Modeling Unintended Personal-Information Leakage from Multiple Online Social Networks

Most people have multiple accounts on different social networks. Because these networks offer various levels of privacy protection, the weakest privacy policies in the social network ecosystem determine how much personal information is disclosed online. A new information leakage measure quantifies the information available about a given user. Using this measure makes it possible to evaluate the vulnerability of a user's social footprint to two known attacks: physical identification and password recovery. Experiments show the measure's usefulness in quantifying information leakage from publicly crawled information and also suggest ways of better protecting privacy and reducing information leakage in the social Web.

Danesh Irani, Steve Webb, and Calton Pu
Georgia Institute of Technology

Kang Li
University of Georgia

Social networks are of great interest to the research community because more than one-third of American adults on the Internet use them.[1] Furthermore, 51 percent of these users participate in multiple social networks, many of which are built for specialized audiences. Social networks are, by nature, platforms for people to share personal information and experiences with friends and to find new friends with common interests. To participate in a social network, users must create profiles; because accurate profile information allows more efficient networking, users have an incentive to use real information in their profiles.

Typically, social networks target particular interests, such as professional networking at LinkedIn, and users are encouraged to share information about their interests and meet others based on commonalities. Consequently, social networks often contain disparate pieces of information about different personal aspects of a user. In providing this information to a social network, most users assume that their profile information will be kept within the social network's boundaries.

1089-7801/11/$26.00 © 2011 IEEE Published by the IEEE Computer Society

Unfortunately, each social network's privacy guarantee to maintain user profile information within the site doesn't protect users from attackers who combine disparate pieces of information about a user from multiple networks. This often exposes more of the victim's information than from a single social network. We call this problem "unintended personal-information leakage" and define a *social footprint* as the total amount of personal information that can be gathered about an online identity by aggregating available social networks.

We've taken two steps to more rigorously understand the problem of unintended personal-information leakage in social networks. First, we defined a measure to quantify the amount of information in a user's social footprint. Second, using approximately 8,200 users' social footprints, we conducted an empirical study to evaluate the effectiveness of two attacks — an identification attack and a password-recovery attack. We found that the amount of information released for the physical identification attack increased from 34 percent for a user with one social network profile to 90 percent when combining six or more social profiles. Results from our study of information released for the password-recovery attack followed a similar trend.

Dangers of Unintended Personal-Information Leak

We focus on the use of personal information for concrete attacks to illustrate the dangers of unintended personal-information leak. The two attacks we investigated are characterized by a set of personal attributes that enable attackers to deduce sensitive information about the user.

Physical Identification Attack

Personal identification information (PII) is "information which can be used to distinguish or trace an individual's identity."[2] Although most social networks don't store or explicitly reveal PII, many social networks reveal a person's birthdate, gender, or location.

According to a study of US Census data,[3] the attribute set `{Birthdate, Gender, Zip}` poses a risk of personal identification because these attributes can uniquely identify 87 percent of the US population. Additionally, `{Birthdate, Gender, Location}`, where `Location` is the city, town, or municipality in which the person resides, can uniquely identify 53 percent of the US population. Because social networks use `Location` more widely than ZIP code, we define the set of attributes to consist of `{Birthdate, Gender, Location}`.

Password-Recovery Attack

Many websites let legitimate users recover their passwords by providing personal "secret" information (such as birthdate and address). Unfortunately, attackers often infer the answers from other sources and use this recovery mechanism to compromise accounts. Users with publicly available social profiles who provide truthful answers to these questions offer a rich source of information for attackers who can mine their profiles for information. A recent example of this attack was the hijacking of Sarah Palin's Yahoo email account after attackers discovered the answer to her password-recovery question.[4]

The exact information needed to recover passwords varies among sites and is sometimes based on the type of questions a user selects. Thus, it's difficult to define a set of attributes that accurately reflects the information that social profiles leak in relation to password-recovery attacks. Based on common password-recovery questions,[5–7] we define the set of attributes for a password-recovery attack as `{Name, Email, Nickname, Location, Gender, Hometown, Homepage, Birthdate}`.

Online Social Footprints

Figure 1 shows a user named Bob Smith and his online social footprint, which is constructed using profile information from three social networks (each represented by a row in the table). Individually, we can use each social network to discover four to six pieces of information or attributes (such as age and gender). By combining the profile information from these social networks, we can discover eight attributes about Bob, which constitute his online social footprint.

Blurring Social Network Boundaries

Because people often use multiple social network services,[1] we must first identify a user's profiles across distinct social network boundaries. This will let us then fetch information from all of those networks to construct a user's online footprint.

	Name	Location	Sex	Relationship status	Hometown	Homepage	Birthdate	About me
Flickr	Bobby Smith	-	Male	Single	-	-	- / - /1980	-
LiveJournal	Bob Smith	-	-	-	-	http://ww...	01/31/ -	My biogr...
MySpace	Bob Smith	30332	Male	-	Chicago, IL	-	- / - /1980	I like priv...

Figure 1. A user's online social footprint. This footprint is an aggregation of attributes from the user's profile information on multiple social networks.

Although users might not specify explicit links to their other profiles, the de-anonymization of numerous users across social networks has proven successful in the past by using friend-network graphs[8,9] or record-matching techniques.[10] We can also use email address lookups[11] or a user's pseudonyms[12] to achieve targeted de-anonymization.

Defining a Social Footprint

A user's profile τ_s^u is a set of attributes representing information obtained about a user u on social network s. Using Figure 1 as a running example, we can represent Bob's Flickr profile as

$\tau_{\text{Flickr}}^{\text{Bob}}$ = {Name: "Bobby Smith", Sex: "Male", Relationship status: "Single", Birthdate: "-/-/1980"}.

A user's online identity T is a list of profiles τ that represents a user's online identity or persona. In Figure 1, Bob's online identity has three profiles:

$\tau_{\text{Flickr}}^{\text{Bob}}$, $\tau_{\text{LiveJournal}}^{\text{Bob}}$, and $\tau_{\text{MySpace}}^{\text{Bob}}$.

Finally, a union of the profiles in their online identity gives a user's social footprint P^u:

$$P^u = \bigcup_{i \in T} \tau_i^u .$$

A subset of Bob's social footprint can be expressed as P^{Bob} = {Name: "Bobby Smith", Name: "Bob Smith", Name: "Bob Smith", Sex: "Male", . . .}.

Defining Aggregate Attribute Leakage

Attribute leakage is a measure of the information that we can discover about a particular attribute given a user's social footprint. In this case, we simply need to check for the presence of an attribute in a user's social footprint.

Definition 1. $\phi(f_a, P)$ — attribute leakage. Given an attribute name f_a and a user's social footprint P, we define attribute leakage as

$$\phi(f_a, P) = \begin{cases} 0 & \text{if } f_a \notin P \\ 1 & \text{if } f_a \in P \end{cases},$$

where $f_a \in$ P checks if P has an attribute with f_a as an attribute key.

In addition to providing a measure of the attribute leakage, we can also use this to calculate the average attribute leakage over a group of users by interpreting it as a count for the number of users revealing this attribute. For example, if the average attribute leakage is 0.5 for a group of users, we can say that half the group reveals the attribute.

Measuring attribute leakage for a particular attribute can reveal useful information, but to quantify a particular attack's threat, we measure the amount of attribute leakage for the set of attributes required to execute the attack. To do this, we normalize the total aggregate amount of attribute leakage for the set of attributes required for an attack.

The advantage of this method is that, as the number of attributes (from a set of attributes defined for a particular attack) increases, the amount of aggregate attribute leakage also increases. Thus, the approach approximates how close a user is to leaking all the attributes required for an attack.

Aggregate attribute leakage is the aggregate amount of attribute leakage for a set of attributes belonging to a user's social footprint. To normalize it to between 0 and 1, we divide by the number of attributes.

Definition 2. $\Psi(F_a, P)$ — aggregate normalized attribute leakage. Given a set of attribute names F_a and a user's social footprint P, we define the aggregate normalized attribute

Table 1. Percentage of profiles for a subset of attributes leaked by social networks.

Social network	Name	Location	Sex	Relationship status	Hometown	Homepage	Birthdate
Del.icio.us	—	—	—	—	53	—	—
Digg	100	67	55	—	—	—	30
Flickr	73	58	82	59	51	74	—
Last.Fm	82	—	87	—	76	77	—
LinkedIn	100	88	—	—	—	—	—
LiveJournal	93	69	—	—	—	68	64
MySpace	94	98	100	72	40	—	100
Technorati	94	—	—	—	—	—	—
Twitter	100	93	—	—	—	89	—
YouTube	68	—	—	—	29	57	73

leakage as

$$\Psi(F_a, P) = \frac{\sum_{f_a \in F_a} \phi(f_a, P)}{|F_a|},$$

where |X| represents the size of the set X.

The aggregate normalized attribute leakage over a group of users provides the average number of attributes released by the set of users for a specific attack.

Experimental Evaluation and Results

We conducted experiments using more than 8,200 social footprints and used the attribute leakage measure to quantify the amount of information discoverable for a subset of attributes. We then used the aggregate attribute leakage measure to quantify the amount of aggregate attribute leakage with respect to the set of attributes required for attacks.

Experiment Setup

Identity management sites let users manage their online identities by providing links to their various social networks (such as ClaimID, FindMeOn, and MyOpenID). On one such site, we crawled and collected 8,268 stored profiles. By parsing the links and subsequently following them to each user's social network profiles, we were able to construct the users' social footprints.

To perform a detailed study on attribute leakages, we wrote parsers for 10 of the 15 most popular social networks (we couldn't parse a few of them due to technical or legal challenges). The parsers performed considerable postprocessing to merge attributes with semantically similar meanings but syntactic differences. An example of this is "age" and "birthdate," in which we represent age as a coarse granularity birthdate. We also standardized certain attributes' values across social networks to semantically equivalent values. For example, if a Flickr user chose "Taken" for their relationship status, we standardized this value to "In a Relationship."

Table 1 shows the amount of information leaked from each site for a small subset of available attributes. For each attribute, we represent the amount of information as a percentage of the number of profiles that publicly displayed this value on a particular social networking site. A "—" indicates that the site didn't reveal the attribute.

Attribute Leakage

We expect the overall attribute leakage of a particular attribute to be related to the number of social networks that reveal that attribute. We thus analyzed the attribute leakage for a subset of attributes. To do this, we calculated the attribute leakage of each attribute in a user's social footprint and then averaged the results across the entire dataset. Figure 2 shows the result of these calculations for a subset of attributes. The y-axis shows the attribute leakage amount, and the x-axis shows the different attributes.

The figure shows that the leakage for attributes such as `Name`, `Location`, and `Homepage` is higher than that for `Relationship status` and `Birthdate`. As Table 1 shows, the attributes with low leakage are the ones revealed in only a few social networks. This confirms our hypothesis that the overall attribute leakage is related to the number of social networks that reveal the attribute.

Although this initial overview is useful, it combines users belonging to different numbers of social networks into one group. To get a deeper understanding of how attribute leakage varies as the number of social networks an attacker finds increases, we also investigate the attribute leakage as it relates to the number of social networks.

To include cases where an attacker might discover only a subset of social networks that a user with x social networks belongs to, we include all users where $y \leq x$ when looking at the attribute leakage for y social networks. Specifically, we pick every combination of y social networks from the user's online identity and calculate the average of the attribute leakage metric across combinations. This results in $\binom{x}{y}$ combinations. We pick this method because it avoids introducing biases that any predetermined ordering might impose. Henceforth, we use this technique in all figures that plot leakage against the number of social networks belonging to a user's social footprint.

Figure 3 shows the attribute leakage, varying based on the number of social networks observed, for the `Name` and `Hometown` attributes. The y-axis represents the average attribute leakage for users, while the x-axis shows the number of social networks. We investigated the `Name` attribute because it's the most often revealed attribute in social networks (see Table 1). Conversely, we studied `Hometown` because few social networks reveal it.

For the `Name` attribute, the ϕ attribute leakage for users with one social network is 0.73, and it quickly rises to 1.0 for users with four social networks. This result is expected because, as Table 1 shows, all but one social network reveals the `Name` attribute. For `Hometown`, the ϕ measure for users with one social network is 0.17; it consistently increases to its maximum of 0.71 with nine social networks. This result is also expected: five social networks reveal the `Hometown` attribute, so as the number of social networks increases, so does the likelihood that the value will be discoverable on a network in the user's social footprint.

Aggregate Attribute Leakage

The aggregate attribute leakage measure offers a combined view of the information leaked by a set of attributes required for a particular attack.

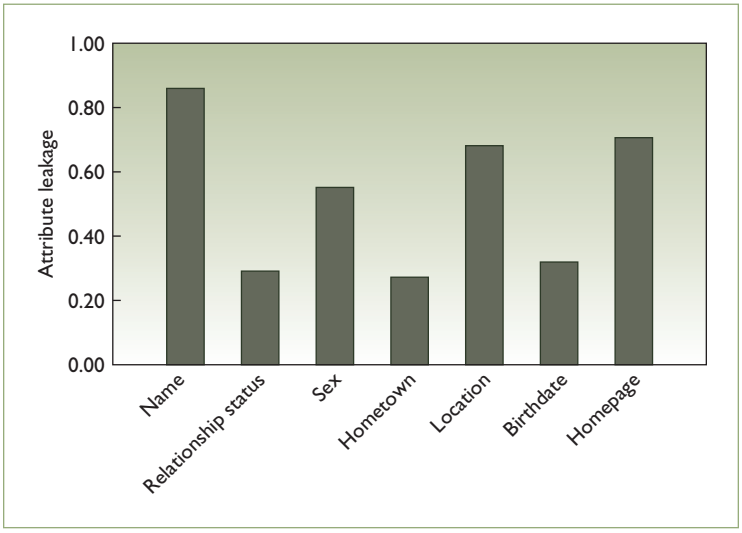

Figure 2. Attribute leakages (ϕ(f$_a$, P)) for a sample of attributes. Each attribute (f$_a$) is represented by a bar on the x-axis.

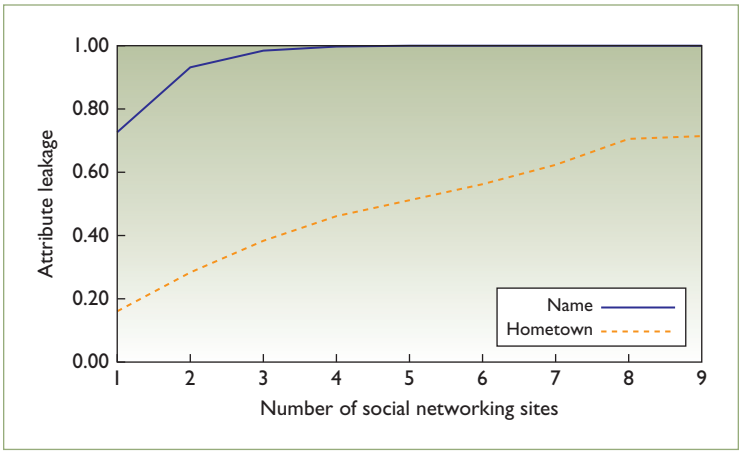

Figure 3. Attribute leakage measures for `Name` (ϕ("Name", P)) and `Hometown` (ϕ("Hometown", P)).

We begin by looking at the overall aggregate attribute leakage measure to quantify the users' average exposure to particular attacks. For the identification and password-recovery attacks, the aggregate normalized attribute leakage is 0.52 and 0.63, respectively. For both attacks, the aggregate normalized attribute value is above 0.5, which means that on average, users reveal more than half the attributes required for those attacks.

Once again, although this initial view provides a good overview of the different attack severities, we wanted to investigate how the aggregate attribute leakage changes as the number of social networks an attacker discovers increases. Doing so would also let us test our hypothesis that the aggregate attribute leakage

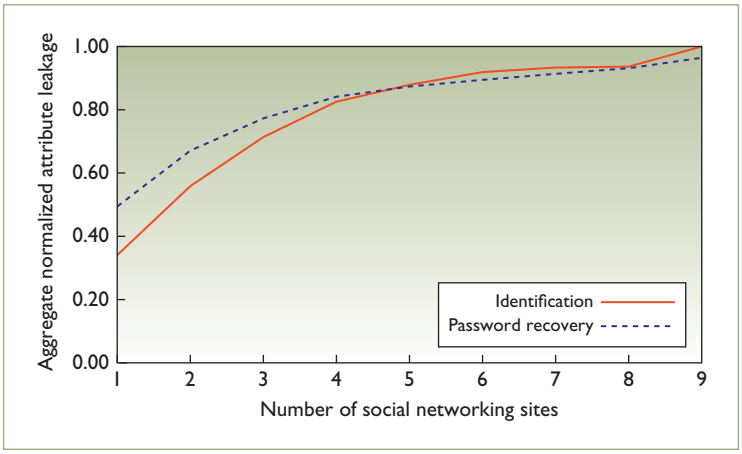

Figure 4. Aggregate normalized attribute leakage ($\Psi(F_a, P)$) for the identification and password-recovery attacks.

increases as the number of social networks increases.

Figure 4 shows the average aggregate attribute leakage (y-axis) as the number of social networks an attacker discovers (x-axis) increases. From this figure, we observe that the identification aggregate normalized attribute leakage is 0.34 (an average of one in three attributes) with one social network, which increases to more than 0.9 (an average of 2.7 in three attributes) with six or more social networks. Because the aggregate attribute leakage checks for an attribute's presence, this result implies that, with one social network, a person reveals one attribute of data required (on average) for the identification aggregate attribute leakage attack. However, with more than six social networks, a person reveals an average of 2.7 attributes. The password-recovery aggregate normalized attribute leakage follows a similar trend.

We used the aggregate normalized attribute leakage measure to quantify the threat of the physical identification and password-recovery attacks. Doing so revealed the value of the aggregate attribute leakage measure in several ways, while confirming our original hypotheses.

Mitigating Information Leakage Risks

Given the nature of their services, solutions to the unintended personal-information leakage problem must extend beyond social networks. Two primary methods mitigate the risk of unintended personal information being used in attacks.

Using Third-Party Services to Increase User Awareness

Increasing user awareness of information leakage and the associated risks should be an active component of the overall solution. We can increase this awareness in part by offering services that let users check their online footprint and assess how much information they're leaking. Users can thus increase their awareness of the risks associated with social networks.

The simplest way to accomplish this goal would be to create an online website that requests a list of the user's profiles and measures the information leaked in relation to particular attacks. Once they see the amount of information leaked and how it could facilitate attacks, the third-party service could advise users on ways to reduce information leakage by altering privacy settings or removing particular information from social networks. To prevent attackers from using this service, it might have to be offered as a tool by banking and credit bureaus or as a pay-per-use service.

Measuring Recovery Services' Susceptibility

Websites that rely on personal user information to provide certain services, such as password recovery, can also mitigate the risk of information leakage. Compared to the effort of increasing awareness in a majority of users, it's relatively easier to provide a service for website owners to evaluate the strength of account-recovery questions. Using a large dataset would make it possible to calculate how much information was leaked in regard to a particular question and use that information to estimate a question's strength.

As social networks move toward becoming a ubiquitous method for people to communicate, the danger of privacy leakage from online social networks grows more severe. As our work indicates, this danger increases substantially with multiple social network profiles. It's critical for users to understand such danger and take defensive measures to prevent personal information leak and defend against information misuse. We're currently investigating approaches that extend traditional mechanisms, such as k-anonymity and p-sensitivities, to model privacy protection in the context of multiple social profiles.

Acknowledgments

The US National Science Foundation (NSF) partially funded this work through its Industry-University Cooperative Research Center, CyberTrust, CISE/Computing Research Infrastructure, and NetSE programs. Other support includes the US National Institutes of Health grant U54 RR 024380-01 and its Clinical and Translational Science Award program's PHS grant (UL1 RR025008, KL2 RR025009, or TL1 RR025010), as well as gifts, grants, or contracts from Wipro Technologies, Fujitsu Labs, Amazon's Web Services in Education program, and the Georgia Tech Foundation through its John P. Imlay, Jr. Chair endowment. All opinions, findings, conclusions, and recommendations are those of the authors and don't necessarily reflect the views of the NSF or other funding agencies and companies.

References

1. A. Lenhart, *Adults and Social Network Websites*, Pew Research Center, 2009; www.pewinternet.org/Reports/2009/Adults-and-Social-Network-Websites.aspx.
2. C. Johnson III, *Protection of Sensitive Agency Information*, US Office of Management and Budget, 2006; www.whitehouse.gov/sites/default/files/omb/assets/omb/memoranda/fy2006/m06-16.pdf.
3. L. Sweeney, *Uniqueness of Simple Demographics in the US Population*, tech. report LIDAP-WP4, Lab. for Int'l Data Privacy, Carnegie Mellon Univ., 2000.
4. K. Zetter, "Palin E-mail Hacker Says It Was Easy," *Wired*, 18 Sept. 2008; www.wired.com/threatlevel/2008/09/palin-e-mail-ha.
5. M. Just, "Designing and Evaluating Challenge-Question Systems," *IEEE Security & Privacy*, vol. 2, no. 5, 2004, pp. 32–39.
6. A. Rabkin, "Personal Knowledge Questions for Fallback Authentication: Security Questions in the Era of Facebook," *Proc. 4th Symp. Usable Privacy and Security*, ACM Press, 2008, pp. 13–23.
7. S. Schechter, A. Brush, and S. Egelman, "It's No Secret: Measuring the Security and Reliability of Authentication via Secret Questions," *Proc. IEEE Symp. Security and Privacy*, IEEE CS Press, 2009; doi:10.1109/SP.2009.11.
8. A. Narayanan and V. Shmatikov, "Robust De-anonymization of Large Sparse Datasets," *Proc. IEEE Symp. Security and Privacy*, IEEE CS Press, 2008; doi:10.1109/SP.2008.33.
9. A. Narayanan and V. Shmatikov, "De-anonymizing Social Networks," *Proc. IEEE Symp. Security and Privacy*, IEEE CS Press, 2009; doi:10.1109/SP.2009.22.
10. V. Verykios, G. Moustakides, and M. Elfeky, "A Bayesian Decision Model for Cost Optimal Record Matching," *Int'l J. Very Large Data Bases*, vol. 12, no. 1, 2003, pp. 28–40.
11. M. Balduzzi et al., "Abusing Social Networks for Automated User Profiling," *Proc. 13th Int'l Symp. Recent Advances in Intrusion Detection (RAID)*, Springer, 2010, pp. 422–441.
12. D. Irani et al., "Large Online Social Footprints – An Emerging Threat," *Proc. Int'l Conf. Computational Science and Eng.*, vol. 3, IEEE Press, 2009, pp. 271–276.

Danesh Irani is a PhD candidate in the computer science program at Georgia Institute of Technology, where his thesis focuses on protecting data content flows in online systems, including preventing propagation of low-quality information and preventing leakage of good information. Irani has a BS in computer science from the University of Toronto. Contact him at danesh@cc.gatech.edu.

Steve Webb is a visiting professor of computer science at the University of Canterbury in Christchurch, New Zealand. His research interests include security threats that target information systems such as P2P networks, email systems, the World Wide Web, and social networking environments. Webb has a PhD in computer science from the Georgia Institute of Technology. Contact him at webb@cc.gatech.edu.

Calton Pu is a professor and the John P. Imlay, Jr. Chair in Software at the School of Computer Science, Georgia Institute of Technology, where he's working on projects in automated system management in cloud environments and document-quality analysis. His research interests include operating systems, transaction processing, systems reliability and security, and Internet data management. Pu has a PhD in computer science and engineering from the University of Washington. He's a senior member of IEEE, a member of the ACM, and a fellow of the American Academy of Arts and Sciences. Contact him at calton@cc.gatech.edu.

Kang Li is an associate professor of computer science at the University of Georgia. His research interests include computer networks, multimedia networking, and operating systems. Li has a PhD in computer science and engineering from the Oregon Graduate Institute of Science and Technology. He's a member of IEEE and the ACM. Contact him at kangli@cs.uga.edu.