Understanding Users' Decision of Clicking on Posts in Facebook with Implications for Phishing

Sovantharith Seng
Rochester Institute of Technology
sovantharith.seng@mail.rit.edu

Mahdi Nasrullah Al-Ameen
Clemson University
malamee@clemson.edu

Matthew Wright
Rochester Institute of Technology
matthew.wright@rit.edu

***Abstract*—Facebook, the largest social networking site (SNS) with over one billion active monthly users, has been woven into the everyday life of many people. While this platform has drastically improved how we interact with one another, it has also opened up a multitude of security and privacy issues. For example, online attackers are increasingly employing phishing attacks on Facebook, seeking to fool their victims by posing as friends using fake or compromised accounts. These attacks are hard to recognize by the Facebook defense system and users alike, and few studies give any insight into how users interact with such attacks. In this study, we take the first step to understand how users react and decide whether to click when they encounter SNS posts with links, including possibly suspicious links. We found that users decide to interact with shared contents based on their relationship with the post author (from whose account the post is shared; perhaps compromised). At the same time, they mostly ignore the location of the shared post (e.g., post author's wall or target user's wall), and any context pointing to a post possibly being suspicious. We also explored the potential of showing a visual warning for suspicious posts. Although our simple warning system failed to prevent users from clicking on suspicious posts altogether, it did reduce the likelihood of users clicking on such posts. Based on our findings, we identified the scope of future work to protect users against phishing attacks in SNSes.**

I. INTRODUCTION

Phishing is a digital security attack that exploits human errors in online navigation [1]. According to the report from Anti-Phishing Working Group [2], 1.2 million phishing attacks occurred in 2016, a 65% increase over 2015. Traditionally, email has been the most common medium to launch phishing attacks. However, social media scams have now become a common type of phishing attack in many countries, including Brazil where it is reported to be the most common form of phishing attack [2].

Social networking sites (SNSes), especially Facebook, have become an integral part of life for many people. Despite the complex use and plethora of information on these networks, many users lack in knowledge and awareness about how to navigate them securely [1], [3], [4]. Phishing attacks on social networking sites are often spear phishing attacks using information on the victim's account, and typically use either fake or compromised accounts of real friends of the victim [5]. The attacker's goal is usually to either collect login credentials from the victim to access her online accounts or have her visit a malicious site with a drive-by download [6]. In SNSes, attackers can improve the chance of their posts being clicked through using a link shortener (e.g., bitly.com) or specialized obfuscation services.[1]

Workshop on Technology and Consumer Protection (ConPro 18)
24 May 2018, San Francisco, CA, USA
www.ieee-security.org

Although there has been some research on phishing attacks in social networking sites, primarily on Twitter [7], [8], we know little about how users interact with potentially malicious posts and what could influence them to click a suspicious link. Our study aims to address this gap by understanding the user's behavior and thought process, so as to devise a more effective defense mechanism that can protect users from sophisticated phishing attacks in SNSes. In particular, we report the impact of different aspects of a post (e.g., types of topics, the sharing location of a post in SNS, etc.) on a user's decision of clicking on that link.

Our study finds that users are not adept at detecting suspicious posts shared on their Facebook newsfeed, where users' decisions to click on the posts are primarily based on their relationship with the *post author* (from whose account the post is shared, possibly compromised). They mostly ignored the location of the post (e.g., the post author's wall or the target user's wall), and any potential mismatches between the post author's interests and the posted topic, where no users tried a mouse-over on the link to look at the destination URL. We also found that the inclusion of a simple visual warning for malicious posts could reduce the likelihood of clicking on such posts, though it did not deter some users. Finding effective solutions to SNS phishing remains an open problem.

II. LITERATURE REVIEW

A. Phishing in SNSes

SNSes provide a fertile ground for attackers to plant sophisticated attacks in front of unsuspecting victims due to the wealth of available personalized information shared, the underlying level of trust in fellow users, and the availability of multiple communication channels. Dhamija et al. showed that the success rate of phishing attacks through email is inversely correlated with the user's knowledge of phishing and security and positively correlated with the level of authenticity

[1] https://apps.lazza.dk/facebook

in the look and feel of the spoofed website [1]. Unfortunately, since SNSes cater to a broad user base and offer a seemingly trustworthy platform in which links can be posted, the risks for users are likely to be high.

Indeed, Vishwanath [9] estimates that attacks on Facebook have an approximately 40% success rate, compared to a success rate of just 1% for traditional email phishing. The findings from his study [9] indicate that attackers typically either post malicious links on a newsfeed, mimicking something of interest to the victims or personally contact the victims through a private message.

Alam et al. [10] noted that the success of targeted phishing is correlated with the amount of information the attacker has. Therefore, if an attacker is a friend with the victim or uses a compromised account of a friend of the victim, he may have little difficulty in fooling the victims without getting noticed. Since SNS users expose a lot of personal information through the site, particularly to their connections, the high success rates reported by Vishwanath [9] may be considered unsurprising.

B. Users' Vulnerabilities

Phishing attacks are successful in SNSes like Facebook due to users' lack of security knowledge and how Facebook is used [5]. For example, some Facebook users gain gratification from the site by either social surfing, finding more information about other people, or expanding their social network [11]. To find others and be found, users may fill out information on their profile and tailor their privacy settings to reach a wider audience. By doing so, these users are providing more information to the phishing attackers and exposing themselves as vulnerable targets [11].

Additionally, users who are receptive to new connections may also be vulnerable to accepting friend requests and messages from attackers posing as legitimate users. Furthermore, users with a large number of friends may be more vulnerable to interacting with unknown strangers or unaware that their friends' accounts have been compromised. The study of Patil [12] found that 40% of users would accept a fake account request. In a separate study, Boshmaf et al. [13] developed the Socialbot Network, a group of adaptive social bots that tricked up to 80% of Facebook users into accepting their friendship requests.

To the best of our knowledge, however, no study has examined whether users treat links from these fake accounts the same as those from accounts connected to them based on relationships that extend beyond Facebook. It is also unexplored whether and how users are looking for indicators of compromised accounts or fake posts. These are essential questions that we begin to address in this work.

III. METHODOLOGY

Phishing posts are designed to trick users, and thus, they usually resemble genuine posts. However, the type and contents of a malicious post shared from a compromised account may not be in line with the interests or sharing behavior of the post author (from whose account the post is shared). In this study, we examine the impact of various factors on the likelihood of a user clicking on a post in her Facebook-like SNS, where we can treat each of these factors as a variable. The variables in the study include: i) The relationship between the participant and the post author, ii) The topic of the post, iii) Whether the topic of a shared post matches the interests of the post author, and iv) Whether the post is on the participant's wall or the post author's (i.e., compromised account's) wall.

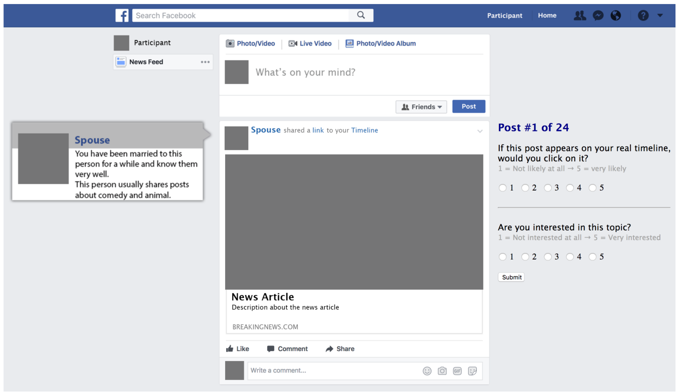

Spouse

You have been married to this person for a while and know them very well.
This person usually shares posts about comedy and animal.

Spouse shared a link to your Timeline

News Article

Description about the news article

BREAKINGNEWS.COM

Post #1 of 24

If this post appears on your real timeline, would you click on it?
1 = Not likely at all → 5 = very likely

○ 1 ○ 2 ○ 3 ○ 4 ○ 5

Are you interested in this topic?
1 = Not interested at all → 5 = Very interested

○ 1 ○ 2 ○ 3 ○ 4 ○ 5

Fig. 1. Mock-Up of the simulated interface for Part I

The study was conducted in two parts, Part I and Part II, which differed regarding interface design and interview questions. The mock-up of the simulated interface resembles the Facebook newsfeed, shown in Fig. 1 for Part I and Fig. 2 for Part II. The posts shown in our interface varied regarding the variables listed above.

A. Data Collection.

We collected data from five different sources in this study, including audio recording, participants interactions with the simulated interface, researchers observations, a survey questionnaire, and a semi-structured interview.

The audio recording provides qualitative data on the thought process of participants while navigating an SNS. We transcribed the audio recordings after the study. The information on participants' interactions with the SNS is primarily a Likert-scale response on the likelihood of clicking on a link in a post, which demonstrates how successful the user is at spotting a potential phishing post using the information provided in our simulated interface. The researchers observations complement the information collected from the interaction data and the audio recording. These observations include non-verbal communication and signs that the participants display during the study, as well as any questions that they ask before and after the study. Finally, the surveys and the interview provide information on participants' demographic and background, SNS usage, their anecdotal experience on social networks, and their self-efficacy in using SNSes and the Web securely.

B. Interface Design

We designed a mock-up of an SNS interface that resembles a Facebook newsfeed (Fig. 1). To serve the purpose of this study, we changed several aspects of the interface: i) Elements on the right-side of the newsfeed, including chat bar, news updates, and the advertisements are removed to make space for Likert-scale questions that the participants are required to answer for each of the shown posts; ii) The profile picture

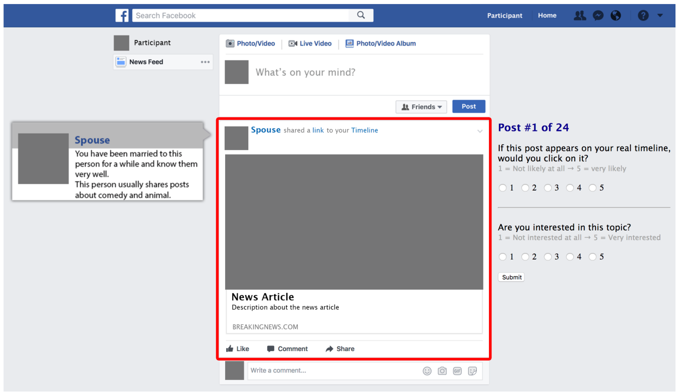

Spouse

You have been married to this person for a while and know them very well.
This person usually shares posts about comedy and animal.

Spouse shared a link to your Timeline

News Article

Description about the news article

BREAKINGNEWS.COM

Post #1 of 24

If this post appears on your real timeline, would you click on it?
1 = Not likely at all → 5 = very likely

○ 1 ○ 2 ○ 3 ○ 4 ○ 5

Are you interested in this topic?
1 = Not interested at all → 5 = Very interested

○ 1 ○ 2 ○ 3 ○ 4 ○ 5

Fig. 2. Mock-Up of the simulated interface for Part II

of each Facebook friend is presented by a solid color box to control for various biases [12]; iii) The name of each Facebook friend is presented in a less specific way, like 'Parent' or 'Best Friend', providing a simplified presentation of the relationship with the participant and thereby controlling for variation in particular relationships; and iv) There is no 'preview details' (e.g., images), 'reaction', or 'comment' shown with any post, since each of them presents a possible variable to study, which is currently out of the scope of this work. Furthermore, preview details are replaced with generic sentences.

In the simulated newsfeed, the participant's name is shown as just 'Participant' with having a solid gray box as the profile picture. There is a dialog box on the left of a post showing the details about the user from whose account the post is shared, to give the participants some background information about the post author. We provided this to simulate the fact that SNS users usually have some background information of the person sharing a post as they are scrolling through their newsfeed.

In our study, the participants could hover over a post to see the link's destination (i.e., URL). In both parts of the study, the phishing and non-phishing posts had a similar appearance. However, phishing posts had a destination link that differed from what was shown in the preview. The topic of each phishing post was also a mismatch with the listed interest of the post author, representing a post from a compromised account.

1) Visual Warnings: In Part II of our study, we explore using a simple warning system to show users that a post is suspicious. Different indicators could be used to identify suspicious posts, including the use of URL redirection and URL shorteners, or mismatches between the post author's interests and the topic of the shared post (stylometry [14], [15]).

In our study, each mismatched post is marked with a colored border (orange or red), as shown in Fig.2, which serves as a warning to the users. The prior study [16] showed that using a colored border motivates users to focus on the content of a warning system in the Web environment. Since we want to observe the initial response of the users rather than to extensively study the effectiveness of the warning system or its components in SNSes, we opted to employ a rudimentary warning indicator inspired by the Web of Trust service [2]. We informed our participants that the warning system is a third-party program currently in development by a university, where a colored border around a post indicates that post to be a suspicious one, and a red border denotes more suspicion than an orange border.

C. Procedure

We conducted the study in a lab setting with a think-aloud protocol, where each part of the study lasted between 30 to 60 minutes per participant and was audio-recorded for transcription and analysis. At the beginning of the study, the participants were given a consent form. Once they had fully understood and signed the form, they responded to a survey questionnaire on their demographics and SNS usage. Afterwards, we explained to the participants how the think-aloud protocol would go and how the simulated interface for our study works. We leveraged deception in our study: at the beginning, we did not inform participants about the goal of our study relating to phishing attacks. Instead, they were told that the study would focus on exploring their Facebook browsing behavior.

Participants were then asked to interact with 24 posts on the simulated SNS interface, where they responded to Likert-scale questions (ranging from 1 to 5; 1: 'Not likely at all', 5: 'Very likely') regarding their likelihood of clicking on each post shown in the newsfeed (see Figures 1 and 2). The posts were presented to the participants in a random order. Once the participants had finished their interaction, they were asked to complete a survey on their real-life SNS behavior and self-efficacy questionnaires on the secure use of SNSes and of the Web, respectively. The session was concluded with a short interview about the user's experience, expanding on their survey feedbacks. Participants were then thanked and compensated with a $20 gift card for their time.

D. Participants

We recruited participants by posting fliers on campus and sending email to a university listserv. The criteria for participation were that the students must be at least 18 years old and have a Facebook account at the time of the study. The participants of Part I were not eligible to participate in Part II.

We had a total of 30 participants, 15 in each part of the study. In Part I, there was eight women and seven men, and their ages ranged from 19 to 37 (median: 23). Eight were majoring in IT or related fields. In Part II, there were four women and 11 men, ranging in age from 18 to 29 (median: 23). We had seven majors from IT or related fields. At the end of the study, all the participants were informed about the phishing aspect of the study. The researchers also answered any questions from the participants regarding the topic.

IV. RESULTS

In this section, we present the results of our study, where the scores are based on a 5-point Likert-scale response, denoting the likelihood of clicking on a link shown in our SNS interface. The higher the score, the more likely the user

[2] https://www.mywot.com

would click on that link. When comparing two conditions, we consider the difference to be significant when the p-value is less than 0.05. We note that the sample sizes are fairly small for statistical tests, such that the lack of a finding of significance could mean that we simply did not find an effect.

A. Mismatch with the Interests of the Post Author

In both parts of our study, half of the posts presented to the participants were a mismatch with the listed interest of the post author. In Part I, there was no apparent visual distinction between the matched and mismatched posts, while in Part II, the mismatched posts were shown with a warning indicator—an orange or red border around the post (see Fig. 2).

In Part I, the average score for the mismatched and matched posts were 3.60 and 3.57, respectively, where only six out of 15 participants even noticed the mismatch. There is no significant difference between the two scores, and we found no significant difference between scores for matched and mismatched posts in any of the splits that we examine below (relationship, sharing location, topic, etc.). While a lack of significance could be due to small sample sizes, the average scores are quite close in absolute terms, and any difference would indeed be surprising for the nine participants who did not notice the mismatch at all. Interestingly, two of the six participants who noticed the mismatch reported *higher* interest in the post because of the mismatch. They said that they would want to know what is so interesting in the post that prompted the post author to share it, even though the content in the post does not match with the post author's general interest.

In Part II, we found a significant difference between the scores of matched (3.3) and mismatched (2.3) posts ($p = 0.0004$), where four of the 15 participants reported explicitly noticed the mismatch in topics. Note that the mismatched posts all had either a red or orange warning border around them. We did not find a significant difference in scores between the posts with orange (2.6) and red warnings (2.0). However, we did find a significant difference in scores between the posts without a warning and the ones with a warning: orange ($p = 0.0288$) and red ($p = 4.4 * 10^{-5}$).

B. Relationship with the Post Author

In Part I, the average score (4.0) for posts shared by people in a close relationship with the participant was significantly higher than the score (3.2) for posts shared by others ($p = 5.7{*}10^{-8}$). We also found a significant difference ($p = 0.0130$) in Part II when comparing the scores for close relationships (3.2) and non-close relationships (2.4).

In Part II, we found significant differences between the scores of matched (no warning) and mismatched posts (with warnings) for both close ($p = 0.0046$) and non-close relationships ($p = 0.0011$).

All of the participants noted that they would click on a post shared by someone who is close to them, especially if they would likely be discussing it later with that person. They also mentioned that certain relationships would add value to the posts, which might make them more likely to click on them, e.g., posts shared by their parents. A few participants mentioned being interested to click on a post shared from a non-close relationship (especially colleagues and classmates) since they might want to know more about them.

C. Sharing Location of a Post in SNS

In this study, we examined the effect of two posting locations in the SNS: i) On the post author's wall and ii) On the user's wall. In Part I, posts shared on the user's walls scored an average of 3.7, compared to 3.5 for posts shared on the post author's wall, with no significant difference between them. Also, we did not find any significant difference between the score (3.0) for posts shared on the user's wall and the score (2.6) for posts shared on the post author's wall.

In Part II, we found significant differences between the scores of the matched (no warning) and mismatched posts (with warning) on the user's wall ($p = 0.0272$) and the post author's wall ($p = 0.0084$).

D. Topic of the Posts

We used three different topics for the posts shown to the participants: travel, news, and entertainment. In Part I, the scores for posts related to travel, news, and entertainment were 3.8, 3.6, and 3.4, respectively, where we did not find any significant difference between them. In Part II, the scores for posts related to travel, news, and entertainment were 2.8, 3.0, and 2.6, respectively, where the differences again were not significant.

E. Demographics, Background, and Self-efficacy

Gender. In both Part I and Part II, we did not find a significant difference between the scores of female (Part I: 3.8, Part II: 3.0) and male participants (Part I: 3.4, Part II: 2.7).

In Part II, although we did not find a significant difference between the scores of matched (no warning) and mismatched posts (with a warning) in the subset of female participants, we did find a significant difference in the male subset ($p = 0.0001$).

Age. The median age of our participants was 23 in both Part I and Part II of our study. In Part I and Part II, respectively, the average scores of the participants was 3.7 and 2.8 for those 23 or older and 3.5 and 2.5 for those younger than 23. We did not find a significant difference in either case.

Technological Background. In both Part I and Part II, we did not find a significant difference between the scores of the participants from non-tech (Part I: 3.8, Part II: 3.0) and tech backgrounds (Part I: 3.3, Part II: 2.7).

In Part II, we did find a significant difference between the scores of matched (no warning) and mismatched posts (with warning) in both subsets of tech ($p = 0.0454$) and non-tech ($p = 0.0093$)backgrounds.

Self-efficacy. We asked participants security self-efficacy questions in using SNSes and the Web. We found an inverse correlation between security self-efficacy in using SNSes and the likelihood of clicking on a post (Part I: $r = -0.6$, Part II: $r = -0.4$). We found similar results as we focused solely on mismatched posts (Part I: $r = -0.6$, Part II: $r = -0.4$), where the higher the self-efficacy, the less likely it is for a participant to click on a mismatched SNS post.

F. Hovering over the Post

Hovering over a post helps users to identify if the actual URL of the post differs from the displayed link, a crucial step to determine a suspicious post. We observed that only three participants hovered over the posts in Part I, and no one did in Part II. Participants said that hovering was not their habit, and instead, they judged the authenticity of a shared link by clicking it and examining the look-and-feel of the corresponding site. Over half of the participants reported being unaware of the risk of drive-by downloads.

V. DISCUSSION

In this section, we first discuss the limitations of our study before presenting the implications of our findings and comparing with prior work.

A. Limitations

The sample size in our study is small and consists of only undergraduate and graduate students through a self-selection recruitment method. Thus, our findings do not generalize to the entire population. The user interface in our study does not represent the exact look and feel of a real Facebook newsfeed, most notably as it lacks the details of the post and real people that the participants would see on their newsfeed. Since the purpose of the study was to determine whether users would react to the available cues (mismatches in topic, interest, and URLs) and our warnings, we sought to limit any confounding information. We also needed to limit the number of independent variables and keep the study design simple for this preliminary work (see §III-B for details).

Furthermore, we are aware that self-reported data may not be an exact reflection of how users react to SNS posts in a real-life setting. This simplified setting represents a best-case scenario for protecting users, so to the extent that our findings indicate that users do not react strongly to such cues and warnings, the real-world case is likely worse. In our future work, we will conduct a field study with a larger and more diverse population set using the real Facebook interface.

B. Factor Influencing Users' Clicking Decisions

The relationship with the post author is a notable factor that influences users' clicking decisions in SNSes, where participants reported a significantly higher likelihood to click on the posts shared by close friends and family. It also demonstrates the advantages an attacker could gain by sharing posts from compromised accounts, instead of sharing a post from a new account of his own.

C. Showing Warnings in SNSes

The comparison between the results from Part I and Part II of our study illustrate that showing a warning (in either red or orange) reduced the overall likelihood of clicking on mismatched posts. Irrespective of the relationship with the post author, the sharing location of a post, or technological background of participants, showing a warning significantly reduced the likelihood of clicking on mismatched posts. We found that the presence of a visual warning slowed down the users and made them pay more attention to the details, which might help them to detect attacks.

Despite the warnings, however, we found that several participants still had a high likelihood to click on mismatched posts, despite the lack of information and details presented in this simulated interface; We note that these users could be even more strongly motivated to click on the posts in real life. Thus, a useful detection and warning system may need to more aggressively stop users to protect them from the most suspicious posts, much as SSL warnings in browsers have evolved to be more aggressive to improve adherence [17]. We plan to explore this in future work.

D. Comparison with Prior Studies

The prior study [18] examined the technological backgrounds of users in understanding their susceptibility to phishing attacks on the Web. Our study presents a preliminary result on the phishing susceptibility of users in online social networking sites (e.g., Facebook), and found similarity with prior findings regarding users' technological background.

In our study, the participants from tech backgrounds reported lower likelihood to click on mismatched posts, as compared to those from non-tech backgrounds, although the difference was not significant in this regard. We also found that the higher the security self-efficacy of a participant, the less likely it is for her to click on a mismatched post. These findings are in line with that from existing literature illustrating the higher susceptibility of non-expert users to phishing attacks on the Web, as compared to expert users [18].

VI. CONCLUSION

Social networking sites (e.g., Facebook) have become a part of everyday life for billions of people all over the world. Unfortunately, they have also become a valuable target for phishing attacks due to several reasons, including users' low security knowledge and the wealth of personal information available for launching targeted attacks. In this study, we took the first step into understanding the factors influencing the decision of a user about clicking on a post in her Facebook newsfeed. We also explored the potential of using a visual warning for suspicious posts. Although our simple warning system failed to stop users from clicking on malicious posts entirely, we found an overall decrease in the likelihood of clicking on such posts in the presence of the warning. In our future work, we plan to develop a robust system to defend SNS users against phishing attacks, by investigate how the warning can be presented to users in an effective and non-intrusive way.

VII. ACKNOWLEDGEMENT

This project is supported by start-up funds provided by Rochester Institute of Technology. We would like to thank all the reviewers for their valuable feedbacks on this work.

References

[1] R. Dhamija, J. D. Tygar, and M. Hearst, “Why phishing works,” in *Proceedings of the 24th SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI ’06. Montreal, Quebec, Canada: ACM, 2006, pp. 581–590. [Online]. Available: http://doi.acm.org/10.1145/1124772.1124861

[2] A. phishing Working Group (APWG), “Phishing activity trends report (fourth quarter, 2016).” [Online]. Available: https://docs.apwg.org/reports/apwg_trends_report_q4_2016.pdf

[3] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer, “Social phishing,” *Communication of the ACM*, vol. 50, no. 10, pp. 94–100, Oct. 2007. [Online]. Available: http://doi.acm.org/10.1145/1290958.1290968

[4] M. Tsikerdekis and S. Zeadally, “Online deception in social media,” *Communication of the ACM*, vol. 57, no. 9, pp. 72–80, Sep. 2014. [Online]. Available: http://doi.acm.org/10.1145/2629612

[5] A. Vishwanath, “Getting phished on social media,” *Decision Support Systems*, vol. 103, no. Supplement C, pp. 70–81, Nov. 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167923617301690

[6] E. Kirda and C. Kruegel, “Protecting users against phishing attacks with AntiPhish,” in *Proceedings of 29th Annual International Computer Software and Applications Conference (COMPSAC’05)*, vol. 1, Edinburgh, Scotland, Jul. 2005, pp. 517–524 Vol. 2.

[7] A. Aggarwal, A. Rajadesingan, and P. Kumaraguru, “PhishAri: Automatic realtime phishing detection on twitter,” in *Proceedings of 2012 eCrime Researchers Summit*, Fajardo, PR, USA, Oct. 2012, pp. 1–12.

[8] S. Chhabra, A. Aggarwal, F. Benevenuto, and P. Kumaraguru, “Phi.sh/$ocial: The phishing landscape through short urls,” in *Proceedings of the 8th Annual Collaboration, Electronic Messaging, Anti-Abuse and Spam Conference*, ser. CEAS ’11. Perth, Australia: ACM, 2011, pp. 92–101. [Online]. Available: http://doi.acm.org/10.1145/2030376.2030387

[9] A. Vishwanath, “Habitual facebook use and its impact on getting deceived on social media,” *Journal of Computer-Mediated Communication*, vol. 20, no. 1, pp. 83–98, 2015. [Online]. Available: http://dx.doi.org/10.1111/jcc4.12100

[10] S. Alam and K. El-Khatib, “Phishing susceptibility detection through social media analytics,” in *Proceedings of the 9th International Conference on Security of Information and Networks*, ser. SIN ’16. Newark, NJ, USA: ACM, 2016, pp. 61–64. [Online]. Available: http://doi.acm.org/10.1145/2947626.2947637

[11] A. N. Joinson, “Looking at, looking up or keeping up with people?: Motives and use of facebook,” in *Proceedings of the 26th SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI ’08. Florence, Italy: ACM, 2008, pp. 1027–1036. [Online]. Available: http://doi.acm.org/10.1145/1357054.1357213

[12] S. Patil, “Will you be my friend?: Responses to friendship requests from strangers,” in *Proceedings of the 2012 iConference*, ser. iConference ’12. Toronto, Ontario, Canada: ACM, 2012, pp. 634–635. [Online]. Available: http://doi.acm.org.ezproxy.rit.edu/10.1145/2132176.2132318

[13] Y. Boshmaf, I. Muslukhov, K. Beznosov, and M. Ripeanu, “The socialbot network: When bots socialize for fame and money,” in *Proceedings of the 27th Annual Computer Security Applications Conference*, ser. ACSAC ’11. Orlando, FL, USA: ACM, 2011, pp. 93–102. [Online]. Available: http://doi.acm.org/10.1145/2076732.2076746

[14] M. Bhargava, P. Mehndiratta, and K. Asawa, “Stylometric analysis for authorship attribution on Twitter,” in *International Conference on Big Data Analytics*. Springer, 2013, pp. 37–47.

[15] S. Vosoughi, H. Zhou, and D. Roy, “Digital stylometry: Linking profiles across social networks,” in *International Conference on Social Informatics*. Springer, 2015, pp. 164–177.

[16] S. Egelman and S. Schechter, “The importance of being earnest [in security warnings],” in *Financial Cryptography and Data Security*, A.-R. Sadeghi, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 52–59.

[17] A. P. Felt, A. Ainslie, R. W. Reeder, S. Consolvo, S. Thyagaraja, A. Bettes, H. Harris, and J. Grimes, “Improving SSL warnings: Comprehension and adherence,” in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2015, pp. 2893–2902.

[18] I. Ion, R. Reeder, and S. Consolvo, “”... no one can hack my mind”: Comparing expert and non-expert security practices.” in *SOUPS*, 2015, pp. 327–346.