Int. J. Human-Computer Studies 82 (2015) 69–82

Contents lists available at ScienceDirect

Int. J. Human-Computer Studies

journal homepage: www.elsevier.com/locate/ijhcs

Why phishing still works: User strategies for combating phishing attacks[☆]

Mohamed Alsharnouby, Furkan Alaca, Sonia Chiasson[*]

School of Computer Science, Carleton University, 1125 Colonel By Drive, Ottawa, ON, Canada K1S 5B6

ARTICLE INFO

Article history:
Received 3 September 2014
Received in revised form
18 March 2015
Accepted 10 May 2015
Communicated by Scott Bateman
Available online 21 May 2015

Keywords:
Phishing
Eye tracking
Usable security
User study

ABSTRACT

We have conducted a user study to assess whether improved browser security indicators and increased awareness of phishing have led to users' improved ability to protect themselves against such attacks. Participants were shown a series of websites and asked to identify the phishing websites. We use eye tracking to obtain objective quantitative data on which visual cues draw users' attention as they determine the legitimacy of websites. Our results show that users successfully detected only 53% of phishing websites even when primed to identify them and that they generally spend very little time gazing at security indicators compared to website content when making assessments. However, we found that gaze time on browser chrome elements does correlate to increased ability to detect phishing. Interestingly, users' general technical proficiency does not correlate with improved detection scores.

1. Introduction

An important aspect of online security is to protect users from fraudulent websites and phishing attacks. *Phishing* is a "criminal mechanism employing both social engineering and technical subterfuge to steal consumers' personal identity data and financial account credentials" (Anti-Phishing Working Group, 2014a). While advances in the automated detection of phishing websites have resulted in improved security, these automated means are not fool-proof and users must be vigilant in protecting themselves in this arms race (Hong, 2012). According to the Anti-Phishing Working Group, phishing attacks remain widespread: 42,890 unique phishing websites were reported in December 2013, with the financial and online payment sectors accounting for nearly 80% of targeted industries (Anti-Phishing Working Group, 2014a).

Modern web browsers provide tools to assist users in making informed security decisions. For example, visual indicators within the URL bar and the SSL padlock have been designed to allow users to judge the legitimacy of websites. Unfortunately, these indicators have been only partially successful at helping to prevent phishing. Poor usability may allow phishing websites to masquerade as legitimate websites and deceive users into divulging their personal information. Earlier browser security indicators have been shown in previous studies to be ineffective, putting users at a higher risk of falling victim to phishing attacks (Whalen and Inkpen, 2005; Lin et al., 2011; Egelman, 2009).

This is compounded by the fact that security is a secondary task for most users (Whitten and Tygar, 1999). Users who are concentrating on the real purpose of their online interaction, such as making a purchase, are unlikely to notice security indicators. Furthermore, some security indicators are visible only when the website is secure. The *absence* of a security indicator, as is possible with phishing websites, is even less likely to be noticed by users. Therefore, developing usable browser security cues to combat phishing attacks remains an important and unsolved problem in usable security, as is understanding how users make determinations about the legitimacy of websites (Purkait, 2012).

To inform the design of improved techniques against phishing, we explored the strategies employed by users to identify phishing attacks. We showed participants a series of websites and asked them to identify whether each one is legitimate or fraudulent. This paper makes several distinct contributions to the literature. First, we evaluate the effectiveness of recent changes that have been made in web browser designs to help users identify fraudulent websites. Secondly, we assess whether users have developed improved detection strategies and mental models of phishing nearly a decade after Dhamija et al. (2006)'s initial phishing study. And finally, we are the first to use eye tracking data to obtain quantitative information on which visual security indicators draw the most attention from users as they determine the legitimacy of websites. Based on our results, we identify aspects in which web browser security indicators have improved in

[☆] This paper has been recommended for acceptance by Scott Bateman.

[*] Corresponding author. Tel.: +1 613 520 4333.

E-mail addresses: malsharnouby@ccsl.carleton.ca (M. Alsharnouby), falaca@ccsl.carleton.ca (F. Alaca), chiasson@scs.carleton.ca (S. Chiasson).

http://dx.doi.org/10.1016/j.ijhcs.2015.05.005
1071-5819/© 2015 Elsevier Ltd. All rights reserved.

modern web browsers, identify areas for potential improvement, and make recommendations for future designs.

The remainder of this paper is organized as follows: Section 2 reviews related work on phishing detection and tools to aid users in identifying phishing websites. Section 3 details our study methodology. Section 4 provides analysis and interpretation of our quantitative and qualitative data. Section 5 discusses some ideas for future web browser designs, while Section 6 concludes the paper.

2. Related work

Research on protecting users against phishing attacks has taken four complementary approaches: automating phishing detection, providing user interface cues to help users detect phishing, educating users about how to protect themselves, and understanding users' susceptibility to phishing to inform the design of protection mechanisms. Our work falls within scope of the fourth area, but we also provide a brief overview of the other areas to give context to our work. For a general introduction, see Hong (2012)'s article, or for a more complete recent review of the phishing literature, see Purkait (2012)'s literature survey.

2.1. Automated phishing detection

The first line of defense against phishing should be automated detection; users cannot fall for phishing attacks if they never see the attacks. Automatic phishing detectors exist at several different levels: mail servers and clients, internet service providers, and web browser tools. Tools may block access to a detected phishing website and/or request that the website's internet service provider takes down the website (Moore and Clayton, 2007).

Automatic email classification tools commonly use machine learning techniques (Fette et al., 2007), statistical classifiers (Bergholz et al., 2010), and spam filtering techniques (Cormack, 2008) to identify potential phishing messages with varying degrees of effectiveness as the threat continues to evolve. Mis-classifications affect the perceived reliability of the service and users are likely to be quite intolerant to “losing” legitimate messages.

Techniques to detect phishing websites include blacklists, machine learning (Whittaker et al., 2010), URL feature classification and domain name analysis, visual similarity assessment (Fu et al., 2006), contextual analysis and user behavioural prediction (Lee et al., 2014), and crowdsourcing (OpenDNS, 2014). Some blacklists, such as Google's (Whittaker et al., 2010), use automated machine learning. PhishTank (OpenDNS, 2014) offers a blacklist for use by other tools through an API. Its blacklist is populated through crowdsourcing volunteers who submit potential phishing websites and vote on the legitimacy of websites.

Web browsers maintain their own blacklists and heuristics for detecting phishing, displaying warnings to users if they reach a known phishing page. Detection rates have improved considerably over the last 5 years. NSS Labs (2013) conducts independent tests and found that the major browsers had an average phishing detection rate of approximately 90%, with zero-hour block rates above 70%. Third-party add-ons are also available. Sheng et al. (2009) evaluated the effectiveness of eight different browser tools and found them generally slow at detecting new phishing campaigns. This is problematic given that the median lifetime of a phishing campaign is about 12 h (NSS Labs, 2013), with many as short as 2 h.

While successful at stopping a large number of attacks from reaching users, automated methods are insufficient as the sole means of protecting users. Secondary methods involving users are necessary for times when automatic detection fails.

2.2. Security indicators

There have been a number of studies regarding phishing and the usability of browser security cues. Herzberg (2009) provides an overview of several studies.

At its core, phishing is a threat because users are unable to verify the authenticity of the website asking for their credentials. Dhamija and Tygar (2005) first proposed Dynamic Security Skins, a browser extension that allows websites to display a secret image and customizes the browser chrome. Variations of this secret image method have now been deployed by banks and major organizations (e.g., Sitekey Bank of America, 2014; Yahoo Sign-in Seals Yahoo! Inc, 2014). Anecdotal evidence suggests that some users may still fall victim to phishing websites who claim that the image database is down for maintenance or who simply leave out this feature since the absence of a cue may not trigger attention. Many browser toolbars (e.g., Chou et al., 2004; Yee and Sitaker, 2006; Li and Helenius, 2007; Kirda and Kruegel, 2006; Kirlappos and Sasse, 2012) have also been proposed to protect against phishing, each with limited success. User studies by Wu et al. (2006), Li and Helenius (2007), and Li et al. (2014) found that security toolbars intended to prevent phishing attacks were ineffective and identified several usability problems. While users may occasionally pay attention to the indicators, accomplishing that their primary task often gets prioritized, and in these cases users look for visual signs reinforcing the website's trustworthiness rather than heeding warnings to the contrary (Kirlappos and Sasse, 2012). Abbasi et al. (2012) compared users' ability to detect phishing given high- or low-performing browser toolbars and found that users were more successful with the high-performing toolbar. However, users still ignored the toolbar's advice 15% of the time, instead believing that their own intuition was more accurate.

Others have explored the browsers' built-in security indicators. Lin et al. (2011) examined the effectiveness of domain highlighting that is now included in most browsers. They found it to be only marginally successful when users' attention was explicitly drawn to the address bar. Egelman (2009) explored various online trust indicators, including web browser phishing warnings and SSL warnings. They found that 97% of users were fooled by at least one attack, but that active warnings which interrupt users' tasks were more effective than passive warnings.

Although addressing a tangential issue, password managers (Yee and Sitaker, 2006; Ross et al., 2005) can offer protection against phishing by storing both the user's credentials and the legitimate URL at which these credentials should be used. Users attempting to use their password manager at a phishing website will either be warned against a suspicious website or the password manager will supply incorrect credentials.

Efforts to reduce phishing at the email level are also popular, but these typically require minimal user involvement beyond needing to occasionally check spam-filtered mail and potentially update spam filters. Email encryption and digital signing can help protect users against phishing and other attacks, but these are plagued with usability issues and are not widely used (Garfinkel et al., 2005).

2.3. Anti-phishing education

Although educational efforts are unlikely to solve the phishing problem on its own, vigilant users form an important part of the defensive strategy. Both research efforts and public education campaigns (e.g., Anti-Phishing Working Group, 2014b; Government of Canada, 2014) have focused on teaching users how to protect themselves against phishing attacks. PhishGuru (Kumaraguru et al., 2007, 2009, 2010) embeds phishing education within the primary task of receiving phishing email and results show that the educational material is most impactful if delivered immediately after users have