Measuring Identity Confusion with Uniform Resource Locators

Joshua Reynolds[†] Deepak Kumar[†] Zane Ma[†] Rohan Subramanian[†] Meishan Wu[†]
 Martin Shelton[‡] Joshua Mason[†] Emily Stark[‡] Michael Bailey[†]

[†]University of Illinois at Urbana-Champaign [‡]Google, Inc.

{joshuar3, dkumar11, zanema2, rcsbra2, meishan2, joshm, mdbaily}@illinois.edu

ABSTRACT

Uniform Resource Locators (URLs) unambiguously specify host identity on the web. URLs are syntactically complex, and although software can accurately parse identity from URLs, users are frequently exposed to URLs and expected to do the same. Unfortunately, incorrect assessment of identity from a URL can expose users to attacks, such as typosquatting and phishing. Our work studies how well users can correctly determine the host identity of real URLs from common services and obfuscated “look-alike” URLs. We observe that participants employ a wide range of URL parsing strategies, and can identify real URLs 93% of time. However, only 40% of obfuscated URLs were identified correctly. These mistakes highlighted several ways in which URLs were confusing to users and why their existing URL parsing strategies fall short. We conclude with future research directions for reliably conveying website identity to users.

Author Keywords

Usable Security; URL; Phishing; Server Identity; Authentication, URL Readability.

CCS Concepts

•**Security and privacy** → **Usability in security and privacy**; *Authentication; Human and societal aspects of security and privacy*; •**Human-centered computing** → *Web-based interaction*;

INTRODUCTION

Users encounter Uniform Resource Locators (URLs) in varied contexts, such as web browsing, text messages, emails, chat applications, and social media. As they decide whether to visit these URLs in a browser, best practices for establishing trust task users with manually parsing URLs in order to determine identity. In the context of URLs, the fully qualified domain name (FQDN) is the only reliable indicator of identity. Unfortunately, users do not always parse URLs accurately, creating opportunities for adversaries to exploit user confusion.

This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI’20, April 25–30, 2020, Honolulu, HI, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6708-0/20/04...\$15.00

DOI: <https://doi.org/10.1145/3313831.3376298>

Consider the URL `https://bank.com.acct.balanc.es`. Users may mistakenly believe the URL refers to the FQDN `bank.com` rather than the potentially malicious FQDN `balanc.es`. Although software can unambiguously identify the FQDN from URLs, users often make mistakes in parsing the FQDN from a URL. These errors are extensively exploited in social engineering attacks such as phishing and typosquatting [15, 42, 32, 26, 20].

These attacks raise a more fundamental question: Why do users incorrectly parse URLs? Prior work has investigated this problem in the context of specific threats (e.g., phishing, typosquatting). Lin et al. [32] studied how different kinds of adversarially crafted URLs trick users in the context of browsing a page; however, they concluded that only 32% of participants even look at the URL bar. Recently, Thompson et al. identified a similar trend, noting that modifications to URL highlighting in the browser bar had no significant effect on helping users detect phishing attacks [43]. The ultimate metric that these studies evaluate, however, is phishing effectiveness, which confounds multiple factors along with URL confusion including website phishing content and browser security/identity indicators. Prior work has not studied how well users comprehend URLs in isolation. This removes confounding factors and enables the extraction of first principles for URL usability to guide URL redesign.

Our work is guided by two primary research questions. First, how accurately can users identify the FQDN from a URL in isolation? Second, what kinds of errors do users make when parsing URLs? To answer these questions, we built a large corpus of both legitimate and obfuscated URLs, informed by prior work and examples of malicious URLs found in the wild. Obfuscated URLs were potentially misleading. We then ran two exercises and one survey designed to assess user ability in parsing URLs and to distill their processes while doing so. We present results from 94 US Mechanical Turk participants.

Our first experiment investigated how well users parse both real and obfuscated URLs. We showed participants a mix of real and obfuscated URLs and found that they are able to correctly identify real URLs 93% of the time, but are misled 60% of the time when faced with obfuscated URLs. Participants’ ability to decipher URLs is demographically widespread and also independent of their security behavior and intentions [19], indicating URL comprehension is a task that challenges a diverse range of users.

We next evaluate how users decide on the identity of URLs. Participants performed an exercise where they were asked to highlight the key identifying portions of a URL. We find that participants differed in the parts of the URL that they highlighted, ranging from full URL highlighting to highlighting just the second-level domain. These findings hint at a misalignment between user mental models and the technical identity of URLs. Broadly, we identify three user models for understanding URL identity: identity tied to the FQDN, identity tied to an organization name, and identity tied to site function. Only 16% of participants consistently looked for the FQDN as the identity of a URL, suggesting that the majority of users extract higher level semantic meaning (e.g. organization, website function) from URLs, which can be easily spoofed by adversaries.

To supplement our quantitative view of user URL comprehension, we survey participants to understand their attitudes, experiences, and strategies when parsing URLs. Participants are generally aware of the risks associated with URL confusion, but—in direct contradiction of our quantitative results—96% of users believe that they are sufficiently protected by their own abilities and strategies for interpreting URLs. Users self-reported a wide range of strategies for determining whether a URL is safe, none of which were shared by more than 30% of participants.

Our results demonstrate that users frequently make mistakes in parsing URLs and offers insights as to why they do so. We conclude with a discussion of our results and outline some directions for future research. Ultimately, we hope this work highlights the issue of user URL misinterpretation and opens the door for future solutions.

RELATED WORK

Our work builds on research from a number of related areas, primarily in how users perceive malicious websites and work in improving security indicators.

URLs and Malicious Websites

Previous work has studied users' ability to identify fake or malicious websites in the context of browsing a page. Early work in this space focused primarily on phishing [16, 15, 32, 4, 17]. Downs et al. specifically studied how users understand URLs, with the conclusion that users who could grasp what URLs meant were less likely to fall victim to phishing attacks [17]. Lin et al. studied the impact of domain highlighting on identifying phishing websites, observing that although some users can benefit from highlighting, it does not have a sweeping effect on phishing detection [32]. Most recently, Thompson et al. launched a survey that investigated how changing the display of the URL in the context of browsing to a phishing site impacts the effectiveness of the attack [43], finding that the URL is often trumped by other, higher-level features of the website. In comparison to these studies, our study focuses specifically on URLs and users' ability to understand them broadly, outside of a phishing context and without the additional help of a browser or other UI elements.

In addition, extensive work has studied how URLs are used in the wild. Kim et al. found that popular URLs tend to be

between 40 and 80 characters long [28], which push the display boundaries of mobile phone browsers. URL shorteners, such as `bit.ly`, both obscure the identity of a URL and are widely used to succinctly reference pages in social media posts [13]. Prior work focusing on malicious URLs has quantified the prevalence and properties of typosquatting [42, 2], combosquatting [29], and other domain look-alike tricks [38].

Improving Security Indicators

To help users properly assess website identity, browser vendors and researchers have proposed and deployed a number of solutions in the realm of browser UI. For example, modern browsers now highlight the fully qualified domain name (FQDN) such that it stands out visually in the navigation bar [25, 33, 32, 36]. Outside of this, many papers have investigated the impact that security indicators have on impacting user behavior [10, 21, 39, 18, 3, 45]. Most recently, Thompson et al. [43] focused on investigating how EV browser UI is ineffective in helping users make adequate identity decisions, especially in the context of phishing.

Anti-Phishing User Education and Support

Phishing email classifiers such as those of Fette et al. [22], detect phishing emails using features unavailable or unknown to the average user. Althobaiti et al. have completed a review of such features used in phishing research [5]. Nevertheless, Erkkilä suggested that user's lack of knowledge was a solvable problem [20], and particularly suggested improving warning messages as a mechanism for user education. Stockhardt et al., Arachilage et al., and Kunz et al. evaluated interactive applications to teach phishing awareness to good success [41, 30, 7]. Volkamer et al. successfully built a short video [47] that empirically succeeded in educating users to detect phishing. Volkamer et al. also built a Thunderbird extension called TORPEDO to give users just-in-time advice about links in emails they are viewing [46]. Petelka et al. also designed a system to warn users in real time while interacting with phishing emails [35]. Althobaiti et al. built a URL-explanation tool to give anti-phishing, rather than generic advice about URLs [6]. However, it did not succeed in educating users to be able to achieve similar success without access to the tool. Our work measures the baseline difficulty users face in understanding a URL to be able to apply what they learn in anti-phishing trainings, and identifies reasons for this difficulty.

BACKGROUND AND TERMINOLOGY

URLs have several components, many of which have the potential to confuse users. A typical URL starts with the *scheme*, which on the web is typically either `http://` or `https://`. URLs next contain the Fully Qualified Domain Name (FQDN), which contains the domain name of the host or web service. FQDNs can themselves be complicated—for example, the FQDN `facebook.com.twitter.com.google.com` is technically a subdomain of `google.com`, in spite of the presence of several other high level “semantic” identities. Typically, the identity of a URL can be distilled to the FQDN, or more specifically, the effective TLD (eTLD) plus one child sublabel.

After the FQDN, URLs typically contain a *path*, which is used to route users to the appropriate resource on a web server. The

URL may also include *query parameters*, which are delimited by a ? in the URL followed, typically, by &-delimited key value pairs. Other possible components include an authentication string terminated by an @, a fragment identifier beginning with a #, and a specific remote port number. URLs also have a convention for escaping characters that may be freely applied to any delimiters or field elements.

METHODOLOGY

In order to learn how users parse and understand URLs, we had participants perform three activities which we describe below. The full instrument is available in the supplementary materials that accompany this paper. Participants were first asked demographic, background, and open-ended questions. They then participated in a URL highlighting activity followed by a URL target identification activity. The instrument concluded by administering the 16-item SeBIS scale [19].

Recruitment and Demographics

We recruited participants from Amazon Mechanical Turk. We recruited adult participants (18+), from the United States, who had a 95% previous task approval rating. Prior work has shown that these criteria provide reliable participants from Mechanical Turk in the context of security surveys and exercises [37, 34]. We offered \$3 in compensation for our activity, which took approximately 20 minutes to complete. Our study was approved by our institution's IRB.

In total, we recruited 120 participants for our study. In coding, we observed that 26 participant responses appeared to be automated or performed in bad faith, and had to be discarded. Our criterion for discarding a participant included: copy-pasting question text as a free response answer, duplicate, vacuous answers across multiple questions, and duplicate answers across multiple MTurk accounts.

Our final participant pool consisted of 94 participants. Participants were mostly male-identifying (62.77%), with a median age group of 26-35. One participant was 66 or older, and 13 (13.83%) of participants were age 18-25¹. Most participants (47, 50.00%) had a Bachelor's Degree, though 41 (43.62%) participants had a degree at the associate level or below. Only 6.38% of participants had a graduate degree. Finally, 86 participants (91.49%) said their primary browsing mechanism was mobile rather than desktop, indicating a heavy skew towards mobile-savvy users.

Background and Demographic Questions

We first asked participants demographic and background questions to gauge their security posture and experiences in understanding URLs. We collected age, gender, education level, and their primary technology platform (e.g., mobile, desktop). We then asked participants to self-report via Likert scales how competent they are at reading and understanding URLs they observe in day-to-day activities. Our Likert scales follow Van Deursen et al.'s suggestion [44] to phrase statements in terms

¹Age was captured as categorical data because we did not expect to find significant difference at a smaller granularity than general decade of age.

of "true of me", rather than traditional "agree/disagree" to attempt to elicit more factual answers.

Participants self-reported their confidence in determining identity from URLs in the following items: "I know how to read a URL." "I know how to tell which website I am on." "I know how to check where a link will take me before I click." "I know how to check where a link will take me before I touch it on a smartphone or tablet."

URL Target Identification

We next asked participants to perform a task we call URL target identification. This task is designed to measure participant efficacy at understanding the identity, or target, of URLs that they may encounter in normal web browsing. To this end we showed them real URLs gathered from pages of sites owned by organizations we expected would be familiar to our participants. We then showed them obfuscated URLs, misleading in several different ways. To avoid simply measuring effects tied to specific instances of obfuscated URLs, we pulled from a set of URLs representing each URL obfuscation technique. Possible learning effects of this ordering in our experimental design are discussed in our Limitations section. Participants attempted to learn the identity of the website indicated in the URL using only the URL text. Participants were shown 19-20 URLs because our timing pilot tests indicated this could be completed within our target 20 minutes for participants overall time commitment.

We created a seed set of 15 URLs by choosing three well-known URLs from the following categories: online commerce, email sites, news, banking, and social media.

Participants were also shown several examples of obfuscated URLs. To build a corpus of obfuscated URLs, we programmatically transformed our seed set of URLs using techniques collected from a combination of prior work on phishing [27, 15, 4, 17, 48, 24, 12, 28], browser defenses [25, 33, 36], thousands of phishing URLs reported on PhishTank [1], and general complexities of the URL [9, 23, 8]. Table 1 shows a full description of the transformations we performed in this study.

To learn how participants understand website identity as conveyed by a URL, we asked them to give the identity of the web site pointed to by the URL in a short free-response box. To try to gauge participant confidence in each answer, participants were instructed to indicate when they felt confident in their answer via a checkbox labeled "I'm certain." Specifically, participants were instructed, "For each of the following URLs, please give the identity of the web site it points to. Some of these URLs will be hard to read. Always make your best guess for each of these URLs, but only check the 'I'm certain' box when you feel confident in your answer. Do not visit these URLs!"

Analyzing Answers

Participants responded to each target identification question with a free-response answer. In order to extract semantic meaning from these responses, two independent coders coded each answer for both its correctness—meaning, if the user correctly identified the identity of the URL—and the category of answer.

URL Obfuscation Technique	Description	Example
Typo-squatting	A domain that looks similar but is spelled differently to one known by the victim.	https://twitter.com
Subdomain as Domain	Places an unrelated but familiar name as the subdomain for a URL	https://bofa.com.sign-in.info
IP Address	Includes Only an IP address	http://127.0.0.1/
IDN Homographs	Use unicode characters that look similar to the true website’s name	https://paypal.com
HTTP credentials as origin	Use HTTP AUTH credentials to precede the FQDN in a URL	https://fb.com@n593.biz
No Apparent Identity	URL contains only unrecognizable strings or a description of function	https://kjgksdsg93528.com
Self-declared secure	Recognizable hostname is prepended with “secure”	https://secure-gmail.com
Ambiguous Delimiter	Puts delimiters (e.g., @) in parts of the URL where they have no effect	https://bbc.com@cnn.com#@google.com
Unfamiliar TLD	Uses an unfamiliar TLD to terminate the FQDN instead of a more common TLD	https://twitter.com-issues.support
Overrunning Subdomain	Uses a long chain of subdomains to obscure the FQDN	https://www.facebook.com.js2awp-1lf8xe89770by5cyxbwewp.gvicw9vl45lie-csmcmcut7z95qcms.etz5811-eiue348wi0li27dh8jtkku.mx
URL Encoded Characters	Encodes characters in the URL to hide important delimiter characters.	https://fb.com%41%41%41%41%2e%41%52
Query parameters or Fragment Posing	Places familiar hostnames in the query or fragment portion of the URL	https://get-help.page?google.com https://192.17.42.13#google.com
Path Posing	Place familiar hostnames in the path portion of the URL	https://connection22.co/facebook.com

Table 1: **URL Identity Obfuscation**—URLs can be obfuscated in several ways, many of which can be confusing to users. These example obfuscations were collected from prior work, real phishing URLs, and our own observations. They guided our automated generation of obfuscated URLs. These given example URLs are not under our control; exercise appropriate caution. While we have included these links as plaintext, some PDF viewers will still convert them to clickable links.

We inductively and iteratively created a codebook to describe. These included a higher level business entity (e.g., Facebook), a site function (e.g., a login site), explicitly written FQDN or eTLD + 1 (e.g., www.google.com), a note that the URL was malicious (e.g., phishing), or if the participant was uncertain about their answer. Responses were coded as belonging to none, one, or multiple identity categories. The inter-rater agreement for this task using Kupper-Haffner agreement [31] was 0.799, indicating strong agreement between coders. Afterwards, coders reviewed each conflict together to resolve them.

URL Parsing Help

In order to help users parse URLs, browser vendors implement various techniques when displaying a URL [36] that are intended to assist user determination of the identity in a URL. The techniques applicable to the isolated URLs in our experimental setup include:

- **FQDN Highlighting** The FQDN of the URL is rendered darker than the surrounding characters to make it stand out.
- **URL Decoding** URL-encoded characters are decoded.
- **HTTP Authentication Eliding** All characters that are part of HTTP Authentication credentials are discarded.
- **Punycode for International Domain Names** Domain names with Unicode characters from multiple language sets are rendered unambiguously in Punycode.

In order to test the effect these techniques have on URL target identification, we applied them to the URLs provided to 40 participants (42.5% of our test population). The imbalance

is due to the removal of low-quality answers described in the following section.

URL Highlighting for Identity

Next, we wanted to learn where in the URL users decide to look when determining its identity. To do this, we designed a highlighting exercise, where participants were each provided four URLs and asked to highlight the parts of the URL that they thought were important to the identity of the URL.

All participants received the same four URLs. One was the legitimate login page of a payment processor. Another was chosen to be normal, but specifying HTTP rather than HTTPS as the protocol. Another was for a legitimate bank login page which used a possibly confusing (but not malicious) subdomain. The last was crafted to use as many of the transforms we identified as possible to see how users would react when their expectations of normal URLs were violated. Participants who were in the treatment group with URL parsing help received the same help in this question.

In addition to highlighting the URL, participants also provided some free-form text as to why that part of the URL is interesting or important to determining identity.

Open Ended Questions

Finally, we wanted to compare how users self-report their process with URLs with their observed behavior. Specifically, we asked participants the following question, “When you see a link or a URL, how do you decide if it is safe to go there?” Three independent coders then coded each response into nine iteratively and inductively generated categories, each

Figure 1: **Target Identification by User**—Users are able to correctly identify the identity of URLs a median 54.1% of the time. Four (4.3%) participants were correct 100% of the time, while 35.3% of participants were correct fewer than 10% of the time.

described in Table 5. Each response could belong to none, one, or several categories. After coding, coders reviewed each conflict and resolved them.

URL TARGET IDENTIFICATION

We begin our analysis by quantifying how well participants are able to correctly identify the target of URLs presented to them. In the context of this paper, we deem “correct” identification of a target to mean that the participant correctly identified either the IP address, fully qualified domain name (FQDN), the domain name (eTLD + 1), the organization/entity that controls the domain (e.g., Google), or correctly flagged a malicious URL as untrustworthy. For example, for the URL <https://www.paypal.com/us/signin> we would have accepted answers like “www.paypal.com”, “paypal.com”, or “PayPal”. In the obfuscated case such as <https://www.super-secure-paypal.com/us/signin>, we would have accepted answers like “www.super-secure-paypal.com”, “super-secure-paypal.com”, “super-secure-paypal”, or “not real paypal”. Due to an error in the URL generation process, we had to discard 81 of 1804 URLs (4.49%) that were syntactically invalid. After removing invalid URLs, each participant was presented between 14 and 19 valid URLs, with a median of 19 URLs per participant. In aggregate, participants were able to correctly identify the target of a URL in 1017 of 1723 cases (59%).

35% of URLs presented to participants were real URLs, and the remaining 65% were obfuscated URLs as described in Table 1. Participants were generally successful at identifying the targets of real URLs—575 (93%) of 618 targets from real URLs were correctly identified. Conversely, participants performed poorly on obfuscated URLs, only correctly identifying the targets of 442 (40%) of the 1105 obfuscated URLs presented.

Properties of the User

Participants varied in their effectiveness at labeling identity. Figure 1 shows a CDF of the fraction of correctly identified target per user. Participants were able to correctly identify a median 54.1% of targets, however, there are extremes at either end. On the low end, 10.6% of participants could correctly identify at most 33% of the targets presented to them. On the other end, four participants (4.3%) were able to correctly

identify the target for every URL provided, indicating strong proficiency in URL parsing ability.

To detail why some participants perform poorly on obfuscated URLs, we distilled the processes that participants used to identify URLs. Three expert coders manually coded each target response into one of three identity categories: whether the user specified a higher level business entity (e.g., Facebook), described a site function (e.g., a login site), or explicitly wrote an FQDN or eTLD + 1 (e.g., www.google.com). Responses were coded as belonging to none, one, or multiple identity categories.

For each user, we analyzed which URL category participants mainly choose when identifying targets. For example, user *A* may typically try to extract an organization or entity from a URL, while user *B* may expressly look for an FQDN. In most cases, we observe participants consistently prefer a single category for URL identification, identifying a median 82.3% of targets with a single top category. In the largest case, 69 (73.4%) participants primarily look for a specific organization or entity in a URL when identifying a target. 15 (16.0%) participants explicitly look for an FQDN or eTLD + 1, and 10 (10.6%) mainly highlight the website function.

Differences in preferred URL categorization also lead to differences in participants’ ability to correctly identify targets. Participants who explicitly look for an FQDN or eTLD + 1 were the strongest at identifying targets of URLs, with a median accuracy of 62.5%. This is in contrast to 55.5% accuracy for participants who look mainly for entities, and 36.8% for those that primarily look for website function. These results align with how URLs are technically designed. As the FQDN is the only technically accurate identity in the URL, it is not surprising that participants who think of URLs in this way are the most proficient at target identification.

Finally, we investigate if any property of the participant can serve as a predictor for identifying the target of provided URLs. Specifically, we look at the relationship between participants’ security behavior and their ability to correctly identify URL targets. As a proxy for security behavior, we had each participant complete the SeBIS Security Behavior Intentions Scale [19]. We correlate their SeBIS score with their rate of correct identification for their targets, using a standard Pearson correlation, and observe no statistically significant correlation between SeBIS score and identification rate ($r = 0.19$, $p = 0.056$). This result indicates that URL identification is a challenging task that impacts users regardless of their security behavior.

Additionally, correlations tested with gender, age, and the time taken to label each URL revealed no significant correlations.

Properties of the URL

Most URLs presented to the user were transformed using techniques found in modern phishing and URL based attacks. Each transform had varying levels of success in tricking users. Table 2 shows the effectiveness of each transform. Users were able to most correctly identify the target of a URL when no transform was applied to the URL (93%) and were relatively successful on other forms of well studied attacks, such as typosquatting (69.8%), and IDN homograph attacks (52.7%).

Transform	% Correct	% Specific Org	% FQDN	% Function
Unaltered	93.0	94.1	90.7	88.9
Typo-Squatting	69.8	68.6	91.3	42.9
IDN Homograph	52.7	56.4	55.6	28.6
No Organization Name in Domain	41.5	37.7	73.3	20.0
False Identity in Path	39.8	37.5	60.0	22.2
False Identity in Fragment Identifier	37.6	36.2	53.3	22.2
False Identity in Query String	37.2	34.8	60.0	20.0
Self-Declared Secure Domain	36.2	31.9	60.0	30.0
IP Address	35.0	38.5	25.0	33.3
Subdomain disguised by using uncommon TLD	34.5	37.5	35.7	11.1
False Identity in HTTP Credentials	32.1	36.8	25.0	14.3
Expected domain as subdomain	31.1	28.6	47.6	23.1
URL-Encoding-Disguised Expected Domain as Subdomain	29.1	33.3	22.2	14.3
Long Subdomain	25.8	28.9	18.2	16.7

Table 2: **Effectiveness of URL Transforms**—This table shows URL identity-parsing effectiveness in terms of overall correctness, as well as broken down by the subpopulations of participants giving each type of identity description. URL transforms had varying effects on user efficacy. 93% of unaltered URLs were correctly identified, compared to 25.8% of long subdomains in the worst case. Users are effective at identifying some kinds of transforms (typosquatting, IDN homographs), but are less effective when the transforms require deeper knowledge of the URL.

Conversely, target identification involving a long subdomain (25.8%), expecting a domain as a subdomain (31.1%), or a false identity in HTTP credentials (32.1%) were the least likely to be correctly identified.

Table 2 also shows the accuracy of identification across users in the three URL identification categories described above. In every case, users that primarily focus on a website function were the least effective. Between the other two classes of participants, some transforms disproportionately affected one or the other. For example, when the transform placed a familiar domain as the subdomain of the URL, only 28.6% of targets were correctly identified from participants who view URLs as an entity or organization versus 47.6% of targets from users who view URLs as an FQDN. In fact, placing a false identity anywhere in the URL was more effective at tricking “organization” participants than FQDN participants. Conversely, participants who primarily look for FQDNs were worse when no clear or recognizable FQDN was available—for example, with IP addresses or unrecognizable TLDs.

Do Browser Parsing Aids Help Participants?

Finally, we investigate how parsing aides deployed in browsers may be useful in helping users identify the targets of URLs. To test this, we performed a proportions t-test on the fraction of targets that were correctly identified across our control group and the group given aid. We find that 61.1% of targets were correctly identified in the group with URL parsing help, and 56.1% of targets were correctly identified in the group without any help. There is a statistically significant difference between the two groups, with $\alpha = 0.05$, ($z = 2.05$, $p = 0.039$), however, the effect size is small ($h = 0.1$). This lends some evidence to the fact that these tools do indeed help participants understand URLs. Even with these highlights, participants were only able to correctly identify 61.1% of targets, indicating there is much more work to be done.

Comparison with Phishing Studies

Erkkila claims that some phishing risk is due to users being unaware of URL syntax to extract domain names, and also identified lack of attention to this task [20]. Althobaiti et al. identify domain reading as one of the highest-used features in human-facing phishing avoidance systems [5]. This includes educating users to recognize what they call “deviated domains”. Thompson et al. find that for a similar highlighting exercise for a “subdomain as domain” transform, 85% of participants were effectively tricked [43], compared to only 38.9% of targets incorrectly identified in our study. Before presenting their training video, Volkamer et al. saw about 40% mistakenly trusted phishing emails and 25% false positive classifications on benign emails [47]. Considering URLs in isolation, our participants experienced far fewer false positives, but also were slightly more misled by our obfuscated URLs.

There may be some benefits to URL parsing that users derive from the context in which they are presented. In the context of website screenshots, Stockhardt et al. found about 60-70% mistakenly trusted sites with phishing URLs before being educated [41]. Kunz et al. also asked participants to evaluate website screenshots with phishing URLs and found users made errors corresponding to various types of obfuscations [30]. In the case of typosquatting and IDN-homographs, our participants performed similarly to pre-education participants in Kunz et al.’s study. However, our participants performed worse than pre-education participants in Kunz et al.’s study for obfuscated domain names and false organization names planted in the path, query parameters, or fragment identifier. Results were similar for pre-education participants in Althobaiti et al.’s study in the context of an instant messaging application [6].

USER URL HIGHLIGHTING

To shed light on how users understand URL identity, we performed an experiment where users show us their process by highlighting portions of several URLs that they deemed useful to determining identity. All 94 users observed the same first three URLs, which were benign URLs for Google

Figure 2: **Heatmap of URL Identifier Highlighting**—Aggregate URL highlights for 94 participants reveals a range of different user conceptions of URL identity.

URL	Type	# Users	# Correct	Highlighted URL portion							
				full	scheme:host	host/path	host	2ld.tld	subd.2ld	2ld	Other
#1	benign	94	83 (88.3%)	13	N/A	N/A	16	17	0	37	11
#2	benign	94	77 (81.9%)	10	2	4	4	29	7	21	17
#3	benign	94	79 (84.0%)	11	7	3	12	17	1	28	15
#4	deceptive	55	0 (0.00%)	3	0	N/A	0	0	0	0	52
#5	deceptive	30	7 (23.3%)	4	1	N/A	3	3	0	0	19

Table 3: **URL Identifier Highlighting**—Users determined the identity of benign URLs more accurately than obfuscated URLs, but highlighted different portions of the URL. The green and red cells represent our permissive definition of *correct* and *incorrect* identity highlighting.

(#1), Chase(#2), and Paypal(#3), respectively. As their fourth URL, 55 users saw URL #4 and 31 saw URL #5, which were the same deceptive URL, with the exception of #5 omitting username information and decoding URL-encoded characters. Nine users did not receive a fourth URL due to operational error. We allowed users to select multiple sections of a single URL, and represent the aggregated heatmap of URL highlighting frequency in Figure 2.

We first analyzed the accuracy of user highlighting: did they correctly highlight the scheme://host, host, host/path, subdomain.2LD, 2LD.TLD, or 2LD? Our criteria for correctness was intentionally permissive. For example, although second-level domains (2LDs) can be mimicked adversarially by using different top-level domains (TLDs), we did not test incorrect TLDs for known 2LDs; in other words, we interpret the selection of google and google.com to be equally correct since we did not test google.cm. Even with a generous definition of correctness, users ranged between 81.9–88% accuracy (Table 3) for the three benign URLs, which were not designed to deceive users in any way.

Users who highlighted the benign URLs displayed a wide range of variation in what part of the URL they highlighted. 21–37 users (22.3–39.3%) highlighted just the second level domain (2ld), i.e. the company name, for the first three URLs. This suggests that they might conceptualize URLs as the se-

mantically meaningful name that matches a real-world entity, and do not emphasize subdomains or TLDs. 29–33 users (30.9–35.1%) highlighted either the full host or 2ld.tld, which hints at an understanding of FQDN structure. And finally, 11–17 users (11.7–18.1%) highlighted incorrect URL portions.

To better understand erroneous highlighting selections, we also collected free response explanations of the highlighted URL sections. For URL #1, three of the eleven users who incorrectly highlighted the URL indicated their knowledge of Google in their explanations, but only highlighted truncated portions of the subdomain (www) and 2LD (google). Six of eleven users highlighted parts of just the scheme (http) when asked to highlight “each group of characters that helps you learn the identity of the website.” For URL #2, eight out of seventeen erroneous users focused on just the subdomain secure05b, with two declaring it secure and six suspicious of it (e.g. “This does not look right in the URL and is possibly a fake website.” or “This is not a company or brand name.”). Finally for URL #3, seven of fifteen mistaken users included the path us/signin along with an incomplete host, and six mentioned that it indicated a specific location.

For the two more complicated URLs, #4 and #5, users tended to highlight the beginning of the URL more frequently than the end of the URL, with the exception of the capitalized “SECURE-BANK-OF-AMERICA-SITE” portion at the end.

This result is likely a product of the left-to-right reading of Western languages, and might be reversed if tested with right-to-left languages.

USER CONFIDENCE, SELF-DESCRIBED STRATEGIES, AND RISK PERCEPTION

Having quantitatively measured participants' URL comprehension and behaviors, we next provide a qualitative analysis of participant attitudes and experiences towards URLs.

Participants were generally aware of the risks associated with URLs, with 67.02% of participants encountering a fake or malicious website at least monthly. In describing the potentials harms of such websites, participants cited phishing, malware, scams, fake news, and bothersome behavior like ads and web-pages that take up the entire screen. Only 17.02% of our participants said they had been harmed by fake websites.

User Confidence and Competency

In spite of participants' poor ability to identify the intended targets of URLs, participants were confident about their ability to parse URLs and identify where on the Internet they were (Table 4). 91 (96.81%) of participants claimed that it was "mostly true" or "very true" that they "know how to read a URL". Only three participants answered neutrally or negatively to the question. Similarly, all but 3 participants responded positively to the prompt "I know how to tell which website I am on." These results indicate a wide gap between observed participant behavior and their perspective on their effectiveness.

User Parsing Strategies

We next asked participants to describe the strategies they employ when parsing a URL. We asked them "When you see a link or a URL, how do you decide if it is safe to go there?" Their coded answers are summarized in Table 5. Broadly, their strategies fall into three large categories: checking for HTTPS, looking for red flags in the URL string, and using external signals or clues. While using external signals is not a URL-parsing strategy, it was a heuristic described by several of our users for outsourcing or bypassing the task.

18 responses coded as "Other" included non-specific statements (e.g. "I decide on a hunch." (P-596)) as well as actions that cannot be taken without actually visiting the URL (e.g. "Whether it is secured or not with a lock icon." (P-640))

Checking for HTTPS

Checking for the presence of HTTPS in the URL was the largest shared strategy participants used (28.7% of participants). Although the presence of HTTPS guarantees certain security properties (e.g., confidentiality, integrity), it does not establish *identity*—the site you are communicating with may still be malicious. Unfortunately, some participants seem to be using HTTPS as a proxy for trust. For example, P-555 felt HTTPS translates directly to safety:

I know it is safe when it reads https, the s stands for secure for me. (P-555)

Other participants indicated a slightly more nuanced view of HTTPS. P-582 and P-568 both note that HTTPS is necessary, but not sufficient for "complete security":

If it's secured, encrypted, or a https:// link. If it's all three, it's safe. If it has 2/3, it's potentially unsafe. (P-582)

The way I usually decide if a site is safe to visit is if it has "https://" at the beginning. This implies at least some measure of security. (P-568)

Finally, one participant (P-644) mentioned checking both for HTTPS and other information from the URL, indicating heightened security awareness:

I first think about if it is a place I know is a legit website. Then I'm looking for HTTPS cert and if the URL just look sensible. (P-644)

Notably, P-644 had a high (89.4%) accuracy when identifying the targets from URLs.

URL Text Parsing

In alignment with our observed models of participant behavior, where participants primarily extract high level organizations and the FQDN of a URL, a total of 49 participants (52.13%) reported using various clues from the URL itself to discern a link's safety.

27% of participants mentioned using familiarity as a proxy for trustworthiness. P-567 noted that URL parsing was an especially easy task:

I check the url for familiarity. It's quite frankly easy to tell if it's an official link to an authentic website. (P-567)

Indeed, P-567 was one of four users in our study with 100% accuracy in target identification.

13 participants look for when the URL text exhibits anomalies that differ from their expectations. P-598 and P-572 noted looking for misspellings and strange characters:

Check to see if it's misspelled[sic] or weird (P-598)

If it looks like crazy letters then I don't click it (P-572)

11 participants mentioned more technical features, including TLDs and link-shorteners.

Check the url see if it has any other characters that are trying to make you click it such as if they used an l were an i should be etc. Also check the prefix of the site and the domain of it. .com .org .ru things of that nature (P-534)

If it not shortened and is one that makes sense. Like if I'm opening company A and the URL is companyA.com/... I would click it. (P-547)

A subset of the users who described scanning the URL text for anomalies provided technical examples.

... or if it has a prefix and country I recognize with the official company name in the center, it doesn't have a long string of numbers or letters anywhere in the main url or tail. (P-630)

Look for discrepancies with the link to what it should be. For example if I want to go to Amazon.com I know that this should be the beginning of the link. Some phishing websites might have Amazonc13.com or something like

Answer	Can ID Cur- rent Site	Can Check Link Target (desktop)	Checks Link Target (desk- top)	Can Check Link Target (mobile)	Checks Link Target (mo- bile)
Very True	68 (72.34%)	58 (61.70%)	36 (38.30%)	32 (34.04%)	25 (26.60%)
Mostly True	23 (24.47%)	29 (30.85%)	37 (39.36%)	28 (29.79%)	25 (26.60%)
Neither True nor Untrue	3 (3.19%)	2 (2.13%)	13 (13.83%)	15 (15.96%)	14 (14.89%)
Not Very True	0 (0.00%)	5 (5.32%)	8 (8.51%)	11 (11.70%)	20 (21.28%)
Not at all True	0 (0.00%)	0 (0.00%)	0 (0.00%)	7 (7.45%)	10 (10.64%)
I don't understand...	0 (0.00%)	0 (0.00%)	0 (0.00%)	1 (1.06%)	0 (0.00%)

Table 4: **Self-Reported Identity Competency**—Participants self-reported whether they know how to identify the website they are on as well as check where links go. We asked them separately about the desktop and mobile platforms because the process differs with the interface.

Strategy	Total	Percent	(Org,Function,FQDN)
HTTPS	27	28.72%	(21,5,1)
Familiarity	25	26.60%	(14,5,6)
Check Link href	14	14.89%	(9,0,5)
Misspelling or Gibberish	13	13.83%	(9,1,3)
Avoid Shortened/Unfamiliar TLD	11	11.70%	(7,1,3)
Passive Tool	7	7.45%	(6,1,0)
Context Clues	6	6.38%	(6,0,0)
Active Tool	4	4.26%	(1,0,3)
Other	18	19.15%	(14,1,3)

Table 5: **URL Evaluation Strategies**—These are the categories of answers participants gave to the question, “When you see a link or a URL, how do you decide if it is safe to go there?” These codes are not mutually exclusive as participants sometimes described using multiple heuristics. For each code, we indicate how many participants were grouped into the three observed URL-parsing strategies from the URL target identification activity.

this. It usually will be close but you can tell if you look in my experience. (P-646)

External Tools and Context Clues

Eleven participants (11.70%) reported using external tools to aid their evaluation of URLs. Seven relied on software they have installed to warn them before they made any dangerous choices.

I have security software in place that warns me if it's unsafe. (P-594)

i have a antivirus scanner, so it will check whether the site is safe or unsafe. (P-592)

Four reported taking active steps to use an external tool for URL validation.

it will show 80 to 90percent in scamadvisor (P-614)

Fourteen participants used the built-in feature of their browser to preview the target of links. Unfortunately, because this feature renders within the render window of the browser, it is vulnerable to spoofing.

I hover my mouse over the link and the website address appears at the bottom of my screen. (P-634)

I hover over the link and see what the address is on my status bar. If it seems sketchy or is something I've never heard of, I won't click. (P-611)

Finally, six participants (6.4%) mentioned they expressly rely on the context in which they received a URL. For example, P-612 noted “sketchy emails” and “spam” as things to avoid:

I consider the context of how it was presented to me. Sketchy email? No thanks. Someone spams a shortened link on a forum advertising something that's too good to be true? No thanks. (P-612)

These user strategies are reasonable, but insufficient for the complex task of detecting obfuscated identity reliably in URLs. The general simplicity of users' URL parsing strategies may explain their common failures in the face of abnormal and obfuscated URLs.

LIMITATIONS

Our work contains several methodological limitations that limit the scope of the results. To begin, our study treats URL identity parsing as a primary activity, rather than in the context of typical user web browsing. Our results thus serve only as a lower bound of user performance at URL identification, as users will devote less attention to URL identification in practice than they did in this study.

Participant familiarity with URLs in our seed-set may bias our results. For example, a participant's familiarity with certain websites (e.g., `secure05b.chase.com`) may have influenced their performance over participants who were unfamiliar with the websites we selected.

Our methodology is subject to learning effects due to the order of exercises. Each participant was provided the same, fixed ordering of tasks—as such, it is possible that participants performed better on the final task if they improved their parsing ability from previous tasks. Also, presenting users first with real URLs and then with their obfuscated counterparts may have primed participants.

Finally, our sample size was relatively small and samples from a single population—Amazon Mechanical Turkers. This

population has been found by Redmiles et al. to reasonably estimate the U.S. population, particularly adults under 50 with some college education. [37] There may still be challenges in interpreting URLs against other populations. Examples include, non-U.S. users, users with non-alphabetical scripts or those that primarily read in a right-to-left language.

FUTURE WORK

Given our results and limitations, there are a number of areas of future research directions. We note that we did not vary our experiments based on the type of device used by participants. Participants may have more difficulty parsing identity of a website on a smaller screen (e.g., a mobile device) than on a desktop computer screen. A future experiment might test a similar metric while also measuring the effect of the underlying screen size per participant.

Another direction of future research is to investigate the wide array of URL parsing aids that exist. In our study, we tested only one of these aids—how Google Chrome (and other Chromium-based browsers) display the URL in the browser bar. Further experiments into different URL aids might inform building better defense mechanisms into currently deployed systems.

Finally, there is a design opportunity to build and test new identity display methods, in the vein of Dhamija et al.’s dynamic security skins [14]. Everything from alternative presentations of URL text to entirely new methods of displaying identity could be explored.

DISCUSSION

User Education

We believe user education may be able to help users be wary of URL obfuscations. Prior studies on user education against phishing have shown that user education can be effective in helping users avoid phishing sites, in context [7, 41, 30]. For example, Volkamer et al.’s phishing awareness video [47], Sheng et al.’s anti-phishing game [40], and the NoPhish android app by Canova et. al [11] all experienced success. There are indications that education can specifically help users parse URLs.

URL confusion stems from a fundamental misalignment between user URL-parsing strategies and technical URL complexity, and solutions to the problem can either educate users (directly or indirectly through UI design) or re-design URL identity to reduce technical complexity and make URLs more user-friendly. Our study revealed two obstacles to user education. First, the current basis of user URL comprehension is varied, since users have a wide range of strategies and models for determining the identity of URLs. These diverse starting points of URL understanding likely require diverse educational approaches. Second, users believe that they are already sufficiently capable of identifying URLs, and overcoming this false confidence will be necessary before effective education can take place.

Limiting Extraneous Information

In our experiments, we observed users looking for identity information in authentication credentials, query strings, and fragment identifiers. Users were relatively adept at noticing mistakes in identity names (i.e. typosquatting and homoglyphs), but struggled when familiar identity names were placed in other parts of the URL, where they could be trivially spoofed. Creating a single place where users can look and find identity information would eliminate distractions. One system following this guideline was tested in 2014, when Google Chrome’s “Origin Chip” experimental feature went as far as to show truncated URL strings in the address bar displaying only the “origin” of the site being visited—but was unfortunately discontinued after strong negative feedback from a set of vocal power users². Gradual feature roll-out and support for user opt-out will likely benefit future changes of this nature.

Limiting the Length of Identifiers

When examining confusing URL transforms, we found that users were least able to understand URLs with long subdomains/FQDNs. Similarly, we observed that users tended to more frequently highlight the beginning of long URLs, with declining frequency towards the end. Even for the real chase.com URL there were a handful of users who only highlighted the “secure05b” subdomain. Two plausible explanations are that users are conditioned to find identifiers at the beginning of URLs, or that users may experience “parsing fatigue” for long URLs. In either case, we recommend shortening URL identifiers, when possible. One especially intriguing technique that could accommodate existing URL usage would be to display the FQDN before other URL components, and invert the FQDN itself. For instance, instead of displaying `https://secure05b.chase.com`, users would see `https://com.chase.secure05b`.

CONCLUSION

In this paper, we studied how accurately users understand and parse URLs. We found that although 96% of participants report confidence in their ability to read and parse a URL, we found that only 40% of obfuscated URLs were identified correctly, indicating a misalignment between participant models and their observed behavior. We classified a wide range of participants’ URL parsing strategies, which we synthesized into three general user models. We concluded with recommendations that we hope will inspire innovation to improve users’ ability to determine identity on the web and form the foundation for future research.

ACKNOWLEDGEMENTS

We would like to acknowledge the contributions of our anonymous shepherds in guiding the presentation of this work, as well as our anonymous reviewers. This work was partially supported by the National Science Foundation (NSF) under grant CNS-1518741. Joshua Reynolds was partially supported by a State Farm Doctoral Fellowship. We would also like to thank Nathan Malkin.

²<https://bugs.chromium.org/p/chromium/issues/detail?id=368725>,
<https://bugs.chromium.org/p/chromium/issues/detail?id=331373>

REFERENCES

- [1] 2019. PhishTank. <https://www.phishtank.com>. (2019).
- [2] Pieter Agten, Wouter Joosen, Frank Piessens, and Nick Nikiforakis. 2015. Seven months' worth of mistakes: A longitudinal study of typosquatting abuse. In *22nd Network and Distributed System Security Symposium (NDSS 2015)*.
- [3] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *22nd USENIX Security Symposium*.
- [4] Mohamed Alsharnouby, Furkan Alaca, and Sonia Chiasson. 2015. Why phishing still works: User strategies for combating phishing attacks. *International Journal of Human-Computer Studies* 82 (2015).
- [5] Kholoud Althobaiti, Ghaidaa Rummani, and Kami Vaniea. 2019. A Review of Human-and Computer-Facing URL Phishing Features. In *IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*.
- [6] Kholoud Althobaiti, Kami Vaniea, and Serena Zheng. 2018. Faheem: Explaining URLs to people using a Slack bot. In *Symposium on Digital Behaviour Intervention for Cyber Security*.
- [7] Nalin Asanka Gamagedara Arachchilage, Steve Love, and Konstantin Beznosov. 2016. Phishing threat avoidance behaviour: An empirical investigation. *Computers in Human Behavior* (2016).
- [8] T. Berners-Lee. 1994. Universal Resource Identifiers in WWW. (1994).
- [9] T. Berners-Lee, L. Masinter, and M. McCahill. 1994. Uniform Resource Locators (URL). (1994).
- [10] Robert Biddle, Paul C Van Oorschot, Andrew S Patrick, Jennifer Sobey, and Tara Whalen. 2009. Browser interfaces and extended validation SSL certificates: an empirical study. In *ACM workshop on Cloud computing security*.
- [11] Gamze Canova, Melanie Volkamer, Clemens Bergmann, and Roland Borza. 2014. NoPhish: an anti-phishing education app. In *International Workshop on Security and Trust Management*. Springer.
- [12] Sidharth Chhabra, Anupama Aggarwal, Fabricio Benevenuto, and Ponnurangam Kumaraguru. 2011. Phi.sh/\$ocial: the phishing landscape through short urls. In *8th Annual Collaboration, Electronic messaging, Anti-Abuse and Spam Conference*.
- [13] Daejin Choi, Jinyoung Han, Selin Chun, Efstratios Rappos, Stephan Robert, and Ted Taekyoung Kwon. 2018. Bit.ly/practice: Uncovering content publishing and sharing through URL shortening services. *Telematics and Informatics* (2018).
- [14] Rachna Dhamija and J Doug Tygar. 2005. The battle against phishing: Dynamic security skins. In *1st Symposium on Usable Privacy and Security (SOUPS)*.
- [15] Rachna Dhamija, J Doug Tygar, and Marti Hearst. 2006. Why phishing works. In *24th ACM Conference on Human Factors in Computing Systems*.
- [16] Julie S Downs, Mandy Holbrook, and Lorrie Faith Cranor. 2007. Behavioral response to phishing risk. In *2nd Anti-phishing working group eCrime researchers summit*.
- [17] Julie S Downs, Mandy B Holbrook, and Lorrie Faith Cranor. 2006. Decision strategies and susceptibility to phishing. In *2nd Symposium on Usable Privacy and Security (SOUPS)*.
- [18] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *26th ACM Conference on Human Factors in Computing Systems*.
- [19] Serge Egelman and Eyal Peer. 2015. Scaling the security wall: Developing a security behavior intentions scale (SeBIS). In *33rd Annual ACM Conference on Human Factors in Computing Systems*.
- [20] J Erkkila. 2011. Why we fall for phishing. In *29th ACM Conference Human Factors in Computing Systems*. ACM.
- [21] Adrienne Porter Felt, Robert W Reeder, Alex Ainslie, Helen Harris, Max Walker, Christopher Thompson, Mustafa Emre Acer, Elisabeth Morant, and Sunny Consolvo. 2016. Rethinking connection security indicators. In *12th Symposium On Usable Privacy and Security (SOUPS)*.
- [22] Ian Fette, Norman Sadeh, and Anthony Tomasic. 2007. Learning to detect phishing emails. In *16th international conference on World Wide Web*. ACM.
- [23] R. Fielding, J Gettys, J. Mogul, H Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. 1999. Hypertext Transfer Protocol – HTTP/1.1. (1999).
- [24] Evgeniy Gabrilovich and Alex Gontmakher. 2002. The homograph attack. *Commun. ACM* 45 (2002).
- [25] Google. 2016. IDN in Google Chrome. <https://www.chromium.org/developers/design-documents/idn-in-google-chrome>. (2016).
- [26] Markus Jakobsson, Alex Tsow, Ankur Shah, Eli Blevis, and Youn-Kyung Lim. 2007. What instills trust? a qualitative study of phishing. In *11th Conference on Financial Cryptography and Data Security*.
- [27] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. "My Data Just Goes Everywhere:" User Mental Models of the Internet and Implications for Privacy and Security. In *11th Symposium On Usable Privacy and Security (SOUPS)*.
- [28] Sungjin Kim, Jinkook Kim, and Brent ByungHoon Kang. 2018. Malicious URL protection based on attackers' habitual behavioral analysis. *Computers & Security* 77 (2018).

- [29] Panagiotis Kintis, Najmeh Miramirkhani, Charles Lever, Yizheng Chen, Rosa Romero-Gómez, Nikolaos Pitropakis, Nick Nikiforakis, and Manos Antonakakis. 2017. Hiding in plain sight: A longitudinal study of combosquatting abuse. In *24th ACM SIGSAC Conference on Computer and Communications Security (CCS)*.
- [30] Alexandra Kunz, Melanie Volkamer, Simon Stockhardt, Sven Palberg, Tessa Lottermann, and Eric Piegert. 2016. Nophish: evaluation of a web application that teaches people being aware of phishing attacks. *Informatik 2016* (2016).
- [31] Lawrence L Kupper and Kerry B Hafner. 1989. *The assessment of interrater agreement for multiple attribute responses*. International Biometric Society.
- [32] Eric Lin, Saul Greenberg, Eileah Trotter, David Ma, and John Aycock. 2011. Does domain highlighting help people identify phishing sites?. In *29th ACM Conference on Human Factors in Computing Systems*.
- [33] Mozilla. 2017. IDN Display Algorithm. https://wiki.mozilla.org/IDN_Display_Algorithm. (2017).
- [34] Eyal Peer, Joachim Vosgerau, and Alessandro Acquisti. 2014. Reputation as a sufficient condition for data quality on Amazon Mechanical Turk. *Behavior research methods* 46 (2014).
- [35] Justin Petelka, Yixin Zou, and Florian Schaub. 2019. Put Your Warning Where Your Link Is: Improving and Evaluating Email Phishing Warnings. In *2019 ACM Conference on Human Factors in Computing Systems*.
- [36] The Chromium Project. 2019. Guidelines for URL Display. https://chromium.googlesource.com/chromium/src/+master/docs/security/url_display_guidelines/url_display_guidelines.md. (2019).
- [37] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2019. How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples. In *40th IEEE Symposium on Security and Privacy*.
- [38] Richard Roberts, Yaelle Goldschlag, Rachel Walter, Taejoong Chung, Alan Mislove, and Dave Levin. 2019. You Are Who You Appear to Be: A Longitudinal Study of Domain Impersonation in TLS Certificates. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2489–2504.
- [39] Stuart E Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. 2007. The emperor’s new security indicators. In *28th IEEE Symposium on Security and Privacy*.
- [40] Steve Sheng, Bryant Magnien, Ponnurangam Kumaraguru, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 2007. Anti-phishing phil: the design and evaluation of a game that teaches people not to fall for phish. In *Proceedings of the 3rd symposium on Usable privacy and security*. ACM, 88–99.
- [41] Simon Stockhardt, Benjamin Reinheimer, Melanie Volkamer, Peter Mayer, Alexandra Kunz, Philipp Rack, and Daniel Lehmann. 2016. Teaching Phishing-Security: which way is best?. In *IFIP International Conference on ICT Systems Security and Privacy Protection*. Springer.
- [42] Janos Szurdi, Balazs Kocso, Gabor Cseh, Jonathan Spring, Mark Felegyhazi, and Chris Kanich. 2014. The long “taile” of typosquatting domain names. In *23rd USENIX Security Symposium (Security ’14)*.
- [43] Christopher Thompson, Martin Shelton, Emily Stark, Max Walker, Emily Schechter, and Adrienne Porter Felt. 2019. The Web’s Identity Crisis: Understanding the Effectiveness of Website Identity Indicators. In *28th USENIX Security Symposium (Security ’19)*.
- [44] Alexander J.A.M. Van Deursen, Ellen J Helsper, and Rebecca Eynon. 2015. Development and validation of the Internet Skills Scale (ISS). *Information, Communication & Society* 19 (2015).
- [45] Melanie Volkamer, Karen Renaud, and Paul Gerber. 2016a. Spot the phish by checking the pruned URL. *Information & Computer Security* 24, 4 (2016), 372–385.
- [46] Melanie Volkamer, Karen Renaud, and Benjamin Reinheimer. 2016b. TORPEDO: tooltip-powered phishing email detection. In *IFIP International Conference on ICT Systems Security and Privacy Protection*. Springer, 161–175.
- [47] Melanie Volkamer, Karen Renaud, Benjamin Reinheimer, Philipp Rack, Marco Ghiglieri, Peter Mayer, Alexandra Kunz, and Nina Gerber. 2018. Developing and evaluating a five minute phishing awareness video. In *International Conference on Trust and Privacy in Digital Business*. Springer, 119–134.
- [48] Min Wu, Robert C Miller, and Simson L Garfinkel. 2006. Do security toolbars actually prevent phishing attacks?. In *24th ACM SIGCHI conference on Human Factors in computing systems*.