The Highly Insidious Extreme Phishing Attacks

Rui Zhao*, Samantha John†, Stacy Karas†, Cara Bussell†, Jennifer Roberts†, Daniel Six†,
Brandon Gavett†, and Chuan Yue*
*Colorado School of Mines, Golden, CO 80401
†University of Colorado Colorado Springs, Colorado Springs, CO 80918

***Abstract*—One of the most severe and challenging threats to Internet security is phishing, which uses spoofed websites to steal users' passwords and online identities. Phishers mainly use spoofed emails or instant messages to lure users to the phishing websites. A spoofed email or instant message provides the first-layer context to entice users to click on a phishing URL, and the phishing website further provides the second-layer context with the look and feel similar to a targeted legitimate website to lure users to submit their login credentials. In this paper, we focus on the second-layer context to explore the extreme of phishing attacks; we explore the feasibility of creating extreme phishing attacks that have the almost identical look and feel as those of the targeted legitimate websites, and evaluate the effectiveness of such phishing attacks. We design and implement a phishing toolkit that can support both the traditional phishing and the newly emergent Web Single Sign-On (SSO) phishing; our toolkit can automatically construct unlimited levels of phishing webpages in real time based on user interactions. We design and perform a user study to evaluate the effectiveness of the phishing attacks constructed from this toolkit. The user study results demonstrate that extreme phishing attacks are indeed highly effective and insidious. It is reasonable to assume that extreme phishing attacks will be widely adopted and deployed in the future, and we call for a collective effort to effectively defend against them.**

I. Introduction

One of the most severe and challenging threats to Internet security is phishing, which uses spoofed websites to steal users' passwords and online identities. To defend against phishing attacks, researchers have proposed various blacklist-based, heuristics-based, and whitelist-based solutions (Section VI), organizations and communities such as APWG [39] and PhishTank [42] have provided phishing reporting and verification services; many vendors have also provided secure browsing systems such as Google Safe Browsing, Microsoft SmartScreen Filter, McAfee SiteAdvisor, and Norton Safe Web. However, phishing attacks have also been quickly evolving to evade the detection and defense [43], and the battle between phishers and defenders will be long-standing.

Phishers mainly use spoofed emails or instant messages to lure users to the phishing websites. A spoofed email or instant message provides the ***first-layer context*** (e.g., asking for account verification or update) to entice users to click on a phishing URL, and the phishing website further provides the ***second-layer context*** with the ***look and feel*** similar to a targeted legitimate website to lure users to submit their login credentials [37]. In terms of the first-layer context, the success of phishing is mainly limited by two constraints [37]. One is that if phishing emails or instant messages are suspicious, users would not click on phishing URLs and visit the phishing websites [8], [16]. The other is that phishing emails captured by spam filters [32] cannot even reach users in the first place. In terms of the second-layer context, the success of phishing is mainly limited by two other constraints [37]. One is that phishing websites will trigger warnings if they are detected by browsers, thus security-conscious users would not visit them and submit credentials [1]. The other is that if the look and feel of the undetected phishing websites are suspicious, security-conscious users would not submit their credentials [7], [8], [12], [14], [26].

In this paper, we focus on the second-layer context to explore the extreme of phishing attacks. In other words, we explore the feasibility of creating extreme phishing attacks that have the almost identical look and feel as those of the targeted legitimate websites, and evaluate the effectiveness of such phishing attacks.

In particular, we design and implement a phishing toolkit that can support both the traditional phishing and the newly emergent Web Single Sign-On (SSO) phishing [37]. In terms of the traditional phishing, our toolkit can automatically construct unlimited levels of phishing webpages in real time based on user interactions; in terms of the Web SSO phishing, our toolkit can allow attackers to easily construct spoofed Web SSO login "windows" for Gmail, Facebook, and Yahoo. The constructed phishing webpages and Web SSO login "windows" are almost identical to their legitimate counterparts, potentially making it very difficult for users to identify if they are interacting with real or spoofed websites.

The toolkit can be used by attackers to easily construct and deploy extreme phishing attacks; it can also be used by researchers to easily construct testbeds for performing phishing related user studies and exploring new phishing defense mechanisms. In particular, we design and perform a user study to evaluate the effectiveness of the phishing attacks constructed from this toolkit. The user study results based on 94 participants demonstrate that extreme phishing attacks constructed by our toolkit are indeed highly effective, i.e., insidious. The questionnaire results show that **87 (92.6%)** of the 94 participants were actually not suspicious about the extreme phishing websites that they visited, and the observation results show that **91 (96.8%)** of the 94 participants submitted their credentials to the extreme phishing websites; meanwhile, most of those "victims" were aware of phishing before participating in this study or had been susceptible to some phishing attacks in the past. Therefore, it is reasonable to assume that extreme phishing attacks will be widely adopted and deployed in the future, and we call for a collective effort to effectively defend against them.

The main contributions of our paper include: (1) we define and explore extreme phishing attacks and investigate the techniques for constructing them (Section III), (2) we

design and implement a concrete toolkit that can be feasibly and easily used by attackers to construct and deploy such attacks (Section IV), (3) we design and perform a user study with 94 participants to demonstrate the effectiveness of such attacks (Section V), and (4) we discuss the impacts of extreme phishing on existing phishing defense mechanisms and provide suggestions to researchers and users for them to better defend against such attacks (Section VI).

II. Related Work

We review the related work on phishing toolkits and testbeds in this section, and defer the discussion of the related phishing detection and defense techniques to Section VI.

Attackers often use phishing toolkits to construct their phishing websites [12]. Cova et al. analyzed a large collection of free underground phishing toolkits [5], and found that those toolkits target not only users but also inexperienced phishers (through backdoors) as victims. They also found that most of those toolkits target only one organization, and include the related resources (e.g., HTML, JavaScript, CSS, image, and PHP files) with a limited page depth for replicating a portion of a targeted legitimate website; meanwhile, the links in the replicated webpages are often unchanged and still point to the targeted website, thus the phishing website may easily lose the control of visitors and fail to collect their login credentials. In contrast, our toolkit can replicate many targeted organizations by automatically constructing unlimited levels of phishing webpages in real time based on user interactions; meanwhile, all the links in the replicated webpages are modified to keep holding visitors on the corresponding phishing website and maximize the chances of collecting their login credentials. In addition, Cova et al. [5] did not report the existence of Web Single Sign-On (SSO) phishing [37] in those toolkits; while our toolkit supports Web SSO phishing as well as the traditional phishing.

Existing phishing susceptibility studies [7], [8], [9], [14], [15], [25], [26] often use some specific, not very realistic, and non-sharable testbeds with a limited webpage depth. For example, in [7], participants were informed of the real purpose of the study (i.e., identifying spoofed websites) in advance; in [8], participants were given a test account to role play; in [9], two specific domains (ebay-login.net and amazonaccounts.net) were registered to spoof Amazon and eBay; in [15], credentials of university students were the targets of a spear phishing test; in [25], one single bank website was used to evaluate the effectiveness of security indicators; in [26], a role-play survey was answered by participants recruited through Amazon's Mechanical Turk. In [14], Jackson et al. used a reverse proxy as the phishing website to intervene between the participants' computer and the legitimate websites; their testbed was designed to study the effectiveness of the extended validation certificate mechanism and the picture-in-picture phishing attacks; their participants were informed of the real purpose of the study in advance similar to [7]. However, our toolkit can be used by researchers to easily construct testbeds for performing various phishing related user studies. The constructed testbeds will be comprehensive and realistic because they support both the traditional phishing and the newly emergent Web SSO phishing, support all the popular browsers and allow participants to use their real login credentials to perform real browsing activities; meanwhile, they will not expose participants to any anticipated risk if properly configured (Section V-A).

III. Extreme Phishing and Our Goal

As introduced in Section I, the success of phishing depends on two layers of contexts [37]. The *first-layer context*, i.e., an email or instant message, is critical to entice users to click on a phishing URL [8], [15], [26], and the *second-layer context*, i.e., a phishing website itself, is critical to lure users to submit their login credentials [7], [8], [12], [14], [26].

Focusing on the second-layer context, we classify phishing attacks into three levels as shown in Figure 1, *simple phishing*, *advanced phishing*, and *extreme phishing*, based on the extent to which their look and feel are similar to their targeted legitimate websites. Intuitively, the more a phishing website is similar to the targeted legitimate website, the more likely it will be effective; researchers indeed found that, users often (1) identify phishing websites based on their suspicious look and feel [7], [12], [14], [26], (2) do not understand security indicators [7], [6], [8], [9], [26], and (3) disregard the absence of security indicators [33].

A. Metrics for Look and Feel

We define the ***look and feel*** of a phishing website using four metrics: its appearance, page depth, support to dynamic user interaction, and phishing types. The three levels of phishing attacks differ in their look and feel based on these four metrics.

The *appearance* including page layouts, text contents, images, and styles of a phishing website gives visitors the first impression. Phishing webpages with low visual similarity to the corresponding legitimate webpages could be easily detected as fake by users [7]. The appearance of *simple phishing* websites is only somewhat similar to that of corresponding legitimate websites, the appearance of *advanced phishing* websites is mostly similar to that of corresponding legitimate websites, and the appearance of *extreme phishing* websites is similar in every way to that of corresponding legitimate websites.

The *page depth* of a phishing website is the levels of webpages that are organized and linked together on the phishing website. Users often visit several linked pages on a website. Phishing webpages with missing or invalid links can potentially reduce the trust from visitors and fail to lure them to submit login credentials, while phishing webpages with valid but unmodified links (i.e., linking to the targeted or other legitimate websites) will lose the control of visitors and fail to attack them. The page depth of *simple phishing* websites is one and the links on the webpage are partially modified, the page depth of *advanced phishing* websites is limited to a certain number and the links on the webpages are partially modified, and the page depth of *extreme phishing* websites is unlimited and the links on the webpages are completely modified to gain the maximum control of visitors.

The *support to dynamic user interaction* of a phishing website means that user interactions such as clicking, searching, and form submission as well as the triggered JavaScript executions such as dynamic URL or other DOM element creation

	Metrics				Technical Challenges in Constructing the Attacks	
	Appearance	**Page Depth**	**Support to Dynamic User Interaction**	**Phishing Types**		
Extreme Phishing	Similar in every way	Unlimited levels of pages with completely modified links	Yes	Traditional & High-quality SSO	Dynamic generation of unlimited levels of pages based on user interactions, accurate link replacement, and sophisticated SSO page construction	**Complexity, Effort, and Effectiveness**
Advanced Phishing	Mostly similar	Limited levels of pages with partially modified links	No	Traditional & Low-quality SSO	Limited levels of page copy, and naive SSO page construction	
Simple Phishing	Somewhat similar	One page with partially modified links	No	Traditional	Single page copy	

Fig. 1. The classification of phishing attacks based on the second-layer context

can all be supported by the phishing website. A phishing website with better support to dynamic user interaction can potentially deceive visitors in a more effective manner. The support to dynamic user interaction is often missing in *simple phishing* and *advanced phishing* websites, while it is provided in *extreme phishing* websites.

The *phishing types* of a phishing website include traditional phishing and Web Single Sign-On (SSO) phishing. Traditional phishing aims to steal visitors' accounts that are created specifically for a website such as a shopping or banking website; Web SSO phishing aims to steal visitors' identity provider accounts such as Gmail, Facebook, and Yahoo accounts, each of which can allow a user to log into multiple relying party websites (Section IV-C). The *simple phishing* websites only support traditional phishing, the *advanced phishing* websites can support both traditional phishing and low-quality Web SSO phishing, and the *extreme phishing* websites can support both traditional phishing and high-quality Web SSO phishing.

B. Existing Phishing Websites

With a careful measurement and inspection of 471 live phishing websites reported on PhishTank [42] in 2015, we found that the majority of existing phishing websites are at the level of simple phishing because they have the corresponding properties of all the four metrics, only a handful of existing phishing websites are at the level of advanced phishing because they have the corresponding properties of some of those four metrics, and none of the existing phishing websites is at the level of extreme phishing because none of them has the corresponding properties of any of those four metrics.

Among those 471 phishing websites, 449 (95%) of them only contain a single phishing webpage which does not link to any other webpage on the same site. Meanwhile, among the landing pages of those 471 phishing websites, 30% of them do not contain any link, 22% of them contain invalid links that do not respond to users' click actions, 17.6% of them contain links to the targeted legitimate websites, and 26.4% of them contain links to other websites. By further manually examining 100 (out of 471) randomly selected phishing websites, we found that 69 of them are only somewhat similar to their targeted legitimate websites, only support traditional phishing, and do not support dynamic user interaction; two Yahoo, eleven Paypal, and three Gmail phishing websites are mostly similar to their corresponding legitimate websites in terms of the appearance; two Paypal phishing websites contain over two levels of webpages; ten phishing websites support low-quality Web SSO phishing.

C. Our Goal

The technical challenges in constructing those three levels of phishing attacks are different. For simple phishing, attackers only need to copy a single login webpage; for advanced phishing, attackers need to copy and link several webpages, and construct low-quality spoofed login "windows" if they want to perform Web SSO phishing. The webpages in these two levels of phishing attacks can be statically constructed and then deployed to a phishing website. For extreme phishing, attackers need to dynamically generate unlimited levels of webpages based on user interactions, accurately replace links on the generated webpages, and properly construct high-quality spoofed login "windows" if they want to perform Web SSO phishing; however, it is very challenging to meet these requirements because essentially the HTML elements, Cascading Style Sheets (CSS), and JavaScript on the legitimate websites must be accurately replicated to phishing websites and then rendered or executed on users' browsers in real time.

Correspondingly, the overall complexity and effort in constructing those three levels of phishing attacks also increase from simple to advanced and finally to extreme phishing. These factors can, to certain extent, explain why the majority of existing phishing websites are at the level of simple phishing and only a handful of existing phishing websites are at the level of advanced phishing. However, the success rate of existing phishing attacks in terms of the second-layer context is about 10% as reported in previous measurement studies [10], [16].

Therefore, our goal in this paper is to explore the feasibility of creating extreme phishing attacks that have the almost identical look and feel as those of the targeted legitimate websites, and evaluate the effectiveness of such extreme phishing attacks by performing a user study. It is reasonable to assume that if extreme phishing attacks can be more effective (i.e., insidious) than existing phishing attacks and can be easily constructed by using some toolkits, they will be widely adopted and deployed by phishers in the future.

IV. Design and Implementation

We now present our design and implementation of a toolkit for extreme phishing. This toolkit has the properties of all the four metrics of extreme phishing illustrated in Figure 1.

A. Overview

A toolkit for extreme phishing needs to automatically construct unlimited levels of phishing webpages in real time

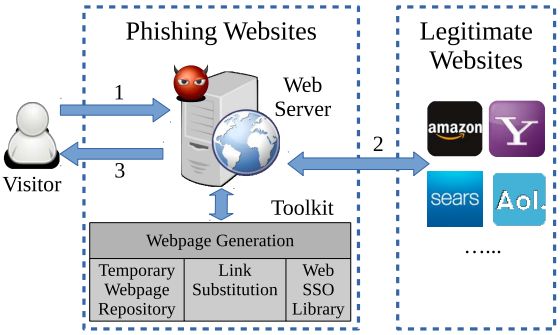

Fig. 2. High level design of the toolkit

based on user interactions. Meanwhile, in general phishers cannot deploy any toolkit on either a user's computer or a legitimate website because they do not have such capabilities in the threat model for phishing attacks. Therefore, it is very appropriate for us to take a proxy-based approach to design a toolkit for extreme phishing.

Figure 2 illustrates the high level design of the toolkit. It is deployed on a phishing website, works together with the phishing Web server, and acts as a proxy between a visitor and legitimate websites. It consists of four components: *webpage generation*, *temporary webpage repository*, *link substitution*, and *Web SSO library*. The webpage generation component intercepts the incoming/outgoing requests/responses to/from the phishing Web server for creating and delivering phishing webpages. The temporary webpage repository component temporarily saves the initially created and the finally modified phishing webpages. The link substitution component locally performs static link substitution and further enables dynamic link substitution on a user's browser by injecting JavaScript. The Web SSO library component constructs spoofed login "windows" based on different Web browsers and Web SSO identity providers. Only with a little configuration and customization effort, phishers can use this toolkit to construct and perform extreme phishing attacks. It is worth mentioning that this toolkit can support the replication of multiple targeted legitimate websites at the same time.

The workflow of the deployed extreme phishing attacks is as follows. Once a visitor clicks on a link to visit a phishing webpage, the corresponding request is sent to the phishing Web server. The webpage generation component of the toolkit intercepts the request, constructs a corresponding request to the legitimate website, saves the retrieved legitimate webpage as a file to the temporary webpage repository, and lets the Web server return the correspondingly constructed phishing webpage to the visitor's browser. Note that our toolkit serves phishing webpages via *http* (not *https*) and does not self-sign or forge SSL certificates [13], thus both avoiding triggering SSL warnings in visitors' browsers and reducing the effort for constructing the attacks; in other words, our proxy-based approach is more similar to SSL stripping than to SSL man-in-the-middle. Using files in the temporary webpage repository is an easy and reliable way to enable the inter-process communication between an independent toolkit and an unmodified Web server for them to complete the process. Once a visitor submits either a traditional or a Web SSO login form, the extreme phishing website ***redirects the visitor's browser to the corresponding legitimate website*** and does not need to further keep controlling the visitor.

B. Link Substitution

To keep holding visitors on a phishing website and maximize the chances of collecting their login credentials, our toolkit needs to ensure that all the links (including the hypertext references for elements such as anchors, buttons, and images) on each phishing webpage will be modified to point to the phishing website. There are two types of links: static links and dynamic links. Static links are contained in a webpage and they do not change after the webpage is delivered to a browser, while dynamic links are created or modified by JavaScript after the webpage is delivered to a browser. Accurate substitution of links especially dynamic links is critical to the success of extreme phishing, but it is also very challenging due to the dynamics of JavaScript. To address such challenges, our toolkit first performs accurate static link substitution on a phishing website, and then injects cleverly crafted JavaScript for performing accurate dynamic link substitution on a visitor's browser.

1) Static Link Substitution: The substitution for static links is relatively straightforward, but some details should be carefully considered. For each *link* element on a given webpage, if its "href" attribute uses an absolute path, (1) the legitimate domain name will be changed to the phishing domain name, and (2) the HTTPS scheme will be changed to the HTTP scheme. Later, once a visitor clicks on any link to a phishing website, a backward domain name replacement (i.e., from phishing to legitimate) will be performed by the toolkit to construct the corresponding request to the legitimate website. Such forward and backward domain name replacement operations will continue while a visitor is still browsing the phishing website.

This substitution cannot replace those static links that are contained in special elements such as <*head*> and <*script*>, for which special substitution operations will be performed by the toolkit. For instance, Yahoo extensively uses the *window.location.replace* method to perform the redirection on a JavaScript-enabled browser, and uses the *http-equiv="refresh"* attribute in <*meta*> to perform the redirection on a JavaScript-disabled browser. Our toolkit uses a set of customizable rules to replace the corresponding URLs in these and other special cases.

2) Dynamic Link Substitution: Unlike static links, dynamic links are created or modified by JavaScript after a webpage is delivered to a browser. Our toolkit injects cleverly crafted JavaScript for performing accurate dynamic link substitution on a visitor's browser. The injected JavaScript intercepts the dynamic link generation and modification events, and changes the legitimate domain name to the phishing domain name for each link.

The rendering of a webpage consists of two phases: *loading* and *event driven*. In these two phases, although many types of events can trigger the generation of dynamic links, eventually DOM (Document Object Model) insertion and modification events (i.e., *DOMNodeInserted* and *DOMAttrModified*) are directly related to the dynamic link generation. Correspondingly, our injected JavaScript code focuses on listening to these two types of events for dynamic link substitution. However, these two types of events often occur frequently in the webpage loading phase, thus intercepting and processing them in this

phase will incur high performance overhead to the browser. Therefore, in our design, the injected JavaScript code first intercepts the *Load* event for the window object (which indicates the completion of the loading phase) and replaces all the links that are dynamically generated in the loading phase in batch, and then intercepts the *DOMNodeInserted* and *DOMAttrModified* events only in the event driven phase. To ultimately ensure that the legitimate domain names in links are replaced by the phishing domain names, out toolkit further injects JavaScript code to monitor the clicked links on a visitor's browser and modify them if needed.

C. Web SSO Login Window Generation

Web Single Sign-On (SSO) systems allow users to log into multiple relying party (RP) websites such as foxnews.com and sears.com using one single identity provider (IdP) account such as a Google, Facebook, or Yahoo account, thus relieving users from the huge burden of creating and remembering many online credentials. In recent years, open Web SSO standards such as the OpenID authentication framework and the OAuth authorization framework have been rapidly and widely adopted by IdPs and RPs [41], [44]. Major IdPs such as Facebook, Google, Microsoft, and Twitter have also re-purposed OAuth for user authentication [3].

Researchers have identified the logic and implementation vulnerabilities of many deployed Web SSO systems [3], [29], [31], and have also highlighted the serious threat and consequence of Web SSO phishing attacks [37]. In comparison with traditional phishing, Web SSO phishing is more profitable and insidious because the value of IdP accounts is highly concentrated, the attack surface area is highly enlarged, and the difficulty of phishing detection by either algorithms or users is highly increased [37].

Web SSO phishing was proposed and manually constructed for a specific RP website by Yue in [37]. Our toolkit extends [37] and achieves the ***automatic and dynamic construction as well as inclusion*** of Web SSO phishing login windows. A Web SSO phishing site contains at least a *base webpage* for displaying the spoofed Web SSO login windows. Our toolkit aims to make the look and feel of the spoofed Web SSO login windows as close as possible to those of the legitimate ones, and make the inclusion of the spoofed Web SSO login windows as easy as possible. On a legitimate RP website, an Web SSO login window is a real browser window with the HTTPS URL address of an IdP (e.g., Google) login webpage; the same-origin policy in Web browsers ensures that a user's login credential submitted on an Web SSO login window cannot be accessed by any RP website. Therefore, the base webpage on a Web SSO phishing site cannot use a real window with an IdP's URL address to steal a visitor's login credentials.

In our design, we use $<div>$ elements to create spoofed Web SSO login "windows" on the base webpage. Our toolkit provides a JavaScript library for populating each $<div>$ element with the corresponding content and style of a real Web SSO login window. This populated $<div>$ element emulates the address bar and buttons of a real browser window using images, emulates the identical HTTPS URL of an IdP and the corresponding security lock icon using images, and emulates the identical content and style of a real Web SSO login page; it also supports all the relevant actions for the spoofed login "window" (e.g., minimize, maximize, close, resize, and drag), the security lock icon (e.g., click for viewing the certificate), and the login form (e.g., submit the login credentials). The appearance of legitimate Web SSO login windows varies on different OSes, browsers, and IdPs. For example, the window icon, the security lock icon, and the certificate viewing interface are different on different browsers, while the window title, the URL address, and the login page content depend on the IdPs. All these differences are properly considered in our JavaScript library. To support Web SSO phishing, attackers only need to embed a few lines of HTML and JavaScript code into the base webpage of their phishing site. The look and feel of our spoofed "window" are identical to those of the real window; it is almost impossible for users to differentiate them as demonstrated in our user study (Section V).

D. Implementation

We implemented our toolkit in Perl and JavaScript. The toolkit runs on an Apache Web server as an external filter [40], and supports the five most popular Web browsers (i.e., Google Chrome, Firefox, Opera, Safari, and Internet Explorer) on visitors' computers. The toolkit allows attackers to easily construct and deploy phishing attacks against different legitimate websites even including very complex ones such as Amazon, Sears, Yahoo, and AOL. It processes requests and responses efficiently, and delivers phishing webpages to visitors' browsers in real time.

V. User Study

To evaluate the effectiveness of extreme phishing, we set up a testbed and conducted a user study with the IRB (Institutional Review Board) approval.

A. Testbed

We used our toolkit to construct a testbed with four extreme phishing websites Amazon, Yahoo, Sears, and AOL, hosted via *http* on a Web server. The legitimate Amazon website only supports traditional sign-on, the legitimate Yahoo website supports both traditional sign-on and Web SSO (using Google and Facebook accounts, from 2011 to 2014), the legitimate Sears website supports both traditional sign-on and Web SSO (using Google, Facebook, and Yahoo accounts), and the legitimate AOL website supports both traditional sign-on and Web SSO (using Google, Facebook, Yahoo, and Twitter accounts). The four phishing websites emulate the corresponding sign-on features of the four legitimate websites, respectively. We assigned domain names *www.amazon.jigdee.com*, *www.yahoo.ibancu.com*, *www.sears.leuxfo.com*, and *www.aol.keirtu.com* to the four phishing websites; this type of phishing domain name composition trick has been used in real phishing attacks as analyzed in [19]. At the client-side, we provided a computer for all the participants. On this computer, we modified the DNS entries in the *hosts* file to have the phishing domain names point to the IP address of our phishing Web server. We also installed the five most popular browsers and configured them to clear the history and cookies for each session. This testbed provides a *realistic* environment for our study because it allows participants to use their real login credentials to perform real browsing activities.

Meanwhile, this testbed with our configuration does not expose participants to any anticipated risk – when a participant submits any login form on our phishing websites, client-side JavaScript code will immediately redirect the participant's browser to the corresponding legitimate website, thus no login credential of any participant will be recorded by our testbed.

B. Participants

We recruited 94 adults (57 younger and 37 older, 62 female and 32 male) from our campus and the local community to participate in this study. The age range of younger participants is from 18 to 38 years, while the age range of older participants is from 50 to 88 years. All the participants ranged in education from 12 to 20 years. 29 participants are/were majoring in psychology, medical, or nursing related fields; 21 participants are/were majoring in education, business, communication, or art related fields; 20 participants are/were majoring in engineering related fields; 24 participants did not provide their major information. All the participants were prescreened for the study eligibility that excludes who had a brain injury or concussion in the last three months, had been diagnosed with a mental disorder that may disrupt cognition, were currently taking any medications that interfered with thinking ability, or had been diagnosed with any cognitive or neurological disorder. Participants who do not routinely use Internet (at least once per month) or do not have enough accounts for this study were also excluded. Older adults received $15 compensation for their participation, and younger adults received either psychology course credits or $15 compensation for their participation.

C. Procedure

At the beginning of the procedure, participants were administered the informed consent, in which they were told that the study examines computer usage patterns in younger and older adults. No mention of phishing was included in the study's introduction. This deception was used so that any observed participant behavior on the websites could be attributed to aspects of the sites themselves and the participants' ability to evaluate them as they typically would. We also provided handout instructions to the participants on using Web SSO, and encouraged them to ask questions if needed.

We then asked each participant to perform four tasks by browsing four different websites, according to his or her list of personal accounts. Of the four websites presented, two were extreme phishing websites, one through traditional sign-on and one through Web SSO; the other two were legitimate websites, similarly, one through traditional sign-on and one through Web SSO. The links to ***the homepages of the four websites*** are provided on a task webpage; they were customized and their sequence was randomized for each participant. All the participants were allowed to use any of their favorite browsers and to leave the study at any time without penalty if they felt uncomfortable.

In each task, a participant was provided with an instruction to browse the corresponding website as he or she usually does (e.g., click on links and submit forms) for a few minutes, log into it using the specified traditional or Web SSO sign-on method sometime during the browsing, and finally sign out. The experimenters left the room before a participant started to perform the set of Web browsing tasks.

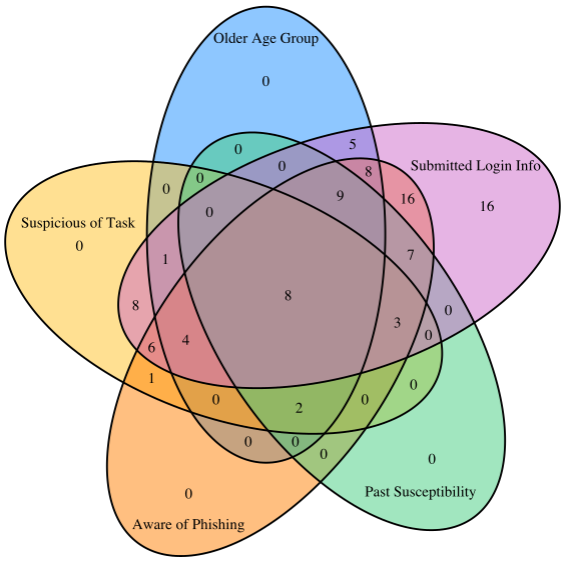

Fig. 3. The Venn diagram for the main questionnaire and observation results

The Internet browsing portion of the study was followed by a questionnaire about Internet safety, security practices, and aspects of the websites just visited. All the participants were debriefed regarding the true nature of the study and their questions were answered before they left.

D. Data Collection

We collected data through behavioral observation and questionnaire. The observation of Web browsing tasks occurred in a separate room, where a second computer was linked to the computer used by a participant via a screen sharing program, TeamViewer. For each participant on each website, experimenters observed and noted the website name, sign-on type (traditional or Web SSO), and user interactions. Experimenters also noted whether a participant exhibited any of a series of behaviors that would indicate his or her suspicion about a website, for example, clicking on the security lock icon, typing the website URL in another tab, searching the website URL or name via a search engine, or refusing to log into the website. After the participant completed the Web browsing tasks, a questionnaire was administered. This questionnaire included questions specific to the participant's behaviors and experience on the Web browsing tasks as well as his or her attitudes, beliefs, and practices regarding Internet security and phishing in general.

E. Results

We present the user study results in terms of the main questionnaire results, the observed & questionnaire results correlation, the Web SSO related questionnaire results, and other questionnaire results. We discuss these results and the Venn diagram for the main questionnaire and observation results (Figure 3) in the following subsections. We report the overall results for all the participants and compare the results between younger and older participants.

1) Main Questionnaire Results: One key question in the questionnaire is "Did you notice anything suspicious about the websites that you visited during this study? Please explain your answer to this question from the following perspectives: (1) Website look and feel; (2) Website URL address; (3) Website

identity certificate; (4) Other (please specify)." As shown in the Venn diagram with the label ***Suspicious of Task***, only 33 (35.1%) participants reported noticing something suspicious about the visited websites. Note that they visited both extreme phishing and legitimate websites; based on their explanations, only 7 participants who reported noticing something suspicious were actually suspicious of the phishing websites due to the look and feel of website (5 participants) or the website URL address (2 participants); 11 other participants were actually suspicious of the legitimate Web SSO websites they had never visited before (e.g., Slickdeals), while the rest did not provide relevant explanation for their suspicion. Therefore, the questionnaire results indicate that only **7 (7.4%)** of the 94 participants were actually suspicious about the extreme phishing websites they visited. There was no significant difference in identification of suspicious websites between older and younger adults ($\chi^2(df = 1) = 0.45, p = 0.5$; Pearson's Chi-squared test with Yates' continuity correction).

Among the 94 participants, 67 of them explained their answers from the perspective of website look and feel, but 35 of them explicitly commented that the websites were what they expected; 50 of them explained their answers from the perspective of website URL address, but 30 of them explicitly commented that they did not look at, understand, or feel unfamiliar with the website URL addresses; 35 of them explained their answers from the perspective of website identity certificate, but 24 of them explicitly commented that they did not look at, understand, or feel unfamiliar with the website identity certificates. So we can see that most people indeed rely on the look and feel of websites to judge their trustworthiness; they also consider the website URL address and identity certificate, but do not really understand or pay attention to them.

One ***Aware of Phishing*** related question is "Are you aware of phishing attacks?" As shown in the Venn diagram, 64 (68.1%) of the 94 participants reported awareness of phishing attacks. Older adults (16.2% reporting lack of awareness) ***were*** relatively more aware of phishing attacks than younger adults (42.1% reporting lack of awareness; $\chi^2(df = 1) = 5.78, p = 0.02$).

One ***Past Susceptibility*** related question is "Have you been susceptible to any phishing attacks in the past?" As shown in Figure 3, 29 (30.9%) of the 94 participants reported having been previously susceptible to phishing. Older adults (51.4% reporting prior susceptibility) ***were*** relatively more likely to have been a victim of phishing in the past than younger adults (17.5% reporting prior susceptibility; $\chi^2(df = 1) = 10.49, p = 0$). The reason for observing greater awareness of phishing in older adults may have been due to their higher rate of victimization relative to younger adults.

2) Observed & Questionnaire Results Correlation: One key behavior that the experimenters were observing is whether participants logged into the two (one traditional and one Web SSO) extreme phishing websites (labeled as ***Submitted Login Info*** in the Venn diagram). Only **3 (3.2%)** participants chose not to enter their username and password on either traditional or Web SSO phishing websites. This observed number is smaller than the number for *Suspicious of Task* obtained from the questionnaire. One reason why some participants would still log into a phishing website while noticing something suspicious may be related to authoritarian attitudes of participants as demonstrated in the well-known Milgram experiment [21], [22]; another reason may be related to trust delegation in lab environments as analyzed in a study performed by Sotirakopoulos et al. [28] to replicate an early study by Sunshine et al. [30] on SSL warning effectiveness.

Because participants were in a laboratory setting, they may have felt the need to obey the experimenters, even if they noticed something suspicious about the websites. We attempted to mitigate this influence by having the experimenters leave the room - thus allowing the participants to browse unsupervised - and by reminding participants that they could discontinue any browsing task at any time. However, it is possible that this tendency toward obedience was magnified because the browsing activities took place in a laboratory setting controlled by researchers. The participants may have trusted that such an environment was a safer place to enter login information than an unsecured environment.

Based on observations made by the experimenters, participants were most suspicious of the Slickdeals website, even though it was not used for phishing attempts. Therefore, another reason why participants submitted their login information to the extreme phishing websites is because the websites were trusted and familiar (e.g., Amazon), rather than unfamiliar (e.g., Slickdeals). This explanation aligns to certain extent with the study by Almuhimedi et al. on showing that users are more cautious on websites that they are not familiar with [2].

As shown in the Venn diagram, among the 64 participants who reported awareness of phishing attacks, 61 (95.3%) of them submitted their credentials to the extreme phishing websites; among the 29 participants who reported having been previously susceptible to phishing, 27 (93.1%) of them submitted their credentials to the extreme phishing websites.

There was no significant difference in this lack of susceptibility (i.e., *Submitted Login Info*) to the entire phishing testbed between those who did and did not report noticing something suspicious about the Web browsing tasks ($\chi^2(df = 1) = 1.49, p = .22$), between those with and without reported awareness of phishing ($\chi^2(df = 1) = 0.05, p = .82$), or between those with and without a past history of reported phishing susceptibility ($\chi^2(df = 1) = 1.83, p = .18$). There was also no significant difference in this lack of susceptibility to the entire phishing testbed between older and younger adults ($\chi^2(df = 1) = 1.06, p = .31$).

Only 5 (5%) participants chose not to enter their username and password when confronted with traditional phishing websites. There was no significant difference in this lack of susceptibility to the traditional phishing websites between those who did and did not report noticing something suspicious about the Web browsing tasks ($\chi^2(df = 1) = 0.57, p = .45$), between those with and without reported awareness of phishing ($\chi^2(df = 1) = 0.00, p = 1.00$), or between those with and without a past history of reported phishing susceptibility ($\chi^2(df = 1) = 0.00, p = 1.00$). There ***was*** a significant difference in this lack of susceptibility to the traditional phishing websites between older and younger adults ($\chi^2(df = 1) = 5.56, p < .05$), with 100% susceptibility in younger adults compared to 86% in older adults.

Only 10 (12%) participants chose not to enter their user-

name and password when confronted with Web SSO phishing websites. There was no significant difference in this lack of susceptibility to the Web SSO phishing websites between those who did and did not report noticing something suspicious about the Web browsing tasks ($\chi^2(df = 1) = 1.45, p = .23$), between those with and without reported awareness of phishing ($\chi^2(df = 1) = 0.24, p = .62$), or between those with and without a past history of reported phishing susceptibility ($\chi^2(df = 1) = 1.11, p = .29$). There was also no significant difference in this lack of susceptibility to the Web SSO phishing websites between older and younger adults ($\chi^2(df = 1) = 0.26, p = .61$).

3) Web SSO Related Questionnaire Results: One Web SSO related question is "Had you heard of *Web Single Sign-On* before coming in today?" Of the 94 participants, 31 (33%) had heard of Web SSO before the study. There was no significant difference in past exposure to Web SSO between older and younger adults ($\chi^2(df = 1) = 2.76, p = 0.1$). In responding to a question "Do you prefer to create a dedicated account for a website to sign into it or do you prefer to sign into the website using your Google, Facebook, or Yahoo account?", 19 (20.7%) participants reported a preference for Web SSO compared to traditional sign-on. There was no significant difference in this preference between older and younger adults ($\chi^2(df = 1) = 0.15, p = 0.7$). In responding to a five-point Likert-scale statement "It is likely that I will sign into websites using the Single Sign-On technique in the future.", 9 (10%) participants indicated that they strongly agreed, 20 (23%) indicated slight agreement, 16 (19%) felt neutral, 16 (19%) slightly disagreed, and 25 (29%) strongly disagreed. There was no significant difference in this expectation for future SSO use between older and younger adults ($\chi^2(df = 4) = 8.69, p = .07$).

In short, the majority (67%) of participants had never heard of the term of Web SSO, only a small percent (20.7%) of participants prefer to use Web SSO rather than traditional sign-on, and only 33% of participants agreed or strongly agreed that they may use Web SSO in the future. Therefore, although open Web SSO standards such as OpenID and OAuth have been rapidly and widely adopted by IdPs and RPs [41], [44] in recent years, users' understanding and acceptance of Web SSO are apparently lagging behind. To bridge this gap, researchers and communities really need to put more effort on educating users about the basic concept and the security practices of Web SSO systems.

4) Other Questionnaire Results: In responding to a five-point Likert-scale statement "Security is a concern when I perform Web browsing activities.", most participants indicated that computer security was important to them when browsing the Web, with 49 (52%) indicating strong agreement, 28 (30%) indicating slight agreement, 9 (10%) feeling neutral, 2 (2%) slightly disagreeing, and 6 (6%) strongly disagreeing. There was no significant difference in the importance placed on Web security between older and younger adults ($\chi^2(df = 4) = 2.94, p = .57$).

Qualitatively, the participants discussed four main themes when answering an open-ended question about their Web browsing security practices: "What security measures, if any, do you use to protect yourself from online identity fraud when using the Internet on your own computer?" These themes included *other personal behaviors* (younger adults, 43.9%; older adults, 37.8%), use of software for protection (younger adults, 42.1%; older adults, 48.7%), examination of a website's security features (younger adults, 19.3%; older adults, 16.2%), and password security practices (younger adults, 15.8%; older adults, 21.6%). Note that *other personal behaviors* include the usage of separate or proxy email addresses (e.g., spam accounts), blocking or avoiding suspicious emails or senders, using a Mac over a PC, avoiding sharing information such as credit card and personal information, and decentralization of accounts (e.g., using different passwords for different websites). Younger adults were more likely to report using one to four of these strategies, whereas older adults tended to report only using one or two of these strategies. Older adults were more likely to abstain from browsing the Internet or using a computer as a safety precaution.

5) Summary: The questionnaire results show that **87 (92.6%)** of the 94 participants were actually not suspicious about the extreme phishing websites that they visited, and the observation results show that **91 (96.8%)** of the 94 participants submitted their login credentials to the extreme phishing websites; meanwhile, most of those "victims" were aware of phishing before participating in this study or had been susceptible to some phishing attacks in the past.

Recall that in Section III-C, we reviewed that the success rate of existing phishing attacks in terms of the second-layer context is about 10% as reported in previous measurement studies [10], [16], thus ***existing phishing attacks do not work sufficiently well***. In addition, we allowed participants to browse extreme phishing websites for minutes, while this type of realistic environment was not observed in existing phishing susceptibility studies that we reviewed in Section II. Therefore, overall, we conclude that extreme phishing attacks are indeed very effective, i.e., highly insidious.

Note that it is not really possible to replicate the exact setup of those previous studies [10], [16] to have a direct comparison between the extreme phishing and existing simple phishing attacks. Also note that an extreme phishing website can use any of its webpages as the landing webpage and does not further control the visitor once the login form is submitted, while a simple phishing website often uses a single login webpage. Therefore, it is ***not really possible to design a new study to directly and fairly compare extreme phishing with simple phishing attacks*** because there will be no difference between them if a login webpage is used as the landing webpage for an extreme phishing website. This is also the main reason why we only measured the effectiveness of the extreme phishing attacks in our study.

VI. Discussion

The extreme phishing attacks that we explored are highly insidious - they can effectively deceive visitors as demonstrated in Section V, and can also effectively weaken many existing phishing defense mechanisms especially heuristics-based detection solutions. In this section, we discuss such impacts and provide suggestions to researchers and users for them to better defend against the extreme phishing attacks.

To detect phishing attacks, researchers have proposed various blacklist-based, heuristics-based, and whitelist-based solutions [38]. Blacklist-based solutions can achieve near-zero

false positives [17], [27], but they do not protect against zero-day phishing attacks [34], [38] because blacklists are updated only periodically and their coverage is often incomplete [27]; moreover, they have been challenged by the "rock phish gang", that uses phishing toolkits to create a large number of unique phishing URLs [34], [35]. As a result, many heuristics-based solutions have been proposed to detect phishing attacks using machine learning techniques with features extracted from URLs [10], [17], [18], [23], [32] and visual or non-visual elements on webpages [4], [17], [20], [23], [32]. Heuristics-based solutions can be used at the client-side to perform phishing detection in real time, and also at the servers-side to detect and supply phishing URLs for serving blacklist-based solutions; they need to achieve low false positives in order to be really usable and useful [38]. Whitelist-based solutions [33], [36] have also been proposed to complement the blacklist-based and heuristics-based solutions. In addition, hashing-based solutions [11], [24] have been proposed to protect against (rather than detect) phishing attacks.

Extreme phishing attacks will directly affect the effectiveness of many existing heuristics-based solutions, will indirectly affect the effectiveness of the existing blacklist-based solutions, but may not affect the effectiveness of the existing whitelist-based and hashing-based solutions.

Any heuristics-based solution that only uses features extracted from visual or non-visual elements on webpages may fail to accurately detect extreme phishing attacks that serve webpages with identical look and feel as those of the legitimate webpages. For example, most solutions heavily rely on the content including text, forms, scripts, and links of a webpage to detect anomalies [4], [17], [23], [32], and some solutions also use images to detect anomalies [4], [20]. Unfortunately, extreme phishing webpages will not produce obvious anomalies to them. Any heuristics-based solution that uses features extracted from URLs may become either inaccurate or incorrect on the detection of our Web SSO phishing attacks. Phishers can simply host the base webpages for Web SSO phishing attacks on their own RP websites or some legitimate websites such as Web forums and blogs, while the spoofed Web SSO login "windows" do not correspond to real URL addresses; therefore, no suspicious URL will be exposed to heuristics-based solutions [10], [17], [18], [23], [32] for performing the detection.

While blacklist-based solutions are not directly affected by extreme phishing attacks, they will be indirectly affected if the construction of their blacklists relies on heuristics-based techniques or anti-phishing communities. For example, the phishing blacklists used in Google Chrome and Mozilla Firefox are constructed and periodically updated by Google's large-scale automatic phishing classification infrastructure [32], which heavily uses heuristics-based techniques. In addition, blacklists often include phishing URLs verified by anti-phishing communities such as PhishTank [42]; it is very difficult for regular users to identify extreme phishing attacks as demonstrated in Section V, and for them to further submit phishing URLs to communities in a timely manner.

So far, whitelist-based solutions [33], [36] and hashing-based solutions [11], [24] are more robust against extreme phishing because they mainly rely on domain names to perform form filling or password derivation operations. However, users may need to pay more attention to properly use those solutions (such as pressing special keys for triggering password protection [11], [24]), while without being tricked by the look and feel of extreme phishing in the first place.

We suggest that researchers should seriously consider extreme phishing in their heuristics-based phishing detection solutions. For one example, anomalies in webpages alone can no longer serve as an effective metric in phishing detection; instead, URL analysis and webpage analysis should be combined together. For another example, identifying the intention (i.e., the intended website) of a user becomes indispensable in detecting extreme phishing, and existing solutions such as [20], [33], [34] are some good examples. Furthermore, researchers should also explore Web SSO phishing detection techniques. For example, the intention of a click action (i.e., the intended Web SSO IdP) on the base webpage could be leveraged to detect if a corresponding real login window or a $<div>$ element for a spoofed login "window" is displayed. However, ***automatic detection of extreme phishing attacks will still not be easier than automatic detection of simple phishing attacks*** especially because many phishing websites are short-lived [5], [17] and may not even be crawled in the first place; in addition, intention-based solutions (such as [34]) are already very effective in detecting simple phishing, and the space for them to further improve on detecting extreme phishing is very limited.

We suggest that Web users should be trained to (1) be aware of extreme phishing, (2) pay more attention to the domain name of a URL displayed in the address bar rather than just the look and feel of webpages, and (3) differentiate the spoofed Web SSO login "windows" from real ones. For example, one technique for detecting a spoofed Web SSO login "window" is to maximize, drag, or resize it because a spoofed "window" can never reach out of the webpage content area. In addition, it could be helpful for users to use some tools such as browser extensions to obtain intuitive information about the domain name in real time, thus potentially making informed decisions.

VII. Conclusion

In this paper, we explored the extreme phishing attacks and investigated the techniques for constructing them. We designed and implemented a concrete toolkit that can be feasibly and easily used by attackers to construct and deploy such attacks. Our toolkit can support both the traditional phishing and the newly emergent Web Single Sign-On phishing, and can automatically construct unlimited levels of phishing webpages in real time based on user interactions. We designed and performed a user study with 94 participants and demonstrated that extreme phishing attacks constructed by our toolkit are indeed highly effective, i.e., insidious. Finally, we discussed the impacts of extreme phishing on existing phishing defense mechanisms and provided suggestions to researchers and users for them to better defend against such attacks. It is reasonable to assume that attackers will adopt and widely deploy extreme phishing attacks using some similar toolkits in the future. Therefore, we urge the research community to pay serious attention to extreme phishing attacks, and we call for a collective effort to effectively defend against such attacks.

ACKNOWLEDGMENT

We sincerely thank anonymous reviewers for their valuable comments and suggestions. This research was supported in part by the NSF grant CNS-1624149.

REFERENCES

[1] D. Akhawe and A. P. Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Proc. of the USENIX Security Symposium*, pages 257–272, 2013.

[2] H. Almuhimedi, A. P. Felt, R. W. Reeder, and S. Consolvo. Your reputation precedes you: History, reputation, and the chrome malware warning. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS)*, 2014.

[3] E. Y. Chen, Y. Pei, S. Chen, Y. Tian, R. Kotcher, and P. Tague. Oauth demystified for mobile application developers. In *Proc. of the ACM Conference on Computer and Communications Security (CCS)*, 2014.

[4] N. Chou, R. Ledesma, Y. Teraguchi, J. C. Mitchell, et al. Client-side defense against web-based identity theft. In *Proc. of the Annual Network and Distributed Security Symposium (NDSS)*, 2004.

[5] M. Cova, C. Kruegel, and G. Vigna. There is no free phish: An analysis of "free" and live phishing kits. In *Proc. of the USENIX Workshop on Offensive Technologies (WOOT)*, 2008.

[6] R. Dhamija and J. D. Tygar. The battle against phishing: Dynamic security skins. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS)*, pages 77–88, 2005.

[7] R. Dhamija, J. D. Tygar, and M. Hearst. Why phishing works. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 581–590, 2006.

[8] J. S. Downs, M. B. Holbrook, and L. F. Cranor. Decision strategies and susceptibility to phishing. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS)*, pages 79–90, 2006.

[9] S. Egelman, L. F. Cranor, and J. Hong. You've been warned: An empirical study of the effectiveness of web browser phishing warnings. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1065–1074, 2008.

[10] S. Garera, N. Provos, M. Chew, and A. D. Rubin. A framework for detection and measurement of phishing attacks. In *Proc. of the ACM Workshop on Recurring Malcode*, pages 1–8, 2007.

[11] J. A. Halderman, B. Waters, and E. W. Felten. A convenient method for securely managing passwords. In *Proc. of the International Conference on World Wide Web (WWW)*, pages 471–479, 2005.

[12] J. Hong. The state of phishing attacks. *Communications of the ACM*, 55(1):74–81, 2012.

[13] L. S. Huang, A. Rice, E. Ellingsen, and C. Jackson. Analyzing forged ssl certificates in the wild. In *Proc. of the IEEE Symposium on Security and Privacy*, pages 83–97, 2014.

[14] C. Jackson, D. R. Simon, D. S. Tan, and A. Barth. An evaluation of extended validation and picture-in-picture phishing attacks. In *Financial Cryptography and Data Security*, volume 4886, pages 281–293. 2007.

[15] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer. Social phishing. *Communications of the ACM*, 50(10):94–100, 2007.

[16] M. Jakobsson and J. Ratkiewicz. Designing Ethical Phishing Experiments: A Study of (ROT13) rOnl Query Features. In *Proc. of the International Conference on World Wide Web (WWW)*, 2006.

[17] C. Ludl, S. Mcallister, E. Kirda, and C. Kruegel. On the effectiveness of techniques to detect phishing sites. In *Proc. of the International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, pages 20–39, 2007.

[18] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker. Beyond blacklists: Learning to detect malicious web sites from suspicious urls. In *Proc. of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2009.

[19] D. K. McGrath and M. Gupta. Behind phishing: An examination of phisher modi operandi. In *Proc. of the Usenix Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, 2008.

[20] E. Medvet, E. Kirda, and C. Kruegel. Visual-similarity-based phishing detection. In *Proc. of the International Conference on Security and Privacy in Communication Netowrks (SecureComm)*, 2008.

[21] S. Milgram. Behavioral study of obedience. *Journal of Abnormal and Social Psychology*, 67(4):371–378, 1963.

[22] S. Milgram. *Obedience to Authority: An Experimental View*. Harper & Row, 1974.

[23] Y. Pan and X. Ding. Anomaly based web phishing page detection. In *Proc. of the Annual Computer Security Applications Conference (ACSAC)*, pages 381–392, 2006.

[24] B. Ross, C. Jackson, N. Miyake, D. Boneh, and J. C. Mitchell. Stronger password authentication using browser extensions. In *Proc. of the USENIX Security Symposium*, 2005.

[25] S. E. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The emperor's new security indicators: An evaluation of website authentication and the effect of role playing on usability studies. In *Proc. of the IEEE Symposium on Security and Privacy*, pages 51–65, 2007.

[26] S. Sheng, M. Holbrook, P. Kumaraguru, L. F. Cranor, and J. Downs. Who falls for phish?: A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, 2010.

[27] S. Sheng, B. Wardman, G. Warner, L. Cranor, J. Hong, and C. Zhang. An empirical analysis of phishing blacklists. In *Proc. of the Conference on Email and Anti-Spam (CEAS)*, 2009.

[28] A. Sotirakopoulos, K. Hawkey, and K. Beznosov. On the challenges in usable security lab studies: Lessons learned from replicating a study on ssl warnings. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS)*, 2011.

[29] S.-T. Sun and K. Beznosov. The devil is in the (implementation) details: An empirical analysis of oauth sso systems. In *Proc. of the ACM Conference on Computer and Communications Security (CCS)*, pages 378–390, 2012.

[30] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. F. Cranor. Crying wolf: An empirical study of ssl warning effectiveness. In *Proc. of the USENIX Security Symposium*, pages 399–416, 2009.

[31] R. Wang, S. Chen, and X. Wang. Signing me onto your accounts through facebook and google: A traffic-guided security study of commercially deployed single-sign-on web services. In *Proc. of the IEEE Symposium on Security and Privacy*, 2012.

[32] C. Whittaker, B. Ryner, and M. Nazif. Large-scale automatic classification of phishing pages. In *Proc. of the Annual Network and Distributed Security Symposium (NDSS)*, 2010.

[33] M. Wu, R. C. Miller, and G. Little. Web wallet: Preventing phishing attacks by revealing user intentions. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS)*, pages 102–113, 2006.

[34] G. Xiang, J. Hong, C. P. Rose, and L. Cranor. Cantina+: A feature-rich machine learning framework for detecting phishing web sites. *ACM Transactions on Information and System Security (TISSEC)*, 14(2):21:1–21:28, 2011.

[35] G. Xiang, B. A. Pendleton, J. Hong, and C. P. Rose. A hierarchical adaptive probabilistic approach for zero hour phish detection. In *Proc. of the European Symposium on Research in Computer Security (ESORICS)*, pages 268–285. 2010.

[36] C. Yue. Preventing the revealing of online passwords to inappropriate websites with logininspector. In *Proc. of the International Conference on Large Installation System Administration: Strategies, Tools, and Techniques (LISA)*, pages 67–82, 2012.

[37] C. Yue. The devil is phishing: Rethinking web single sign-on systems security. In *Proc. of the USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, pages 1–4, 2013.

[38] C. Yue and H. Wang. Bogusbiter: A transparent protection against phishing attacks. *ACM Transactions on Internet Technology (TOIT)*, 10(2):6, 2010.

[39] Anti-Phishing Working Group (APWG). http://www.antiphishing.org/.

[40] Apache External Filters. http://httpd.apache.org/docs/2.2/mod/mod_ext_filter.html.

[41] OAuth 2.0. http://oauth.net/about/.

[42] PhishTank. http://www.phishtank.com/.

[43] Symantec Internet Security Threat Report. http://www.symantec.com/security_response/publications/threatreport.jsp.

[44] What is OpenID? http://openid.net/get-an-openid/what-is-openid.