FriPEL: Friendship Privacy Educational Labware

Na Li
nali@pvamu.edu
Department of Computer Science
Prairie View A&M University

Lin Li
lilin@pvamu.edu
Department of Computer Science,
Prairie View A&M University

Yuhong Liu
yhliu@scu.edu
Department of Computer Engineering
Santa Clara University

***Abstract*—As social media grows increasingly in popularity, so do people's concerns about their privacy disclosure. Considering the large amount of time the younger generations spend on online social networks, educational activities are needed to promote privacy protection. This paper presents an initial effort on exploring a hands-on learning approach to educate students on social media privacy. Specifically, a labware was developed to educate students that their query behaviors in online social networks may disclose the private relationships of other users on the site. The labware aims at making students aware of the privacy issues on social media, understanding the costs of privacy protection, stimulating their interests, and improving students self-efficacy. This paper discusses the design and implementation of the labware, which was evaluated among a group of student volunteers through the lecture and hands-on activities. Student feedback was very positive and encouraging.**

***Keywords*—friendship privacy, social networks, friend search engine, sociability**

I. Introduction

With the spiral growth of social media platforms (e.g., Facebook and Twitter), people have spent tremendous time online, enjoying interacting with their families and friends. However, due to the serious data breaches that happened in recent years, for instance, Facebook-Cambridge Analytica data scandal [2], people became more and more concerned about their privacy disclosure. Many have realized that their online activities could generate a massive amount of private data including personal profile and unique online footprints. Therefore, they are worried about what data is collected, who will use their data, and how the data will be used.

A large group of researchers have been dedicated to developing effective technologies of protecting users' privacy on social media for about a decade. They have addressed different social media privacy issues, such as image privacy [15] [14], data sharing privacy [12], and online service/application privacy [11] [13]. However, there have been no well-established curricula for educating younger students (e.g., high school students and undergraduate students) on such a topic. The ACM's Computer Science Curriculum 2013 has acknowledged the importance of privacy education [8], but the teaching content is still lacking.

Recently, a team of cross-disciplinary members, including computer scientists, educators, and social scientists, at the International Computer Science Institute (ICSI) and UC Berkeley, developed an online privacy curriculum which targets younger students [9]. This team designed and implemented ten principles with the purpose of spreading the awareness of protecting privacy among younger students and helping them better understand what happens to personal information when it goes online, how it might be used to negatively affect users, and how they can defend their privacy by limiting what they share. They focused on online privacy in general.

In this paper, the authors focused on the privacy protection on social media, which is a very important aspect of privacy education, especially considering such a large population of younger users surfing on Online Social Networks (OSN) on a daily basis. The goal is to make students aware that their behaviors through OSN applications may compromise other users' privacy. Since hands-on labs proved to be effective in engaging students into the learning process, an interactive labware was also developed for teaching privacy attacks and defenses, particularly the vulnerability of friend search engines. A friend search engine is an application which is often integrated with social media platforms, such as Facebook, and used to retrieve friend lists of individual users. Through the lab activities, students can learn that a friend search engine may open a door to attacks which expose other users' private relationship information by making simple queries.

The outline of this paper is as follows: Section II discusses the design of the friend search engine, particularly focusing on the vulnerability, the tradeoff between preserving users' relationship privacy and the cost, as well as three display strategies corresponding to user queries. Section III details the design and implementation of our web application, followed by a case study. Section IV evaluates this labware by showing the feedback from a group of student volunteers who attended the pilot test session and spent time playing with the web application. A conclusion is made in Section V.

II. Friend Search Engine

This section first discusses the tradeoff in designing a friend search engine and then introduces three display strategies for responding queries through a friends search engine.

A. Design of a Friend Search Engine

In the design of a friend search engine, OSN operators are motivated to display the entire friends list in response to each query, as they believe that showing more information helps to increase the site sociability by stimulating users' interactions and creating more activities online. For instance, given the full list, the requestor may see some of his friends have already made friends with the queried user. Such observation will encourage him to make friends with the queried user as well. However, not all users feel comfortable to share their friend lists to the public, due to their privacy concerns. Therefore,

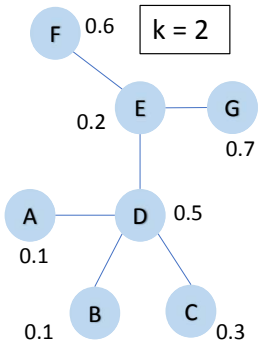

Fig. 1: A toy example for illustrating display strategies

there exists a tradeoff between increasing site sociability and protecting users' privacy. The tradeoff must be handled carefully; otherwise, people may feel discouraged to continue their use of social media platforms.

A straightforward solution is to partially display the list. Individual users may specify how many friends they would like to show in response to a query. A parameter k is determined by the user's privacy preferences. If a user is sensitive about sharing his friendship information, he can assign k with a small value; otherwise, a large value can be set. Currently, the privacy setting options provided on OSNs, such as Facebook, are all-at-once, either hiding or exposing all friends of a user. These are two extreme cases of the partial displaying strategy where k is set to zero or the total number of friends. A further question to ask is which k friends should be selected for display. Three display strategies will be introduced and compared in the following subsection and these strategies were implemented in the developed web application.

B. Comparison of Three Display Strategies

Random K – randomly selects k friends from the entire friends list. This strategy is not secure because if one queries the same user multiple times, he may get different sets of k friends of the queried user, thereby discovering more than k of his friends in total. This problem is named *Direct Exposure*.

Rank K – always picks the k friends who have the most impact on the site sociability for display. This strategy does not have the direct exposure problem since it shows the same set of k friends no matter how many times the user is queried. But it is still vulnerable. Suppose when user x is queried, his friendship with user y is not displayed. Their friendship is exposed, however, when user y is queried. In this case, user x's privacy is compromised as its $k+1$ friends are disclosed with these two queries. The reason for this privacy breach is that each friendship is associated with two end users and can be exposed by a query upon either of them. This problem is called *Mutual Effect*.

Top K – ensures that any user appearing in the query results displays only k friends (or the maximum number of his friends if the user has less than k friends in total on the OSN). This strategy was first proposed in the work [11]. The Top K strategy can well handle the tradeoff between preserving privacy and increasing the sociability of an OSN. The strategy consists of two primary steps: (1) sorting friends of the queried user in the descending order of their impact on the site sociability; and (2) checking each friend in the sorted list to determine whether it can be displayed as a friend of the queried user taking into account of the privacy restriction. This strategy can defend against the two problems addressed above, direct exposure and mutual effect.

Figure 1 shows a toy example of an Online Social Network, where k is set to 2. The nodes and edges are used to represent users and friendship between users. The impact of each user on site sociability is normalized to a value between 0 and 1. A larger value indicates a bigger impact. If one wants to query D first and then E. With Rank K, C and E are displayed after D is queried. Then F and G are displayed after E is queried. With Top K, C and E are displayed after D is queried, but D and G are displayed after E is queried. This is because the relationship between D and E is discovered already when D is queried. Although F's impact value is higher than D's, displaying his relationship with E will expose three of E's friends, namely D, G, and F, resulting in privacy breach. Therefore, F cannot be displayed.

III. Web Design and Case Study

The objective of this research is to develop an interactive web application with which students can conduct the following activities: (1) making queries through a friend search engine; (2) visualizing query results in a graph; (3) detecting the occurrence of privacy breach; (4) comparing the responses of different display strategies; and (5) analyzing the tradeoff between preserving users' privacy and enhancing the sociability of the OSNs.

The web application was developed on Eclipse using the frameworks of Spring [7] and Hibernate [5] and the Twitter4J [10] library, which is the unofficial Java library for accessing the Twitter APIs. The application interacts with a MySQL database and queries Twitter users for their followers through the Twitter API "GET followers/list" [3]. The MySQL database stores the queries made by requestors, the followers lists returned from the Twitter API, and the followers displayed on the web site as query results.

A. Log-in Page

The log-in page is intended for authenticating requestors and authorizing them to access the web application. Instead of creating accounts for individual requestors on the web site, users were directed to the Twitter site for authentication, where they use their own Twitter username and password to log in. Then Twitter generates a unique PIN which the requestors need to type into the corresponding field on the web page to sign in. Twitter uses OAuth [6] for authenticating users and authorizing the application to access Twitter API for queries on behalf of the requestors who are Twitter users. The screenshots of corresponding pages are displayed in the Figure 2.

Once a requestor logged into the web site, he will see the status of his queries from his previous visit. The status information includes the display strategy he selected previously and the results of the queries he made. Such query results are presented with different components on our query page which will be particularly introduced in the following subsection.

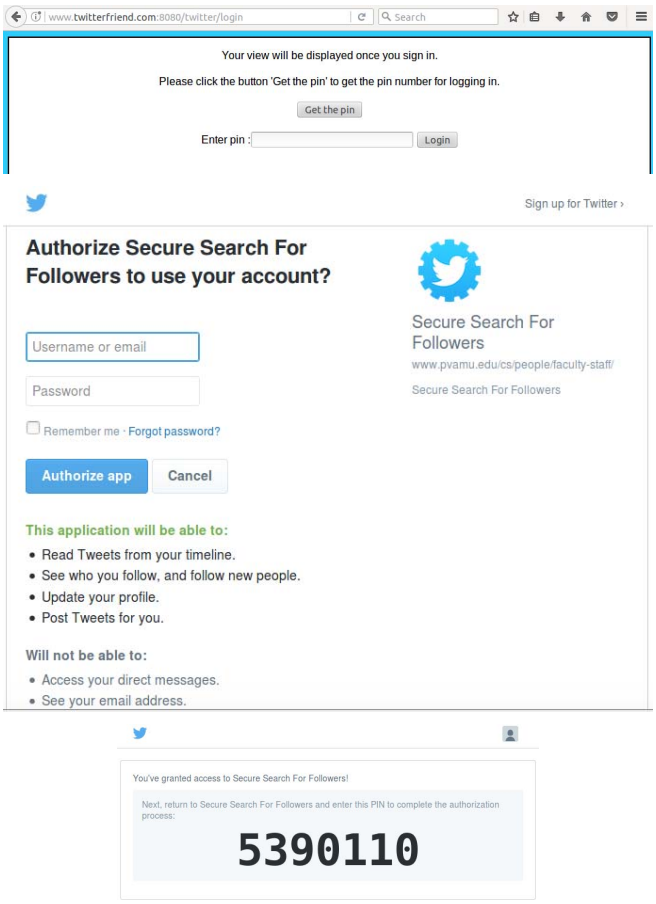

Fig. 2: User authentication through Twitter account

B. *Query Page*

As shown in the Figure 3, the query page can be divided into two sections: user input section and visualization section. In the user input section, a requestor can choose any of the three display strategies discussed in Section II, Random K, Rank K and Top K. The selected strategy is applied to the friend search engine to determine which k followers will be displayed to the requestor. The requestor can switch between those strategies anytime, but the switch will clear the requestor's query history and start new query. This design is to prevent different strategies from interfering with each other. Specifically, the query results from Random K should not be kept when Top K was selected; otherwise, nodes compromised with Random K will retain in the visualization, thereby causing confusion. The other component in this section is the search interface. The requestor can type the Twitter ID of a Twitter user and then start the query. Twitter ID instead of Twitter username was used for the purpose of anonymizing users to protect their privacy.

The visualization section contains three components: the Twitter "following" relationship graph, the statistical data chart for evaluating attack effectiveness, and the impact of a display strategy on site sociability according to the queries already made. Whenever a new query is issued, these components will be updated accordingly.

The Twitter "following" relationship graph: The application used Cytoscape Web [1] to visualize the query results in a graph pattern, where nodes and edges represent Twitter users and their "following" relationship. Additionally, some features on the relationship graph were programmed to facilitate the interactions of a requestor with the web application.

First, if a requestor attempts to query a user who is already in his query results, he can simply click on the user node and then the selected user's Twitter ID will be added to the search box automatically. This feature saves the requestor's effort on typing the user ID for query.

Second, the most recently queried node is colored in red. Since Cytoscape Web automatically adjusts the size of nodes and the length of edges, as well as the position of the graph, in order to accommodate all nodes and edges in the visualization window nicely, coloring the recently queried node will help the requestor track his query focus.

Third, the compromised nodes are colored in green, where compromised nodes are defined as those users who have more than k friends exposed in the query results. Coloring these nodes differently will distinguish them from other nodes exposed in the results, and make it easier for the requestor to see any privacy breaches.

The statistical data chart: The application used *google.visualization.LineChart* [4] to plot the number of users whose relationship privacy has been compromised by the current requestor. This statistical data tells the seriousness of privacy breach with the selected display strategy. With the Random K and Rank K strategies, one can see the update of the chart with more private "following" relationships discovered by queries, due to direct exposure and mutual effect as discussed in Section II. No increase would be observed with Top K since it can protect users' relationship privacy, and thus, no user's privacy can be violated.

The impact on site sociability: As discussed in Section II, there is a tradeoff between preserving users' privacy and facilitating the sociability of online social sites. In reality, each user's impact on the sociability of an OSN should be measured by multiple factors, including but not limited to his online activities and social connections. For instance, if a user participates in a great number of group discussions, writes hundreds of twittees per day and actively retweets, his impact on the site sociability is huge. This is because his activities on Twitter will drive data flow, for example, other users may read and retweet his tweets. However, due to the restriction on the usage of Twitter APIs, only a certain number of queries can be made in a time window through Twitter APIs, which hinders the data collection process. Therefore, real-time evaluation was not doable. Instead, random values from the range $[0, 1000)$, called node weights, were assigned to nodes as impact indicators. The total weight of nodes visible in the result graph are calculated to show the impact of the nodes on the site sociability. This totoal weight keeps increasing with more queries.

A flow chart of the web query page is shown in Figure 4, which illustrates how the web application interacts with the local database, Twitter API, and web visualization interface. Due to the query restriction set by Twitter API, users are only allowed to send a certain number of queries within a short time

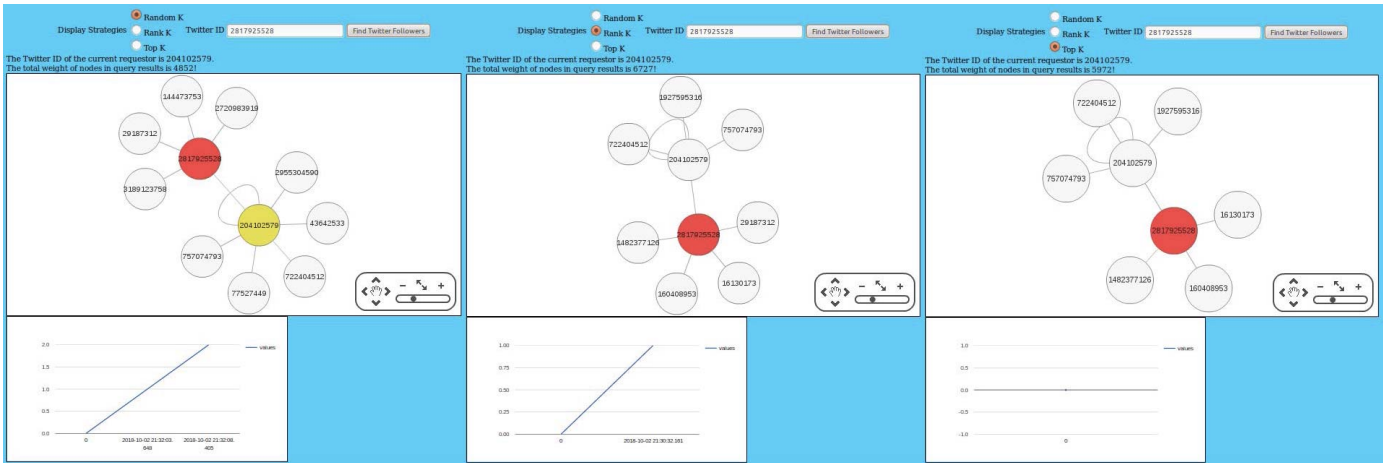

(a) Random K (b) Rank K (c) Top K

Fig. 3: Comparison of three display strategies

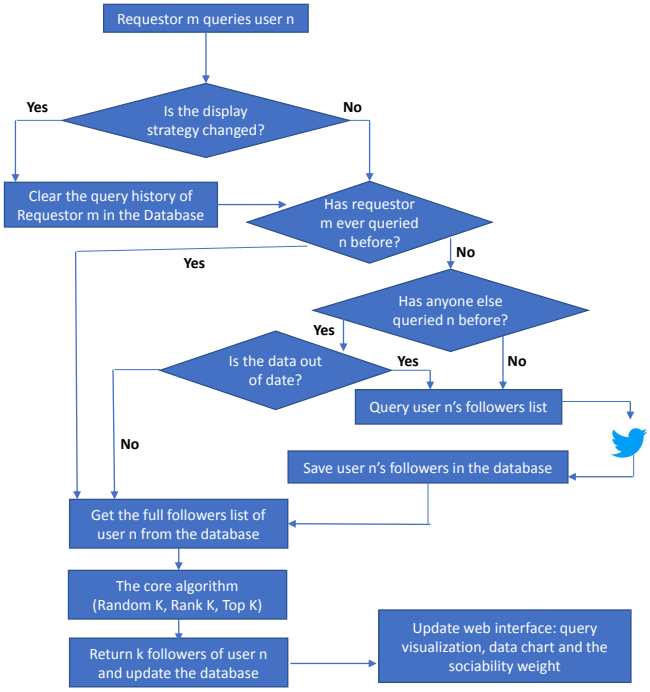

Fig. 4: The flow chart of our web application

window. Therefore, reducing the query number through the Twitter API is necessary. In order to do that, a local database is used to store the query results returned from Twitter to all requestors. In this way, the k followers displayed for one requestor may be derived from the full following list responded by Twitter for a query made by another requestor.

C. Case Study

The web application was thoroughly tested with different scenarios to assure its functionality and accuracy. A case study to demonstrate privacy disclosure and preservation using this application is as follows: first, a single Twitter account was created to play the role of the requestor. k was set to be 4 and the weights assigned to nodes were randomly selected from the range $[0, 1000)$. Then, the same set of nodes were queried in the following order: 204102579, 204102579, 2817925528, with three different display strategies. 204102579 was queried twice to compare the results. The screenshots of the web interface with different query results are depicted in Figure 3, which shows that two nodes were compromised with Random K; one node was compromised with Rank K; and no nodes were compromised with Top K. The total weight of nodes in the query results generated with Top K is closer to that weight generated with Rank K, which is much higher than the weight generated with Random K.

IV. Pilot Lab and Student Feedback

In order to expose students to privacy preservation in social networks and inspire their learning interests, the authors pilot tested this educational tool at the beginning of fall 2018. Teaching slides and lab instructions were also developed. A total of thirty students volunteered to participate in the evaluation. Among the students, eighteen were graduate students and twelve were undergraduate students. All participants were minority students.

The learning and evaluation activities fell into five categories: (1) lecture to introduce social media privacy attacks and protections, including data sharing privacy and application privacy; (2) class presentation to introduce the web tool; (3) lab environment setup; (4) hands-on labs using the tool to examine the three display strategies, and evaluate their performance in terms of privacy protection and site sociability; and (5) pre and post surveys to get students' feedback.

A. Pilot Lab

The students conducted all the defined activities which spanned about two hours. To enhance students' learning, participants were given time to collaborate and work out on all tasks including courseware functionalities, privacy protection

mechanisms, and data analysis, after lab discussions. Prior to the lab, a VirtualBox image of the lab environment were prepared and distributed to the students. The VirtualBox image has pre-configured Apache HTTP Server and MySQL database.

The lab activities the students took include:

- Activity I. Use students' own Twitter accounts to log into the web site.
- Activity II. Verify the security assurance of the three display strategies. Students may conduct queries randomly to see the change of the chart or they may intentionally take advantage of the vulnerabilities of direct exposure and mutual effect for verification. (Take screenshots of the query results and record the query sequences.)
- Activity III. Evaluate the three display strategies. Students need to conduct three groups of experiments. In each group of experiments, use three strategies individually to query a sequence of nodes, and write down the number of nodes compromised and the total weight. Try to query the same set of nodes and in the same order. Students may also take screenshots in each step and embed it to a Microsoft Word document. Plot the results into two bar charts, one for privacy and the other for weight. In the charts, x-axis is three groups of experiments, each group with three bars for three strategies, and y-axis is the number of nodes compromised or the total weight.

B. *Suvery Feedback*

Pre and post surveys were conducted at the beginning and the end of the testing session to evaluate the labware and analyze the outcomes. Table I presents the survey questions. All questions except for the last use a rating scale of 1 to 5 with 5 being the greatest deal or the most positive.

The students' feedback was positive and encouraging. For questions 1-6, the authors analyzed the average ratings with regard to the discrepancy of the pre and post surveys. The result showed a significant increase of students' awareness of and interest in privacy disclosure, privacy in social media, and anonymization after the lab, as depicted in Figure 5 and Figure 6. Figure 7 compares the rating change of the six questions between the graduate students and the undergraduate students. An interesting finding is that although both cohorts had similar increase in their ratings on the awareness of the three subjects, the undergraduate students showed much stronger increase in the ratings on their interests of the subjects. For questions 7-9, as depicted in Figure 8, the student average ratings are high. The percentage of participants who said that they gained a lot or a great deal in understanding of social media, anonymization mechanisms, and privacy disclosure and preservation is 70.0%, 53.3%, and 66.7%, respectively. For questions 10-12, over ninety percent of the participants said that they understood the possibility of compromising user's relationship privacy through the lab and knew the tradoff between privacy preservation and service usage. Almost all the participants felt that the lab should be taught in a security course. The percentages of students who rated either agree or strongly agree in the questions are 96.7%, 96.7%, and 90.0%,

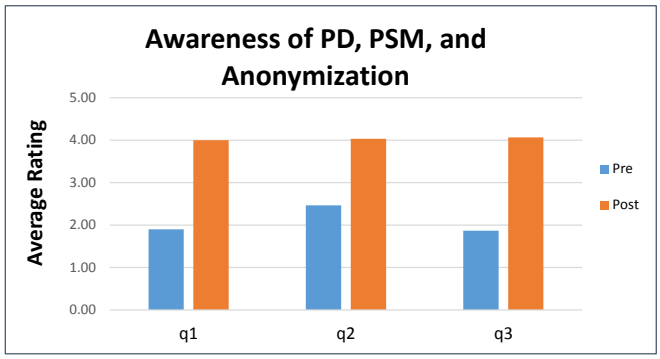

Fig. 5: Students' awareness of privacy disclosure, privacy in social media, and anonymization

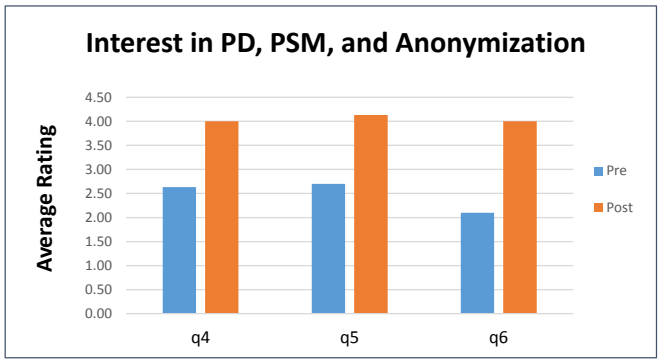

Fig. 6: Students interest in privacy disclosure, privacy in social media, and anonymization

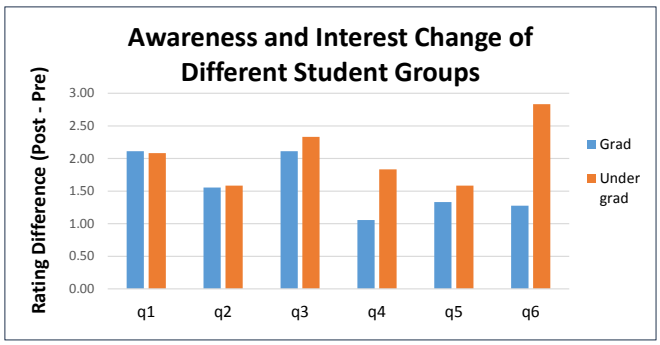

Fig. 7: Comparison of graduate and undergraduate students' awareness and interest change pre and post lab

respectively. Besides, students gave very positive comments with regard to question 13. Many said that the labware enhanced their learning in privacy preservation and social networks. Most participants felt that more exciting learning materials in privacy protection and educational tools like this should be developed in the future. Moreover, no significant difference was observed in the performance between males and females. Although the sampling space was small, this may reveal that, if appropriate instructional methods are adopted, more talented female students can be attracted into security fields and they can achieve the same as the opposite gender.

TABLE I: Pre and post survey questions

#	Survey Questions	Type
1	How do you rate your awareness about privacy disclosure and preservation?	Pre & Post
2	How do you rate your awareness about privacy in social media?	Pre & Post
3	How do you rate your awareness about anonymization mechanism?	Pre & Post
4	How do you rate your interest in privacy disclosure and preservation?	Pre & Post
5	How do you rate your interest in privacy in social media?	Pre & Post
6	How do you rate your interest in anonymization mechanisms?	Pre & Post
7	How do you rate your gains in understanding of social media?	Post
8	How do you rate your gains in understanding of anonymizaiton mechanism?	Post
9	How do you rate your gains in understanding of privacy disclosure and preservation?	Post
10	Understand the possibility to compromise users' relationship privacy when using apps in online social networks.	Post
11	Knowing there is a tradeoff between privacy preservation & using services in online social networks.	Post
12	I would like this lab and privacy preservation to be taught in a computer security course.	Post
13	How could students' learning about privacy preservation in social media be improved in this lab?	Post

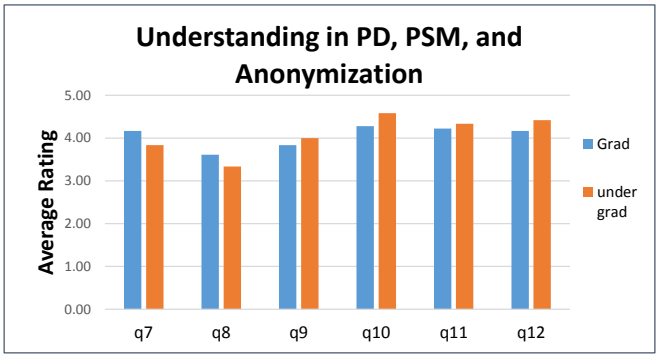

Fig. 8: Feedback of different student groups on the effectiveness of the learning tool

V. Conclusion

With the recent occurrence of serious privacy breaches on social media, people became more and more concerned about their privacy disclosure. The authors made an initial effort on developing hands-on learning materials to educate younger generations about privacy protection on social media. Particularly, this paper discussed the implementation of an interactive labware — a web application which aims at making younger students understand the vulnerability of friend search engines and the tradeoff of privacy protection and its cost. The labware was tested through an a pilot lab among a group of students. Feedback confirmed the effectiveness of the lab materials. Future work includes developing more hands-on labware to teach different aspects of privacy protection on social media, such as data sharing.

Acknowledgment

This project is supported in part by the National Science Foundation under grant no. 1712496. Any opinions, findings, and conclusions expressed in this paper are those of the authors, and do not necessarily reflect the views of the National Science Foundation. The authors would like to thank Adil MD Nowshad Hussain for helping with the experiments.

References

[1] Cytoscape Web. http://cytoscapeweb.cytoscape.org/.

[2] Facebook-Cambridge Analytica. https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html.

[3] GET followers/list. https://developer.twitter.com/en/docs/accounts-and-users/follow-search-get-users/api-reference/get-followers-list.html.

[4] Google Line Charts. https://developers.google.com/chart/interactive/docs/gallery/linechart.

[5] Hibernate Framework. http://hibernate.org.

[6] Oauth. https://dev.twitter.com/oauth.

[7] Spring Framework. http://projects.spring.io/spring-framework.

[8] T. J. T. F. on Computing Curricula (Association for Computing Machinery and I.-C. Society). Computer science curricula 2013.

[9] Teaching Privacy. http://teachingprivacy.org.

[10] Twitter4j. http://twitter4j.org/en/index.html.

[11] N. Li. Privacy-aware display strategy in friend search. In *Proceedings of IEEE International Conference on Communications (ICC), Communication and Information Systems Security Symposium*, pages 951–956, 2014.

[12] N. Li, N. Zhang, and S. Das. Preserving relation privacy in online social network data. *IEEE Internet Computing*, 15(3):35–42, 2011.

[13] Y. Liu and N. Li. Retrieving hidden friends: A collusion privacy attack against online friend search engine. *IEEE Transactions on Information Forensics and Security*, 2018.

[14] J. Yu, Z. Kuang, B. Zhang, W. Zhang, D. Lin, and J. Fan. Leveraging content sensitiveness and user trustworthiness to recommend fine-grained privacy settings for social image sharing. *IEEE Transactions on Information Forensics and Security*, 13(5):1317–1332, 2018.

[15] J. Yu, B. Zhang, Z. Kuang, D. Lin, and J. Fan. iprivacy: Image privacy protection by identifying sensitive objects via deep multi-task learning. *IEEE Transactions on Information Forensics and Security*, 12(5):1005–1016, 2017.