Your Privacy, My Privacy? On Leakage Risk Assessment in Online Social Networks

Ruggero G. Pensa and Livio Bioglio

Dept. of Computer Science, University of Turin, Turin, Italy
{ruggero.pensa,livio.bioglio}@unito.it

Abstract. The problem of user privacy enforcement in online social networks (OSN) cannot be ignored and, in recent years, Facebook and other providers have improved considerably their privacy protection tools. However, in OSN's the most powerful data protection "weapons" are the users themselves. The behavior of an individual acting in an OSN highly depends on her level of privacy attitude, but, in this paper, we show that user privacy is also influenced by contextual properties (e.g., user's neighborhood attitude, the general behavior of user's subnetwork) and define a context-aware privacy score to measure the objective user privacy risk according to the network properties.

Keywords: privacy metrics · online social networks · information spread

1 Introduction

The problem of user privacy in the so-called "Big Data Era" cannot be ignored and many companies are realizing the necessity to consider it at every stage of their business. In practice, they have been turning to the principle of *Privacy by Design* [2] by integrating privacy requirements into their business model. Online social network (OSN) providers are embracing this model as well. However, in OSN's the most powerful data protection "weapons" are the users themselves. The behavior of an individual in these situations highly depends on her level of privacy awareness: an aware user tends not to share her private information, or the private information of her friends, while an unaware user could not recognize some information as private, and could share it without care to her contacts, even to untrusted ones, putting her privacy or the privacy of her friends at risk. Users' privacy awareness then turns into the so-called "privacy attitude", i.e., the users' willingness to disclose their own personal data to other users, that can be measured by leveraging the way users customize their privacy settings in social networking platforms [5].

A new question may arise now: how safe is the privacy of a social network user who is mostly surrounded by friends with a good privacy attitude? The question is not trivial, since the way most people set their privacy settings is based on the notion of closeness: close friends are usually allowed to see all user's updates, while acquaintances can only see "public" or less sensitive updates. The common assumption is that closed friends are trusted ones and thus will not disclose friends' posts to other friends. In this paper, we model the effects of privacy attitude on information propagation in social networks with the aim of studying what happens to information diffused to friends with different levels of privacy awareness. The outcomes of this study lead to the intuition that privacy risk in a social network may be modeled similarly as page authority in a hyperlink graph of web pages. In fact, it is a well-known fact that more authoritative websites are likely to receive more links from other authoritative websites. Our hypothesis is that we may transpose the concept of "importance" of a web-page into the concept of "privacy risk" of users in a social network as follows: the more an individual is surrounded by friends that are careless about their privacy, the less the individual her/himself is likely to be protected from privacy leakage.

With the final goal of enhancing users' privacy awareness in online social networks, in this paper we propose a new context-aware privacy score based on personalized Pagerank [4], one of the most popular algorithms to rank web pages based on a personalized view of their importance (or authority). We show the effectiveness of our privacy measure on a large network of real Facebook users.

2 Information diffusion vs. privacy

We consider a social graph G composed by a set of n nodes $\{v_1, \dots, v_n\}$ representing the users of G. We represent the social network as a directed graph $G(V, E)$, where V is a set of n nodes and E is a set of directed edges $E = \{(v_i, v_j)\}$. Given a pair of nodes $v_i, v_j \in V$, $(v_i, v_j) \in E$ iff there exists a link from v_i to v_j (e.g., users v_i is in the friend list/circle of v_j or v_j follows v_i). We define the neighborhood of a node $v_i \in V$ as the set of nodes v_k directly connected to the node v_i, i.e., $\mathcal{N}(v_i) = \{v_k \in V \mid (v_i, v_k) \in E\}$. Finally, we consider a set P of privacy classes, representing the propensity of a user of the class to disclose her own or other's information. Each user of G belongs to a privacy class $p \in P$.

We employ an extension, proposed by us in [1], of the SIR model for considering the explicit or implicit privacy polices of an individual during the spread of information on a social network. A privacy class in the set $P = \{p_0, p_1, \dots, p_N\}$ is assigned to the susceptible (S) and infectious (I) compartments, representing the privacy attitude of an individual belonging to the compartment, and consequently her behavior on information spreading, from less aware (p_0) to more aware (p_N). This behavior is reached by assigning different values to the parameters λ (infection probability) and μ (recovery probability) of each privacy class: every privacy class $p \in P$ is linked to a different pair of values λ_p and μ_p. We also introduce a novel parameter $\beta_p \in [0, 1]$ in the SIR transmission model, symbolizing the interest in information of the users in privacy class p. We denote with $c(v_i, t) \in \{S, I, R\}$ the compartment of user v_i at time t. The evolution probabilities are obtained as follows. Let $p(v_i) = p \in P$ be the privacy class of an individual v_i. If it belongs to the susceptible compartment, it may be infected at time $t+1$ with probability: $P_{inf}(v_i, t+1) = \beta_p \cdot \left(1 - \prod_{p' \in P}(1 - \lambda_{p'})^{n_I(v_j, t)}\right)$, where $n_I(v_j, t) = |\{v_j \in \mathcal{N}(v_i) \mid c(v_j, t) = I \ \wedge \ p(v_j) = p'\}|$ is the number of individuals in infectious (I) compartment and privacy class p' at time t among the

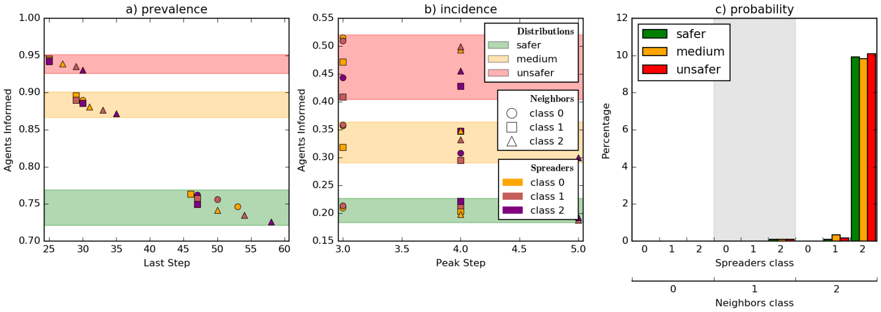

Fig. 1. a) Median of prevalence of informed individuals (ratio) in each set of simulations. b) Median of incidence of informed individuals (ratio) in each set of simulations. c) Probability that information does not reach more than 1% of the population.

neighbors of individual v_i. Otherwise, if the individual v_i of privacy class p belongs to the infectious (I) compartment at time t, it may recover with probability μ_p at time $t + 1$.

We now provide the results of our experiments performed over a Facebook-like synthetic network generated using LDBC–SNB Data Generator which produces graphs that mimic the characteristics of real Facebook networks [3]: in particular, we generate a network with 80,000 nodes, but here we consider only the greatest connected component of such network, composed by approximately $75,000$ nodes and $2,700,000$ edges. See [1] for the details of our experiments.

The features extracted from our simulations are graphically summarized in Figure 1(a) for prevalence curves and in Figure 1(b) for incidence curves. For each point of these graphs, the marker color identifies the privacy class of the initial spreader, its shape identifies the privacy class of the neighbors of initial spreader node, while its background color identifies the distribution of privacy classes of the nodes in the whole network. Each point shows the median value resulting from 100 simulations performed on 10 initial spreaders.

The most noticeable result is the role of the attitude towards privacy of the initial spreader and its neighbors: a safer attitude of the node and its neighbors decreases the portion of informed population, and extends the duration of information diffusion, but its impact is not as influential as the behavior of the whole network. For an aware user, even if the probability of diffusing information to a friend is low, the number of friends is so high that a small number of friends can become spreaders themselves. As soon as information reaches a node out of the neighborhood, its diffusion depends only on the attitude on privacy of the whole network. For this reason we decide to analyze the portion of simulations where information has reached only a small portion of the population, lower than 1%, before dying: our results are reported in Figure 1(c). In this case we notice that the attitude of the network is irrelevant, and only the privacy classes of the spreader and its neighborhood is crucial. Interestingly, a safe attitude of the spreader or of the neighbors is not sufficient on its own to block information diffusion. An information item on a safer user with unsafer friends, and vice versa, can easily overtake the circle of friends.

These observations lead to the following crucial consideration: to assess the objective privacy risk of users, one cannot simply take into account their propensity to self-protection. The attitude of their neighbors also counts, as well as the attitude of the whole subnetwork in which their interact. In the next section, we will present a measure that quantifies the privacy leakage of users considering the risks due not only to their attitude towards privacy but also to the attitude of their friends and subnetwork.

3 A context-aware privacy metrics

We consider the social graph $G(V, E)$ defined in Section 2 and associate, to each user $v_i \in V$, an *intrinsic privacy risk* $\rho_p(u_i)$, which is defined as the user propensity to privacy leakage. The assumption is that some users are more prone to disclose their personal data than others and, thus, they are intrinsically more at risk. In the following, we instantiate $\rho_p(u_i)$ according to the *P-Score* proposed by Liu and Terzi [5], an established and reliable definition of privacy score.

As shown in Section 2, if a user is mostly surrounded by friends (or friends of friends) that do not care that much about privacy, then she should be more exposed to privacy leakage than a user who is principally connected to friends that care about their own (and others') privacy. This consideration leads to the intuition that privacy risk in a social network may be modeled similarly as page authority in a hyperlink graph of web pages. Hence, we transpose the concept of "importance" of a web-page into the concept of "privacy risk" of users in a social network as follows: the more an individual is surrounded by friends that are careless about their privacy, the less the individual her/himself is likely to be protected from privacy leakage. Hence, we correct the *P-Score* by using *personalized Pagerank* [4], one of the most popular algorithms to rank web pages based on their centrality (or authority). It is used to create a personalized view of the relative importance of the nodes. We can now introduce our context-aware privacy score (called *CAP-Score*), defined by the following distribution:

$$\boldsymbol{P}^{\rho} = d\boldsymbol{A}^{\top}\boldsymbol{P}^{\rho} + \frac{(1-d)}{\sum_{k=1}^{n} \rho_p(u_k)}\boldsymbol{\rho} \tag{1}$$

where $\boldsymbol{\rho} = [\rho_p(u_1), \ldots, \rho_p(u_n)]^{\top}$, $\boldsymbol{P}^{\rho} = [p^{\rho}(v_1), \ldots, p^{\rho}(v_n)]^{\top}$ is the *CAP-Score* vector ($p^{\rho}(v_i)$ being the *CAP-Score* associated to vertex v_i), $d = [0, 1]$ is the damping factor (the $1 - d$ quantity is also known as restart probability) and $\boldsymbol{A}$ is a $n \times n$ matrix such that each element $a_{ij} = 1/deg^{+}(v_i)$ ($deg^{+}(v_i)$ being the outdegree of v_i) if $(v_i, v_j) \in E$ ($a_{ij} = 0$ otherwise).

An example of context-aware score computation is given in Figure 2. In Figure 2(a), we provide an example of graph where an aware user (the central one) is surrounded by unaware users (i.e., users with high intrinsic risk). Figure 2(b) represents the same network with the computed *CAP-Scores*: the score value of

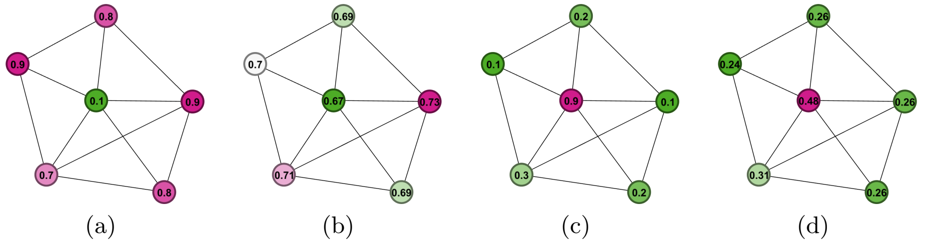

Fig. 2. Context-aware scores (b and d) in two differently aware networks (a and c).

the central user is adjusted according to the context and it is higher than in Figure 2(a). Instead, in Figure 2(c), we provide a network with the same topology but different intrinsic risks. In particular the unaware central user (with high risk) is surrounded by rather aware users (with low privacy risk). In this case, our measure for the central user is revised upwards (see Figure 2(d)), according to a context in which all other users form a kind of barrier protecting the privacy of the central users.

4 Experimental results

In this section, we show experimentally the improved effectiveness of our *CAP-Score* w.r.t. information diffusion phenomena. In our experiments we use a snapshot of the Facebook graph consisting on the ego-networks of real Facebook users gathered leveraging an online experiment described in [6]. It is a graph with 75,193 nodes and 1,377,672 edges, with the largest connected component consisting of 73,050 nodes and 1,333,276 edges. The degree distribution of the network is given in Figure 3(a). For 101 users, who replied to a specific survey (specified in [6] as well), we have computed the *P-Score* [5].

We study the relationship between the two definitions of privacy score (*P-Score* and *CAP-Score*) and the effects of information propagation across the network. A good privacy score should take into account the number of nodes that may potentially access and diffuse some information coming from other nodes in the same network. For this reason, we perform several Monte Carlo simulations of an information propagation scenario within our snapshot of Facebook. In particular, we adopted the Susceptible-Infectious-Recovered (SIR) epidemic model. In our experiments, for all nodes we set an infection probability $\lambda = 0.5$ and a recovery probability $\mu = 0.3$. Then, we select N seed nodes that, in turn, are considered as the individuals that start the infection (i.e., information diffusion process) and measure the number of nodes (*prevalence rate*) that are either infected (I) or recovered (R) after each step of the simulation. The seed nodes are the 101 Facebook users for which we have measured the *P-Score*. Finally, for each simulation step we compute the Spearman's ρ coefficient between the prevalence rate and the two privacy scores. The results are reported in Figure 3. Interestingly, the gap between the *P-Score*'s ρ and our context-aware score's ρ

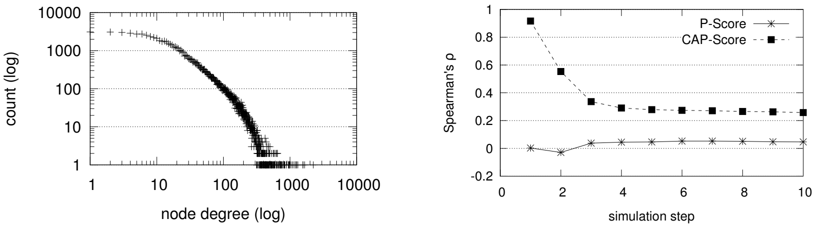

(a) (b)

Fig. 3. Degree distribution (a) and correlation between prevalence and privacy (b).

in the very first iterations is significantly large. Assuming that the *P-Score* correctly measures the privacy risk based on users' privacy preferences, a possible explanation is that the users underestimate their centrality within the network. Although our Facebook snapshot cannot be considered a statistically valid sample of the entire Facebook graph, the huge difference in terms of correlation with the prevalence rate confirms that privacy leakage metrics should not ignore the context in which the users operate within the social network.

5 Conclusions

In this paper, we have shown how privacy attitude can affect the diffusion of information on social networks and, with the final goal of supporting users' privacy awareness in online social networks, we have proposed a context-aware definition of privacy score. This measure, as shown in our experiments, is a good estimate of the objective privacy risk of the users. Based on these results, we believe that our framework can be easily plugged into any domain-specific or general-purpose social networking platforms, thus inspiring the design of privacy-preserving social networking components for *Privacy by Design* compliant software [2].

References

1. Bioglio, L., Pensa, R.G.: Impact of neighbors on the privacy of individuals in online social networks. In: Proceedings of ICCS 2017. pp. 28–37 (2017)
2. Cavoukian, A.: Privacy by design [leading edge]. IEEE Technol. Soc. Mag. 31(4), 18–19 (2012)
3. Erling, O., Averbuch, A., Larriba-Pey, J., Chafi, H., Gubichev, A., Prat-Pérez, A., Pham, M., Boncz, P.A.: The LDBC social network benchmark: Interactive workload. In: Proceedings of ACM SIGMOD 2015. pp. 619–630. ACM (2015)
4. Jeh, G., Widom, J.: Scaling personalized web search. In: Proceedings of WWW 2003. pp. 271–279. ACM (2003)
5. Liu, K., Terzi, E.: A framework for computing the privacy scores of users in online social networks. TKDD 5(1), 6:1–6:30 (2010)
6. Pensa, R.G., Di Blasi, G.: A privacy self-assessment framework for online social networks. Expert Systems with Applications 86, 18–31 (2017)