A study of the personalization of spam content using Facebook public information

Enaitz Ezpeleta[1], Urko Zurutuza[1], and José María Gómez Hidalgo[2]

[1] Electronics and Computing Department, Mondragon University
Goiru Kalea, 2, 20500 Arrasate-Mondragón, Spain
{eezpeleta,uzurutuza}@mondragon.edu,
[2] Pragsis Technologies
Manuel Tovar, 43-53, Fuencarral - 28034 Madrid, Spain
jmgomez@pragsis.com

Abstract. Millions of users per day are affected by unsolicited email campaigns. Spam filters are capable of detecting and avoiding an increasing number of messages, but researchers have quantified a response rate of a 0.006% [1], still significant to turn a considerable profit sending millions of emails, as the spammers do. While research directions are addressing topics such as better spam filters, or spam detection inside online social networks, in this paper we demonstrate that a classic spam model using online social network information can harvest a 7.62% of click-through rate. We collect email addresses from the Internet, complete email owner information using their public social network profile data, and analyze response of personalized spam sent to users according to their profile using a fake website. Finally we demonstrate the effectiveness of these profile-based emails to circumvent spam detection and we compare results between typical spam and personalized spam.

Keywords: spam, security, Facebook, personalized spam, online social networks

1 Introduction

The mass mailing of unsolicited e-mails has been a real threat for years. Spam campaigns have been used both for the sale of products as well as online fraud. Researchers are investigating many approaches that try to minimize this type of malicious activity that reports billions of benefits.

Within the spam problem, most research and products focus on improving spam classification and filtering. According to Kaspersky Lab data, the average of spam in email traffic for the year 2015 stood at 54.2%[3].

With the rise of online social networks (OSNs), specifically Facebook, which has more than 1.59 billion monthly active users as of December 2015[4], the extraction of personal information that users leave public on their profiles multiply

[3] https://cdn.securelist.com/files/2015/11/Q3-2015_Spam-report_final_EN.pdf
[4] http://newsroom.fb.com/company-info/

spam success possibilities. Facebook provides a great opportunity for attackers to personalize the spam, so a much lower volume of messages would get a higher return on investment.

We made a preliminary experiment regarding personalized spam campaigns in [2]. In this work we presented a more detailed view and analyzed precisely the results obtained during the execution of the campaigns.

The main objective of this paper is to measure the consequences of displaying information publicly in OSNs. It also aims to demonstrate that advanced techniques for generating personalized email that evade current spam detection systems while increasing the click-through rate can be developed. These techniques can enable new forms of attacks. First we extracted email addresses while crawling the Internet. These addresses were then checked on Facebook to look for related profiles. Once obtained a considerable quantity of user addresses, we extracted all the related public profile information and temporally stored it in a database. Then this information was analyzed in order to design user profiles based on their main activities in Facebook. Email templates were generated using common information patterns. Finally, to demonstrate the effectiveness of these templates when systems circumvent spam detection, different experiments have been performed. We collected sufficient evidence to confirm that the goal was successfully achieved.

The remainder of this paper is organized as follows. Section 2 describes the previous work conducted in the areas of personalized spam, and social network spam. Section 3 describes the process of the aforementioned experiments, regarding data collection, data processing, and personalized spam testing. In Section 4, the obtained results are described, comparing typical spam results with the personalized ones. Section 5 gives a discussion of the countermeasures that can be applied to prevent personalized spam. Section 6 describes the ethical considerations about this research. Finally, we summarize our findings and give conclusions in Section 7.

2 Related Work

2.1 Personalized spam

During the last years several works about the possibilities to create personalized spam or collect personal information from different OSNs have been proposed. For instance, in [3] authors launch targeted and non-targeted attacks on different channels using information from Facebook accounts.

In [4], researchers at University of Cambridge and Microsoft analyzed the difficulty of extracting user information from Facebook to create user profiles. They described different ways of collecting user related data, and they demonstrated the efficiency of the proposed methods. Authors conclude that the protection of Facebook against information crawlers was low. They also proved that big scale collection of data is possible. While it is true that Facebook has improved its systems' security since then, like limiting its own query language, the research proved that data extraction was effective.

In [5] researchers found a Facebook vulnerability giving attackers the possibility of searching people through email addresses in OSNs. Starting from a list of different emails, they managed to connect those email addresses with the account of their owners. After that, they collected all the information they could, and created different user profiles. This work gave a baseline for allowing attackers to launch sophisticated and specific attacks, but still did not realize about the potential of creating personalized spam campaigns. In the same direction, Polakis et al. demonstrated in [6] the real possibilities to create personalized spam campaigns in different OSNs.

2.2 Online social network spam

Over the last few years, social networking sites have become one of the main ways to keep track and communicate with people. Sites such as Facebook and Twitter are continuously among the top 10 most-viewed Web sites on the Internet[5]. The tremendous increase in popularity of OSNs allows them to collect a huge amount of personal information about users as authors prove in [7]. Unfortunately, this wealth of information, as well as the ease with which one can reach many users, also attracted the interest of malicious parties.

Researchers from the University of California proved that spam is a very big issue for OSNs [8]. In their research they created a large and diverse set of false profiles on three large social networking sites (Facebook, Twitter and MySpace), and logged the kind of contacts and messages they received. They then analyzed the collected data and identified anomalous behaviors of users who contacted their profiles. Based on the analysis of this behavior, they developed techniques to detect spammers inside OSNs, and they aggregated their messages in larger spam campaigns. Results show that it is possible to automatically identify accounts used by spammers, and block these spam profiles.

In Gao et al. [9] authors carried out a study to quantify and characterize spam campaigns launched from accounts on OSNs. They studied a large anonymized dataset of asynchronous "wall" messages between Facebook users. They analyzed all wall messages received by roughly 3.5 million Facebook users, and used a set of automated techniques to detect and characterize coordinated spam campaigns. This study was the first to quantify the extent of malicious content and compromised accounts in a large OSN. While they cannot determine how effective these posts are at soliciting user visits and spreading malware, their result clearly showed that OSNs are now a major delivery platform targeted for spam and malware. In addition, their work demonstrates that automated detection techniques and heuristics can be successfully used to detect social spam.

While most of the research focus on spam campaigns that might appear inside OSNs[10], we still think that a combination of typical spam and OSN spam exposes serious threats that needs to be addressed.

[5] http://www.alexa.com/topsites

3 Creating a personalized spam campaign

As shown in figure 1, our study followed four different phases. First, we collected a large amount of public information from Facebook. To do this we used email addresses that were publicly available when crawling the Internet. Then we computed a number of interesting statistics from the collected information that will be shown later. As a result of the data analysis, different user profiles were identified, and used them as customizable email templates. Once we had defined these templates, we developed an automatic email sending system and conducted two different experiments. Finally we analyzed the results obtained in the experiments.

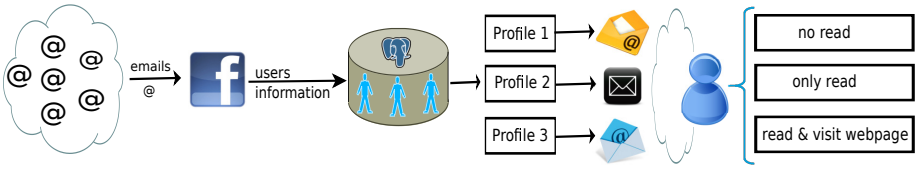

Fig. 1. Full process of personalized spam campaign creation

3.1 Collection of data

This process has been performed in three steps:

Email address harvesting. In this task we considered two options: the first one, obtaining the email addresses using the techniques as explained in [6], where they get e-mail addresses using various combinations of public information from OSN users. The second, using publicly available applications that automatically harvest email addresses from simple search queries over known search engines. The one used by the authors[6], generates a query for the search engine using a given keyword, and extracts email address patterns from the search result. We used a set of common keyword patterns such as "facebook", "hotmail", "gmail", "yahoo", "msn", and used both Google, Ask and Yahoo as search engines. Those patterns harvest email addresses from popular email service sites, which are at the end commonly used as user related data for online social networks.

Email address validation. Facebook offers the option to find the profile that is associated to a given email address. we developed an application that first authenticates a user to the OSN, and then searchers for a user corresponding to each email address harvested before.

[6] http://www.fast-email-extractor.com/

Once the user profile was found, we extracted and saved the user's ID and their full name.

Next URL is used to check if a given email (changing the word "EMAIL" with the email addresses) corresponds to a specific Facebook user.

```
http://www.facebook.com/search.php?init=s%3Aemail&q=EMAIL&type=users
```

Collection of the information. Facebook allows extracting information from the source code of all its web site pages. In order to do it, it is mandatory to access directly to the page from which we want to extract the information. Therefore, in this program we have used a user identifier from Facebook to connect directly to the user information page. Thus, we have visited all users pages and we have been able to extract all the public information that users have in the Facebook database.

Below is the address where all public information of each user can be found. 'USERUID' correspond to the ID that Facebook gives each user, which is stored in the database.

```
http://www.facebook.com/profile.php?id=USERUID&v=info
```

Results. We found that a 19% of the collected email addresses have a corresponding Facebook account associated to it. We found 22,654 Facebook accounts using 119,012 email addresses (19.04%).

3.2 Data Processing

At this stage the aim has been to treat the data stored in the database to extract user profiles. We have summarized the most useful information about each user, which is related to: id, sport, team, sportman, music, book, movie, tv, game, activity, interest, studies, languages, religion, politic, man, woman, partner and company. Most of the information stored is numeric, namely, the number of sports played by a person or number of activities that users entered in their own profile. Additionally each user also has logical or *boolean* data types, which indicates for example if the user in question is a male or not.

Using those user profile-based features, we gathered interests and user-related attributes to generate a set of statistics that could describe the behaviour of OSN profiles.

To obtain the more representative variables, we created a tag-cloud where we use every variable introduced by user, were appearing frequency increments the variable name.

Figure 2 shows that the most commons variables are Music, Gender and Studies. Table 1 gives a detailed view of those.

Fig. 2. Tag-cloud

Table 1. Number of users who have given each feature (total users: 22,654)

Feature	Amount	Percentage
man	8,786	39%
woman	6,189	27%
music	5,788	26%
titles	5,612	25%
company	5,149	23%

Results. With the extracted statistics, we can draw the following conclusions:

- As a result of descriptive analysis of the collected data, we found that 25.21% of the users do not insert any type of personal information, and 82,25% of the users have entered 3 or less types of variables.
- 66% of the collected users leave their gender public.
- Taking into account only the users that have at least one public variable, those are the percentage of the most common variables: gender 88%, music 34%, studies 33% and company 30%.
- The variables related to personal information that are most added by the users (gender, music, studies and company), are still in very low ranges as to be processed and used for clustering user profiles, as shown in table 1.

3.3 Personalized spam

The objective of this phase is to create different email templates that will be later used. With this templates it is possible to send personalized mail to all Facebook users stored in the database. Once the templates were designed and implemented, a strategy to count the number of users that "bite the hook" of the spam was designed. For this we have implemented a website.

Mail templates. Before any other action, the first step was to define a template through which we were able to send personalized emails to the people.

As shown in table 1, the most abundant variables are those related to the gender. Although these data cannot be used for creating templates, it can be used for implementing a formal greeting according to their gender. Following the process defined in figure 3, we have used the other three most common variables to create spam templates. That is, if a user has entered his favourite music group in the profile, it will receive a personalized **music template**. However, if the person has not added any singer or group in Facebook and has added the university in which he or she has studied, she will receive a personalized message with the **studies template**. And if none of those two had been added but the information refers to their current job or a company in which the user works, she will receive a personalized the email using the **company template**.

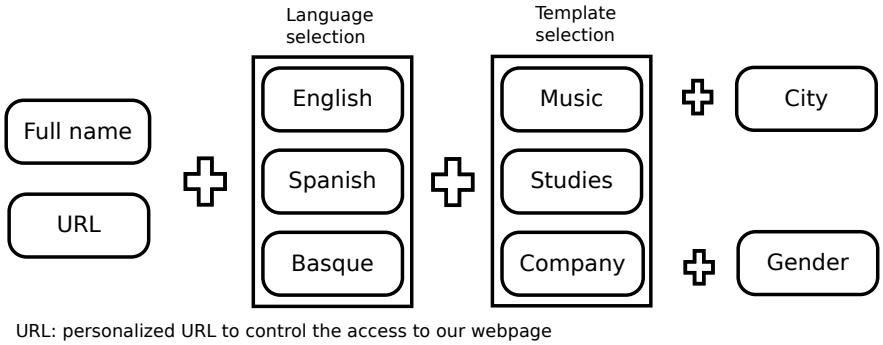

Fig. 3. Email templates

For better customization, we have also used some profile fields such as the language, the name of each user, the gender, and the city in some of the templates. Each template includes a customizable URL that will track the action made by the spam recipient.

Website. Access to this site should only come from the users personal email. The Website is defined to store information about which user, and from which specific spam message has reached the site. Considering all these details, we decided that the most appropriate way was to introduce parameters in the URL which will be included in emails. When the user clicks on the URL and reaches the site, these parameters are stored in our database. The Web also gives the user the possibility to write a comment or to unsubscribe from the system so that she will no longer receive emails. Maintaining the user subscribed to our system gives us the possibility to perform future experiments.

4 Experimental Results

We have performed two separate experiments. In order to generate a baseline, we sent typical spam from a classical spam text in order to measure the success rate, taking into account that spam could have been detected and filtered by the email service, Internet Service Provider, email client in the users computer, or it could have been ignored or deleted by the user. In the second experiment, we focused on personalized spam, in order to prove the click through rate obtained, sending a bigger amount of personalized spam. In both experiments the users were assigned randomly to each of them. The results of each experiment, and explanation thereof will be explained in the next two sections. The comparison of the results is discussed in the last subsection.

4.1 First experiment: typical spam

Using multiple email accounts and sending a total of less than one hundred emails per day in order to avoid mail client's restrictions, we sent a typical spam

email. The account change is due to a strategy to make things more difficult to spam detection systems. We sent one of those emails where spammers try to draw the receiver's attention to enter a web address. To write this email, we read different emails received in our personal email address and we wrote a similar one. In total, we sent 972 typical spam emails, and results are shown in table 2.

Sent emails	Website visits	Percentage
972	4	0.41%

Table 2. Results of the first experiment

As it can be seen, only four users reached our website address. This means that the click-through of the typical spam in our experiment is 0.41%.

4.2 Second experiment: personalized spam

In this case, instead of sending typical spam, we sent a personalized emails to 2,889 Facebook users' email addresses. We used the same experiments setup as in the first experiment for the message delivery, and we sent each template from different email accounts. In order to avoid source address blocking, we sent less than one hundred emails per day and account.

As we mentioned previously, we used three different templates in our study. Those templates had a personalized URL to obtain details of each sent email. Note that the website described the experiment, apologizing for the inconvenience caused, and left space for users comments.

Type	Amount	Percentage
Music	1,787	61.85%
Studies	843	29.18%
Company	259	8.97%
Total	2,889	100%

Table 3. Number of sent emails

The previous table shows the amount of emails sent, and the their distribution among the generated profile templates. A 'Music' profile-based template was the most commonly used. More than the 60% of the personalized mails encouraged the user to visit an URL regarding their favourite music preferences. Our personalized spam campaign model first checked if the user had music preferences added to her profile. If not, it checked for a past studies profile, and so forth.

As we can see in the table 4, 220 users have accessed the website. This is 7.62% of the people that received a personalized email. Also note that 1.38% of people have discharged from the study.

	Amount	Percentage of total shipments
Users who have accessed the website:	220	7.62%
Users who have been unsubscribed:	40	1.38%
Users who have left comments:	11	0.38%

Table 4. Website data

Moreover, we break down the answers taking into account the different templates. Table 5 shows the website accession from each of the templates sent.

	Access to the website	Percentage of total accesses	Click-through
Music	111	50.45%	6.21%
Studies	81	36.82%	9.61%
Company	28	12.73%	10.81%

Table 5. Information according to each template

As we can see in the table, most of the users who acceded our website, received a music-related spam message. This can be considered as expected, as music-based templates involve the 61% of the whole campaign. But it is worth highlighting that the 'Company' or work experience-based template got higher click-trough rate, while the music-based one obtained the lowest one. Otherwise, the musical template had the lowest click-through rate.

4.3 Comparison between experiments

	Sent	Answered	Percentage
Typical spam:	972	4	0.41%
Personalized spam:	2,889	220	7.62%

Table 6. Comparison between results

Table 6 summarizes the response rates obtained while using the different spam types. If we analyze these data, the first interesting information that emerges is that only 4 people have gone through the typical spam. In contrast, 220 other people have come through personalized email. I.e. 0.41 percent compared to 7.62 percent. Authors hypothesize that one reason might be that typical spam can be filtered by most of the spam detection systems. Even so, 0,41% is still many times higher than the rate shown in [1].

4.4 Statistical comparison between genders

Finally, in order to see if there are differences between the behaviour of men and women, several statistical comparisons are carried out.

First, we analyzed the gender of the users related to a certain Facebook account to know how many people leave this information publicly available on OSNs. Second, using the emails sent, we are able to extract the information about the amount of women and men that received our email (we didn't send emails to every Facebook account owner, only to a randomly selected ones). Third, we obtain the gender of the users that accessed to our website following the parameterized URL inserted in the emails. And finally, we extract the website visits over the total emails sent per gender. All the mentioned information is presented in the following table.

Gender	FB account	Sent	Website visits	Relative click-through
Man	38.78%	51.96%	56.70%	6.33%
Woman	27.32%	32.47%	26.34%	4.70%
Unknown	33.90%	15.58%	16.96%	6.32%

Table 7. Comparison between genders

The table above shows that the number of men that we could correlate with a Facebook account is significantly higher. Specifically, number of men is 12 percentage points larger than women on our study.

Moreover, if we compare the percentage of sent emails and the visits to our website, it is possible to see that men click more on the URL than women (6.33% vs 4.70%).

Once the general behaviour is shown, information of the sent emails and visits to the website are presented 8 divided into the previously defined three different templates.

	Sent			
Gender	Total	Music	Studies	Company
Man	1502	57.39%	33.49%	9.12%
Woman	937	69.37%	23.91%	6.72%
Unknown	450	61.11%	25.78%	13.11%
	Click-through (%)			
Man	8.26%	6.61%	9.74%	13.14%
Woman	6.30%	4.46%	10.27%	11.11%
Unknown	8.22%	9.09%	7.76%	5.08%

Table 8. Comparison between genders divided in templates

The difference between the behaviour of men and women is reflected in each template, while the click-through of men is bigger using 'Company' (13.14% vs 11.11%) and 'Music' (6.61% vs 4.46%) templates, using the 'Studies' template the opposite is shown (9.74% vs 10.27%). Presented results show the huge difference in the behaviour of women depending on the template. While the click-

through of 'Studies' and 'Company' templates are similar (10.27% and 11.11%), the result obtained with the 'Music' template is 5 points lower (4.46%).

5 Countermeasures

After the whole experimentation and results discussion, we consider three ways to avoid spam customization. Two from the OSNs point of view, and the other from the users perspective.

- *Limiting users public information:* OSNs may limit public information from users. Thus, it might be more difficult to extract information from users. And the attackers can not use this information in their attacks. This is an obvious countermeasure, but authors consider that goes in the opposite position of what OSN owners seek to attract more attention.
- *Changing the code of the website:* At the time of writing this paper it is possible to collect information from the source code of the Facebook web page. If they change the website and do not leave the user information in a extractable format (for instance: images), it will be more difficult to obtain information for attackers. An interesting research line could be the use of code randomization that could evade automatic web page scrapping.
- *Raising Awareness:* We must teach people how dangerous it can be to leave personal information publicly. If people minimize their profiles public information will be much more difficult to customize the spam.

6 Ethical Considerations

Some actions taken in this paper are ethically sensitive. For some people, collecting information from the Internet is not ethically correct. But as was discussed in [11, 12] and more recently in [5], the best way to do an experiment is to do as realistically as possible. We defend this mode of action for the following reasons.

First, we must be clear that we work to improve the safety of users, we use users information to protect them in the future. Second, we only use information that users displayed publicly in OSNs. This means that we never attacked any account, password or private area. Third, attackers use this information, if we use the same information and act in the same way, we will defend users better.

Finally, we have consulted to the leadership of our university and they have given us the approval. For this, we proposed our intentions to the general direction of the university before the experiments took place (spam campaigns), where we showed them the ethical considerations for conducting the study. We also explained them the procedure we had designed to collect personal data and the way we had thought to send emails. Once the R&D Manager gave us the approval, we started with the experiment phase.

7 Conclusions

This work makes clear the issue that could exist if spam campaign creators turn their spam templates into a personalized text based on user characteristics, interests, and motivating subjects. Attackers have millions of email addresses stored. We have demonstrated that a 19% of the collected email addresses have a corresponding Facebook account associated to it. Moreover, basic public information can be extracted from those users, which is sufficient to create personalized email subject and bodies. These emails can have a click-through rate higher than a 7.62%, being this more than 1,000 times higher than typical spam campaign rate s as shown in [1]. It is obvious that in parallel to the research of new techniques for spam detection inside OSNs, it is necessary to perform research beyond the state of art of classic spam filtering, taking into account the possibility of personalized spam campaign success.

Regarding the behaviour of OSN users analyzed, we found that most of the Facebook users choose their favourite music band and leave it public. We could also see that 30% of users who have some data in their Facebook profile, have at least one company with which they have been connected.

Another interesting fact is that there are more men than women associated with their email to Facebook. The number of men is 12 percentage points higher than the number of women. There is also a difference between genders in the click-through rate, while women react in a rate of a 4.70%, men do it in a 6.33%. Consciousness differences and gender psychological reasons might arise to explain this fact.

The main conclusion to be drawn is that we can develop advanced techniques for generating personalized mail that circumvent current spam detection systems. Clear examples of this are the results shown in the results section. In the first experiment, we can see that only the 0.41% of users have bitten the bait. Whereas in the second 7.62% of the users have entered to the project website. The second result rate is more than 18 times higher than the first one.

We can see that it is not a large number of people, but as a steady stream of visitors, which means that personalized emails reach their destination. Then, once the message is on the user's email inbox, it depends on each person's behaviour to click on the link that is sent in the mail. This shows that spam is not blocked as it's customization have not been detected.

Acknowledgments. This work has been partially funded by the Basque Department of Education, Language policy and Culture under the project SocialSPAM (PI_2014_1_102).

References

1. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G.M., Paxson, V., Savage, S.: Spamalytics: an empirical analysis of spam marketing conversion. In: Proceedings of the 15th ACM conference on Computer and communications security. CCS '08, New York, NY, USA, ACM (2008) 3–14

2. Ezpeleta, E., Zurutuza, U., Hidalgo, J.M.G.: An analysis of the effectiveness of personalized spam using online social network public information. In: International Joint Conference - CISIS'15 and ICEUTE'15, 8th International Conference on Computational Intelligence in Security for Information Systems / 6th International Conference on EUropean Transnational Education, Burgos, Spain, 15-17 June, 2015. (2015) 497–506
3. Gupta, S., Gupta, P., Ahamad, M., Kumaraguru, P.: Abusing phone numbers and cross-application features for crafting targeted attacks. CoRR **abs/1512.07330** (2015)
4. Bonneau, J., Anderson, J., Danezis, G.: Prying data out of a social network. Social Network Analysis and Mining, International Conference on Advances in (2009) 249–254
5. Balduzzi, M., Platzer, C., Holz, T., Kirda, E., Balzarotti, D., Kruegel, C.: Abusing social networks for automated user profiling. In: Proceedings of the 13th international conference on Recent advances in intrusion detection. RAID'10, Berlin, Heidelberg, Springer-Verlag (2010) 422–441
6. Polakis, I., Kontaxis, G., Antonatos, S., Gessiou, E., Petsas, T., Markatos, E.P.: Using social networks to harvest email addresses. In: Proceedings of the 9th annual ACM workshop on Privacy in the electronic society. WPES '10, New York, NY, USA, ACM (2010) 11–20
7. Erlandsson, F., Nia, R., Boldt, M., Johnson, H., Wu, S.F.: Crawling online social networks. In: Network Intelligence Conference (ENIC), 2015 Second European. (Sept 2015) 9–16
8. Stringhini, G., Kruegel, C., Vigna, G.: Detecting spammers on social networks. In: Proceedings of the 26th Annual Computer Security Applications Conference. ACSAC '10, New York, NY, USA, ACM (2010) 1–9
9. Gao, H., Hu, J., Wilson, C., Li, Z., Chen, Y., Zhao, B.Y.: Detecting and characterizing social spam campaigns. In: Proceedings of the 17th ACM conference on Computer and communications security. CCS '10, New York, NY, USA, ACM (2010) 681–683
10. Zheng, X., Zeng, Z., Chen, Z., Yu, Y., Rong, C.: Detecting spammers on social networks. Neurocomputing **159** (2015) 27 – 34
11. Jakobsson, M., Johnson, N., Finn, P.: Why and how to perform fraud experiments. IEEE Security and Privacy **6**(2) (2008) 66–68
12. Jakobsson, M., Ratkiewicz, J.: Designing ethical phishing experiments: a study of (ROT13) rOnl query features. In: WWW '06: Proceedings of the 15th international conference on World Wide Web, New York, NY, USA, ACM (2006) 513–522