Privacy Scores: Assessing Privacy Risks Beyond Social Networks

Michal Sramka

***Abstract*—Assessing privacy risks arising from publishing private information on social networks is challenging for the users. Privacy Scores were proposed in the past to provide each user with a score – a measurement of how much sensitive information a user made available for others on a social network website. We present the privacy scores, discuss their shortcomings, and show several research directions for their extensions. We propose an extension that takes the privacy score metric from a single social network closed system to include background knowledge. Our examples and experimental results show the need to include publicly available background knowledge in the computation of privacy scores in order to get scores that more truthfully reflect the privacy risks of the users. We add background knowledge about users by means of combining several social networks together or by using simple web search for detecting publicly known information about the evaluated users.**

***Index Terms*—Privacy, Risk analysis, Inference algorithms.**

I. INTRODUCTION

RECENTLY there was an explosion of popularity of web sites that allow users to share information. These sites – social-network sites, blogs, and forums such as Facebook, Twitter, LinkedIn, and others – attract millions of users. The users publish and share information about themselves by creating online profiles, posting blogs and comments. Such information often contains personal details. Quantifying the individuals' privacy risk due to these information-sharing activities of the individuals is a challenging task.

Securing individuals' privacy in these environments and protecting users against threats such as *identity theft* and *digital stalking* becomes an increasingly important issue. Both users and service providers recognize the need for users' privacy. The sites may provide some privacy controls. However, the users are faced with too many options and too many controls, and lack the understanding of privacy risks and threats or are unable to accurately asses them. This all contributes to the confusion for the users, and often results in skipping the complicated and time-consuming tasks of setting the privacy controls that should protect them.

Even with properly configured privacy settings for a user profile, some privacy concerns remain. Take for example discussion forums, where tenths or hundreds contributions to multiple discussions of various topics are written by a user. Although the user is careful not to disclose any personally identifiable information in his/her individual posts, personal, sensitive, and private information may be disclosed by looking at the set of all posts by the user. From the cumulative set of all posts, it may be then possible to profile the user, infer the user's opinions or even identity.

Michal Sramka is with the Institute of Computer Science and Mathematics, Faculty of Electrical Engineering and Information Technology, Slovak University of Technology, Ilkovicova 3, 812 17 Bratislava, Slovakia, e-mail: sramka@stuba.sk

Privacy Scores by Liu and Terzi [1] were proposed to quantify the privacy risks of individuals posed by their profiles in a social-network site. Focus here is on privacy risks from the individuals' perspective. In the proposed framework, each user in a social network is assigned a privacy score based on the information in his/her profile compared to all available information in all profiles. The score then measures the user's potential privacy risk due to having his/her profile available on the social-network site.

The main drawbacks of this proposal of privacy scores are the concentration only on users' profiles and inconsideration of other publicly available information about the users on the same social network and beyond it. In particular, *background knowledge* about a user is not included in the computation of the privacy score. Background knowledge (sometimes referred to as external knowledge or auxiliary information) is some information about an individual that by itself is not a privacy disclosure, but combined with other information it becomes one.

A. Our Contribution

We propose a new concept for privacy scores. We explore the idea of presenting users with a new privacy score that measures their overall potential privacy risk due to available public information about them. Compared with the original privacy scores by Liu and Terzi, we overcome the drawbacks of concentrating only on users' profiles in a single social network, and we include publicly available background knowledge in computation of the new privacy scores. Our new privacy scores metric better represents the potential privacy risks of users and thus helps them make better decision in managing their privacy.

Our results are twofold. Firstly, in Section II-A we discuss the shortcomings of the privacy scores. We present several opportunities for extending the original privacy scores. With the extension of including background knowledge in mind, we identify some background knowledge that is publicly available but that cannot be easily extracted by computers in an automatic manner. Secondly, we proposed an extension of the privacy score metric that takes it from a closed system evaluating privacy over a single social network to a metric that includes information about the users that comes from outside the social network. In Section III, we present examples and experimental results showing paradoxes that may happen when the computation is over only a single social network. Next, in Section IV, we extend the computation of privacy scores to

include two or multiple social networks. Our final proposal, in Section IV-B, uses web searches to include all available public indexed human knowledge in the computation of the privacy score of a user. Thus, our new privacy score reflects the privacy risks of combining user's profile information with available knowledge about the user represented by the web.

Our proposed method for making web search inferences while scoring the privacy risks of individuals can also be seen as a privacy attack. However, we do not explore this direction, as there are already too many attacks, some of them referenced later in Section I-B. Our contribution rather focuses on helping users achieve their privacy needs and lower their privacy risks. The extended privacy score helps the users to make more informed decisions about their online activities.

B. Related Work

Our work is influenced by the approach by Liu and Terzi [1], which provides users with a quantification of privacy risks due to sharing their profiles in a social network. Each user is assigned a privacy score based on their and all other users' profile items. The proposal is for a single social-network site, that is, a closed system evaluation of privacy that lacks the consideration and inclusion of background knowledge in computation of the privacy scores. We overcome this shortcoming by including background knowledge in the computation of privacy scores, see Section III and IV.

PrivAware [2] is a Facebook application that scores privacy settings of a user based on the user's profile and profiles of his/her direct friends, which are implicitly available to any Facebook application. The score represents individuals' privacy risks arising from using third-party Facebook applications. In addition to [1] and [2], there exist several other scoring systems that somehow evaluate and rank users in social networks, but not their privacy. For some of them, please refer to [1]. However, none of these systems measures privacy risks for individuals.

The privacy risks of social-network sites are summarized in [3]. Several papers present privacy attacks in social networks [4], [2], [5], [6], [7], [8] or try to lower privacy risks and prevent privacy attacks in social networks [9], [10]. In addition, there are privacy risks from being tracked while browsing these websites [11].

Some form of background knowledge is usually considered in privacy attacks and is very likely available to attackers. Absolute privacy is impossible, because there will be always some background knowledge [12]. Inference techniques can then be used to attack or to help protect private data. In particular, web-based inference detection [13], [14] has been used to redact documents and prevent privacy leaks.

II. PRIVACY SCORES

Privacy Scores by Liu and Terzi [1] were proposed to quantify the privacy risks of individuals posed by their profiles in a social-network site. The privacy score is a combination of each one of user's profile items, labelled $1, \ldots, n$, for example, real name, email, hometown, land line number, cell phone number, relationship status, IM screen name, etc. The contribution of each profile item to privacy score is based on sensitivity and visibility. The *sensitivity* β_i depends on the item i itself – the more sensitive the item is, the higher is the privacy risk of it being revealed. The *visibility* of an item i belonging to a user j is denoted $V(i,j)$ and captures how far this item is known in the network – the wider the spread in the network, the higher the visibility.

The privacy score of an item i belonging to a user j is simply $\text{PR}(i,j) = \beta_i \times V(i,j)$. The overall *privacy score* for a user j with n items is then computed as

$$\text{PR}(j) = \sum_{i=1}^{n} \text{PR}(i,j) = \sum_{i=1}^{n} \beta_i \times V(i,j) \ . \qquad (1)$$

To keep the privacy score PR a non-decreasing function, in order for it to be a nicely behaving score, both the sensitivity β_i and visibility $V(i,j)$ must be non-negative functions. In practice, the sensitivity and visibility are determined from an $n \times m$ matrix R that represents n items for m users of a single social network. The value of each cell $R(i,j)$ describes the willingness of the user j to disclose the item i. In the simplest case, the value of $R(i,j)$ is 0 if the user j made the item i private and 1 if the item i is made publicly available. From this, the (observed) visibility can be defined as $V(i,j) = R(i,j)$. In a more granular approach, the matrix R can be defined by $R(i,j) = k$, representing that the user j disclosed the item i to all the users that are at most k jumps away in the graph of the social network. Regarding the sensitivity of an item, β_i can be computed using Item Response Theory (IRT) [1]. The IRT can be also used to compute the true visibility of an item for a user.

The privacy score is computed for each user individually. It is an indicator of the user's potential privacy risk – the higher the score of a user, the higher the threat to his/her privacy.

A. Shortcomings and opportunities of privacy scores

The privacy score is no doubt a useful metric for each and every user of a social network. Nevertheless, there exist several shortcomings of the originally proposed privacy scores. We list a few of them here. Some of these were already noticed and identified by the authors of the privacy scores, others are just observations, and some are our proposals for further exploration, research, and enhancements of privacy scores.

Regarding the items of a user profile, one can immediately notice hardship in quantifying the items themselves:

- The granularity of profile items is of particular concern. For example, the profile item "personal hobbies" can cover a range of non-private and private information and so its true sensitivity cannot be really established for the general case required by the privacy scores.
- Different profile items have different life-cycles. Some profile items may have a time attribute attached to them – for example, a cell/mobile phone number or an address are temporary information, while the date of birth or the mother's maiden name are permanent for life. The proposed privacy score, as defined, ignores these facts. We believe that implicit time relevance should be taken

into account for more precise evaluation of a user's privacy.

- Impossibility or hardness of including all, possibly private, information in privacy score computation. For example, consider photos: It may be hard or impossible in some cases to (automatically or even by a human involvement/assessment) establish relationships from photos. Or whether a person is drinking alcohol in a photo. Another example are discussion forums: Information is exhibited in natural language form. Determining a political orientation of a user from a single post may not be possible, yet looking at the cumulative set of the user's posts, private information can be inferred about the user (see Section III).

Of more concern and interest is the computation and use of sensitivity β_i for item i. As proposed, sensitivity is computed from the matrix R, that is, sensitivity is based only on the users and items in the single social network. When considering a single social network represented by a matrix R, it is easy to get a wrong perception of privacy due to the limited information about the users.

- The sensitivity β_i computed for an item i would reflect the true real-world sensitivity of this item only if the distribution of people in the social network would mirror the real-world distribution. Obviously, many social networks are not like this, and so a paradoxes are likely because this fact. For example, take a date of birth that most people consider a sensitive and private information. However, if everybody in a social network reveals his/her date of birth, then this item will be considered as not sensitive at all (because everybody reveals it). Paradoxes on the other side of the spectrum are possible, too. For example, if an item in a social network is filled only by one or a few users, because the other users are too lazy to fill it in, then the item will be considered sensitive (by the computation of sensitivity), although the item is far from being considered sensitive or private in the real life.
- No background knowledge inclusion, and so no inference detection or control: A privacy metric should include "background knowledge" (auxiliary information or external knowledge) in establishing a score for a user. Speaking more generally, a single social network or any closed system evaluation is not sufficient for real and proper privacy evaluation of a user.
  For privacy scores, this means that the computation of the score should not depend only on the matrix R coming from a single social network. Several extensions of the original privacy score metric are possible based on the background knowledge type and source. In Section IV we propose a new method to compute privacy scores, one that considers information about users beyond the ones in the social network, namely from a second/other social networks or more generally from the web.

Finally, it needs to be mentioned that the proposed privacy score metric measures only some aspect of privacy, namely attribute (item) disclosure and identity disclosure arising from the attribute disclosure. There are several other aspects that may be of concern to the users of a social network, such as:

- the risks of identity disclosure that is not based on attribute disclosure – for example, based on behavioral observations,
- the risk of identity theft,
- the risk of link or relationship disclosure,
- the risk of group membership disclosure, or
- the risk of digital stalking.

How to measure these risks and help the users making informed decisions by presenting them a score reflecting these risks is still an open problem.

III. Discussion forum

The computation of privacy scores proposed by Liu and Terzi [1] introduced in Section II assumes the analyzed information to be readily available for inclusion in the matrix R. As we noted in Section II-A, non-structured information cannot be always easily included for analysis. It may be either information that is hard to extract – for example, relationships from photos – or previously not defined information – for example, non-structured text in natural language may contain multiple private items some of which may not be pre-defined as items of the matrix R.

Together with my Master's student Ján Žbirka we performed a few experiments [15], where simple natural language analysis was used to determine if some private information has been included in discussion comments on a news website. Since the users usually post multiple comments, they may contain multiple private information that must be looked-up for inclusion in the privacy scores. In our experiments, shown in the next Section, we concentrated on information about political orientation before election and religious believes.

A. Experimental results

Discussions of the Slovak news web site www.sme.sk were analyzed just before the government election in March 2012. From all the users that posted comments on the website, 5,268 users who posted more than 500 comments over the lifetime of the website were considered as the most active ones. In the three weeks before the election, these 5,268 most-active users posted 43,035 comments that were analyzed. Almost 20% of the analyzed users revealed in their comments which political party in particular they were or were not going to vote for.

Summary of the findings are in Table I and all the other details about the experiment can be found in the Master's thesis of Ján Žbirka [15].

Since discussions on this website about religion and church are very heated, we also analyzed whether it is possible to find out the faith/religious beliefs of the users from their comments. The experiment that was done on the same sample of the users and comments have shown that simple natural language analysis can determine faith, although the users were more conservative in revealing their religious believes compared to the political orientation. In total, 133 (2.5%) users were found to disclose their religion, and 106 (2.0%) users were found to disclose that they are atheists.

TABLE I
THE NUMBER OF THE USERS (FROM THE TOTAL OF 5,268) WHO WERE FIND TO DISCLOSE THIS INFORMATION IN DISCUSSION COMMENTS

Users who will	vote		not vote	
at all	763	(14.5%)	173	(3.3%)
for a right wing party	209	(4.0%)	194	(3.7%)
for a left wing party	59	(1.1%)	46	(0.9%)
for a particular party	688	(13.1%)	335	(6.4%)

IV. PRIVACY SCORE EXTENSION

The biggest disadvantage of the privacy scores that were outlined in Section II is the non-consideration of background knowledge. *Background knowledge* (sometimes referred to as external knowledge or auxiliary information) is some information about an individual that by itself is not a privacy disclosure, but combined with other information it becomes one. We propose two possible extensions of the original privacy score metric that take public background knowledge into account.

It needs to be noted that the reason to include of background knowledge in the computation of the privacy score is two-fold. On the one hand, such extended privacy score will more precisely present users with privacy risks arising from publishing their information. On the other hand, using background knowledge also reduces another shortcoming of the original privacy scores. Namely, the more background knowledge is considered, the closer is the sensitivity of items to the true sensitivity. In other words, adding background knowledge to the privacy score computation also reduces or eliminates sensitivity paradoxes – see Section II-A.

Our extended privacy score metric uses the same formula as in the equation (1) with sensitivity and visibility as the original privacy scores, but the information that is used to compute these – the matrix R – is extended by additional knowledge. We discuss two instances of this extension. The first one, presented next, combines information from two or several social networks when evaluating privacy risk of a user. The second instantiation of the extended privacy score metric, which we present in Section IV-B, uses "all the human knowledge" in privacy risk evaluation.

Our proposal of a simple inclusion of additional information in the privacy score computation is based on users' information (items) from multiple social networks. Let N be the number of considered social networks, and let R_t be as the already defined matrix R for a social network t, with $t = 1, \ldots, N$. Hence, R_t is a $n \times m$ matrix, where $R_t(i, j)$ represents the publicity of an item i for a user j – that is, non-disclosure when $R_t(i, j) = 0$ or disclosure when $R_t(i, j) > 0$ and possibly how far from the user j is the item public in the (graph of the) social network t.

It is likely in practice that not all the users are in every social network and that every item is in each of the corresponding profiles. Here we assume that the range of the items $i = 1, \ldots, n$ and the range of the users $j = 1, \ldots, m$ are the supersets over all the social networks, and so $R_t(i, j) = 0$ if an item i or user j do not exist in the social network t. We define the matrix R used for sensitivity and visibility computation as $R(i, j) = \max_t R_t(i, j)$ and use the formula from the equation (1) to compute the privacy score. Such privacy score better estimates the risk of privacy disclosure.

Together with my Master's student Lucia Maringová we performed a few experiments [16], where the same users on two social networks were evaluated for their privacy risks. The two social networks were of different type, so it was expected that the users would behave differently and therefore would disclose different amount of information about themselves on each social network. In our experiments, shown in the next Section, we focused on computing privacy scores from each social network individually and then comparing the behavior of people in the terms of private information disclosure on two social networks.

A. Experimental results

The purpose of the experiment is to show that privacy risks, as measured by the original and extended privacy score, are higher when two social networks are combined. Specifically, this means that some users tend to be conservative in one social network while publicly disclose private information in another social network.

For the experiments, profiles from the same users on two social networks were downloaded and analyzed. The social networks (websites) were Pokec.sk and Zoznamka.sk. They both belong to the same content provider, and so use the same user authentication, which facilitated the pairing of the users from the two networks. Zoznamka.sk is a dating website, where a profile can contain up to 5 items: age, body type, weight, height, and contact. Pokec.sk is a website about chatting, messaging, and picture sharing. A profile on Pokec.sk can contain up to 34 items.

A sample of 3,923 users was selected. From all these users, there are only 23 users (<1%) who completely filled all profile items on both websites. These people can be considered very open minded and/or not understanding or ignoring the risks of disclosing private information. Roughly 32% of the users shared the same information on both sites.

Because of the nature of the website, users on the dating website Zoznamka.sk revealed more personal information about themselves. This is likely due to the fact that the users tried to create interest and attract the users who viewed their profiles. No user had less than 2 items (out of 5) filled on Zoznamka.sk. Conversely, many users on Pokec.sk left their profiles empty. What is of interest to us are the users who had empty profiles on Pokec.sk and non-empty profiles on Zoznamka.sk. Table II summarizes these users. All the details can be found in the Master's thesis of Lucia Maringová [16].

In the terms of the privacy score, the users on Pokec.sk who had empty profiles would receive the score of 0, because they do not share or disclose anything. However, this would be awfully wrong in any privacy risk analysis, because private information about these users is publicly available and linkable to these users. At least two additional items can be learned about roughly 44% of the users with empty profiles on Pokec.sk when considering Zoznamka.sk, so the extended privacy score computed over both networks for these users

TABLE II
THE NUMBER OF USERS WHO SHARED NOTHING ON POKEC.SK, BUT HAD NON-EMPTY PROFILES ON ZOZNAMKA.SK. NOTE THAT MINIMUM ITEMS FILLED IN ON ZOZNAMKA.SK WAS 2.

On Pokec.sk	On Zoznamka.sk	# of users
0 items	2 items	542
0 items	3 items	107
0 items	4 items	961
0 items	5 items	119
0 items	> 0 items	1,727
		44%

will be non-zero. This simple experiment itself shows the need to extend the original privacy scores from analyzing information over one social network to analyzing also auxiliary information.

B. Using all the human knowledge in privacy score computation

Extending the original Privacy Scores by Liu and Terzi [1] to multiple social networks certainly helps in privacy risk evaluation. The selection of social networks included in the extended privacy score computation, presented above, strongly impacts the quality and truthfulness of the score. The most truthful privacy risk evaluation can be achieved if all the human knowledge is used in the computation of the privacy score.

Including "all the human knowledge" in any computation is obviously impossible, so an approximation would have to suffice for all practical purposes. To effectively include the knowledge, we need to be able to quickly search for particular information or relation. Thus, we should use all the *indexed* human knowledge. Private databases and the "deep web" are believed to contain much more information than what is publicly available. In general, private information is out of reach for privacy adversaries as well as for privacy evaluators. Hence, we foresee to use all the *indexed public* human knowledge in the privacy score computation. Currently, the best instance and the best source of all the indexed public human knowledge is (Google) web search. In fact, there exists a proposal, namely Web-Based Inference Detection [13], [14], which takes advantage of the assumption that the web search is the proxy for all human knowledge.

Our idea is as follows: If an item of a user is not disclosed in the social network, we want to determine if the item has been disclosed elsewhere by using an inference detection based on the other disclosed items for the user. Our inference detection method is heavily influenced by the Web-Based Inference Detection [13]. So, our idea rewritten in the terms of inference detection is: If there is a privacy-impacting inference detected for an undisclosed item, then this detected inference should be included in the privacy score computation.

More formally, we propose the following method to compute the privacy score:

Consider a social network of m users each having a possibility to fill a profile of n items. Let R be, as before, the $n \times m$ matrix over $\{0, 1\}$ with $R(i, j)$ representing whether the user j has (or has not) disclosed the item i. Let P be $n \times m$ array of strings with $P(i, j)$ being the value of the item i for the user j, in case this value has been disclosed. Let the set D_i be the domain of the item i. Finally, let β, γ, and δ be positive integers, where β and γ are parameters of the proposed algorithm that control the search depth, and δ is the parameter that controls the number of the most frequent words to be considered. Then the algorithm to extend R and determine the users' disclosures outside the social network is as follows:

For each user j, $j \in \{1, \ldots, m\}$

1) Let $S_j = \{k \mid R(k, j) = 1,\ k = 1, \ldots, n\}$ be the set of all disclosed items for the user j.
2) For each undisclosed item i, that is, for all $i \in \{1, \ldots, n\}$ with $R(i, j) = 0$
   a) Let T be an empty multiset.
   b) Take every subset $S'_j \subseteq S_j$ of size $|S'_j| \leq \beta$.
   c) For every such subset $S'_j = \{i_1, \ldots, i_\ell\}$ with $\ell \leq \beta$
      i) Use a web search engine to search for keywords $P(i_1, j), \ldots, P(i_\ell, j)$
      ii) Retrieve the top γ most relevant documents containing these keywords
      iii) Extract the top δ most frequent words from all these γ documents
      iv) Add the top δ most frequent words to T together with their frequencies
   d) Take the most frequent word from T that is also in D_i, if it exists.
   e) If there is such word, let $R(i, j) = 1$.

After this, the newly enhanced matrix R contains the users' disclosures not just from the social network itself, but also from the web. Visibility and sensitivity values can be then computed from this matrix R, and the privacy score can be computed for each user using the equation (1).

The parameters β, γ, and δ can be tuned to achieve different trade-offs between the running time of the algorithm and the completeness and quality of the disclosure detection. In fact, these values can be different for different users, perhaps based on the number of items disclosed in the social network. Additional tuning can be achieved by performing the steps of the algorithm only for those users that have disclosed a "sufficient" number of items in the profile that would allow the web search to identify additional items.

V. CONCLUSIONS

As more and more users are joining and using social-network web sites, they become more heavily used and their owners look for new ways to share different content, including private information and information that may lead to unwanted privacy leakages. It becomes increasingly difficult for individuals to control and manage their privacy in the vast amount of information available and collected about them.

Privacy Scores were proposed as a metric that presents users with a score that reflects their privacy risks arising from disclosing information in their profiles on a social network. We presented several shortcomings of the privacy scores as research opportunities for extending the privacy score metric.

Next, we supported the need for extensions by experimental results from different websites and social networks. Finally, we proposed two extensions of the privacy score metric that consider additional background information about the users in the computation of the scores. Our approach provides a better decision support for individuals than the original privacy scores. Based on our extended privacy score metric, the users can compare their privacy risks with other fellow individuals and make informed decisions about whether they share too much potentially private and sensitive information.

Michal Sramka has received PhDs in Mathematics and Applied Mathematics from Florida Atlantic University and Slovak University of Technology, and also holds MSc in Computer Science. He has worked as a researcher at universities in the USA, Canada, Spain, and Slovakia. His main research interests are in information security, specifically in data privacy and cryptology. Michal Sramka is the holder of the Werner von Siemens Excellence Award and a lifelong member of the Pi-Mu-Epsilon honorary society.

REFERENCES

[1] K. Liu and E. Terzi, "A Framework for Computing the Privacy Scores of Users in Online Social Networks," in *Proceedings of the Ninth IEEE International Conference on Data Mining, ICDM 2009*, 2009, pp. 288–297.

[2] J. Becker and H. Chen, "Measuring Privacy Risk in Online Social Networks," in *Web 2.0 Security & Privacy 2009 Workshop of 2009 IEEE Symposium on Security and Privacy, W2SP 2009*, 2009.

[3] D. S. Rosenblum, "What Anyone Can Know: The Privacy Risks of Social Networking Sites," *IEEE Security & Privacy*, vol. 5, no. 3, pp. 40–49, 2007.

[4] A. Narayanan and V. Shmatikov, "De-anonymizing Social Networks," in *Proceedings of the 30th IEEE Symposium on Security and Privacy, S&P 2009*, 2009, pp. 173–187.

[5] B. Krishnamurthy and C. E. Wills, "On the leakage of personally identifiable information via online social networks," in *Proceedings of the 2nd ACM Workshop on Online Social Networks, WOSN 2009*, 2009, pp. 7–12.

[6] E. Zheleva and L. Getoor, "To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles," in *Proceedings of the 18th International Conference on World Wide Web, WWW 2009*, 2009, pp. 531–540.

[7] A. Korolova, R. Motwani, S. U. Nabar, and Y. Xu, "Link privacy in social networks," in *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM 2008*, 2008, pp. 289–298.

[8] L. Backstrom, C. Dwork, and J. M. Kleinberg, "Wherefore Art Thou R3579X?: Anonymized Social Networks, Hidden Patterns, and Structural Steganography," in *Proceedings of the 16th International Conference on World Wide Web, WWW 2007*, 2007, pp. 181–190.

[9] A. Felt and D. Evans, "Privacy Protection for Social Networking Platforms," in *Web 2.0 Security & Privacy 2008 Workshop of 2008 IEEE Symposium on Security and Privacy, W2SP 2008*, 2008.

[10] E. Zheleva and L. Getoor, "Preserving the Privacy of Sensitive Relationships in Graph Data," in *Proceedings of the First ACM SIGKDD International Workshop on Privacy, Security, and Trust in KDD, PinKDD 2007*, 2007, pp. 153–171.

[11] K. McKinley, "Cleaning Up After Cookies," iSEC Partners, Technical report, 2008. [Online]. Available: https://www.isecpartners.com/files/iSEC_Cleaning_Up_After_Cookies.pdf

[12] C. Dwork, "Differential Privacy," in *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming, ICALP 2006*, 2006, pp. 1–12.

[13] J. Staddon, P. Golle, and B. Zimny, "Web-Based Inference Detection," in *Proceedings of the 2007 USENIX Annual Technical Conference, USENIX 2007*, 2007, pp. 71–86.

[14] R. Chow, P. Golle, and J. Staddon, "Inference detection technology for Web 2.0," in *Web 2.0 Security & Privacy 2007 Workshop of 2007 IEEE Symposium on Security and Privacy, W2SP 2007*, 2007.

[15] J. Zbirka, "Privacy risks arising from publishing private information on the web (in Slovak)," Master's Thesis, Advisor: Michal Sramka, Slovak University of Technology, 2012.

[16] L. Maringova, "Privacy risks arising from publishing private information in social networks (in Slovak)," Master's Thesis, Advisor: Michal Sramka, Slovak University of Technology, 2012.