system: no feedback vs. feedback given) x 2 (context: dating vs. professional) ANCOVA was run. The number of facts submitted in participants' initial profile was used as a covariate. Results showed, after controlling for the number of facts in the initial profile, a marginal effect of privacy system, $F (1,36) = 3.05$, $p = .09$, $\eta^2 = 0.08$. Those given feedback on their initial profile submitted marginally fewer facts ($M = 6.83$) than those who received no feedback ($M = 8.43$) about their potential online exposure. However, there was no effect of context, $F (1,36) = 0.87$, $p = .36$, $\eta^2 = 0.24$, nor was there an interaction between privacy system and context, $F (1,36) = 0.25$, $p = .62$, $\eta^2 = 0.01$.

The privacy settings selected by participants were coded as 1=Only You; 2=Approved Contacts; 3=Joined Groups and 4=Everyone; for the accessibility of each text box within the social network. These settings were averaged across the five text boxes for each participant's profile.

Feedback	Dating	Professional	Total *M*
No Feedback	2.10	2.62	2.39
Feedback Given	2.26	2.50	2.34
Total *M*	2.20	2.58	

Table 3. Mean privacy settings collapsed across the five profile text boxes by privacy system feedback and context.

The direction of means listed in Table 3 suggest that those given feedback about their exposure during the review of their profile set their privacy settings slightly more stringently than those given no feedback. In addition, participants in the dating context set their privacy settings more stringently than those in the professional context. However, statistical comparison showed no significant effects of feedback or context on privacy settings selected (all effects, $F (1,37) < 1$).

Conclusions and future work

The results from the initial evaluation of the privacy warning system described suggests a promising start to addressing issues of privacy and security via over-disclosing in SNS. The information provided in our mock social network profiles suggest that different types of facts are typically disclosed across different SNS contexts. By further understanding these variations, we can begin to adapt feedback systems to particular SNS. Furthermore, those who were given privacy feedback relevant to the information they were disclosing submitted fewer identity facts to SNS communities. Whereas, those who receive no feedback actually submit more facts about themselves while reviewing the profile they created. For future work, we aim to run further iterations to evaluate this warning system with larger sample sizes. For instance, a sample of 100 would provide the power necessary to detect a significant effect of the warning system (at 0.83 probability) on the number of facts disclosed given the reported effect size ($\eta^2 = 0.08$). Likewise, further data collection would also allow us to take a more in-depth look at what people are willing to disclose and the type of facts individuals are more likely to alter or protect via more stringent privacy settings. This type of fine-grained analysis as well as collection of user experience measures would support further development of our privacy feedback system. For example, can we identify common feedback, or output facts, that result in more privacy-seeking behavior? Do individual differences in personality and attitudes towards online privacy influence the extent to which users act on minimizing their information exposure following feedback from the warning system? These are only a few of the questions that need to be addressed in order to refine the design and interface of, and eventually fully automating, this privacy warning system. Discussing the potential for this system with the wider HCI community will result in a more efficient and beneficial tool which allows SNS users to make an informed decision about their privacy online.

Acknowledgements

References

[1] Bevan, C., Emanuel, L., Jamison-Powell, S. Neil, G, Stanton-Fraser, D., Stevenage, S., Whitty M. Twenty Statements Test: The role of cyber-physical environment, context and multiple-selves. (*In prep*).

[2] European Commission. Attitudes on data protection and electronic identity in the European Union, *Eurobarometer 359 Special Report*, (2011).

[3] Hodges, D., Creese, S. and Goldsmith, M. A Model for Identity in the Cyber and Natural Universes. *IEEE EISIC,* (2012).

[4] James, J. How much data is created every minute? Retrieved December 13, 2012. http://www.domo.com/blog/2012/06/how-much-data-is-created-every-minute/

[5] SuperIdentity project. Retrieved January 9, 2013. www.superidentity.org

[6] Rose, C. The security implications of ubiquitous social media. *International Journal of Management & Information Systems* 15.1 (2011).

[7] Rosenblum, D. What anyone can know: The privacy risks of social networking sites. *Security & Privacy, IEEE* 5.3 (2007), 40-49.

[8] Zhu, F., Carpenter, S. and Kulkarni, A. Understanding identity exposure in pervasive computing environments. *Pervasive and Mobile Computing* 8.5 (2011), 777-794.