Site Inspector: Improving Browser Communication of Website Security Information

ERIC SPERO, School of Computer Science, Carleton University, Ottawa, Canada and School of Computer Science, The University of Auckland, Auckland, New Zealand
ROBERT BIDDLE, School of Computer Science, Carleton University, Ottawa, Canada

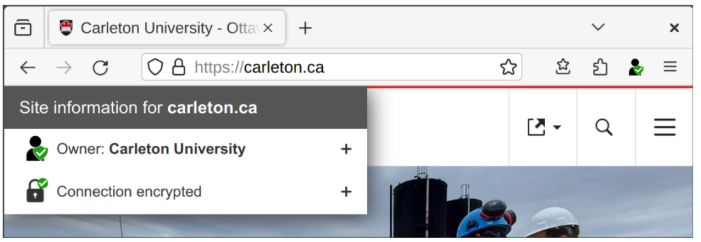

Fig. 1. Site inspector in situ.

Phishing sites exploit users' limited understanding of website identity to mimic legitimate sites. While X.509 certificates can provide crucial cues regarding a website's identity, current browsers fail to effectively communicate this information to users, even as phishing becomes an increasingly serious issue. To address this, we developed Site Inspector (SI), a UI tool that conveys website identity and connection encryption information, along with brief explanations of the relevant underlying security concepts. SI is implemented as a Mozilla Firefox browser extension, but the basic design could be integrated into any web browser.

SI organizes content in a three-tiered abstraction hierarchy, drawing on Ecological Interface Design. The top level presents an indicator of the website owner, if known, and also whether the connection is encrypted. The second and third levels offer progressively detailed explanations of the verification process. SI adheres to design principles aimed at educating users about security through the UI while overcoming associated challenges. Its text is concise and direct, respecting limitations in users' attentional resources and motivation to engage with security matters.

As a proof of concept for SI's principled design, we conducted a user study with 30 participants to evaluate its effectiveness in helping users differentiate real from fraudulent websites. Results suggested that SI improved users' ability to identify fraudulent sites. Future work will involve further testing with a larger user base, integrated SI directly into browsers, and ultimately a more widespread and improved validation process for certificates, with stronger verification and transparency.

Eric Spero acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) Alexander Graham Bell Canada Graduate Scholarship. Robert Biddle acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), RGPIN-2022-04887.
Authors' Contact Information: Eric Spero, School of Computer Science, Carleton University, Ottawa, Ontario, Canada and School of Computer Science, The University of Auckland, Auckland, New Zealand; e-mail: eric.spero@carleton.ca; Robert Biddle, School of Computer Science, Carleton University, Ottawa, Ontario, Canada; e-mail: robert.biddle@carleton.ca.

ACM 2471-2566/2025/08-ART34
https://doi.org/10.1145/3726867

CCS Concepts: • **Security and privacy** → **Usability in security and privacy**; • **Human-centered computing** → Empirical studies in HCI; *Interactive systems and tools;*

Additional Key Words and Phrases: Phishing prevention, web certificate indicators, user-centred security, security mental models, security indicators and warnings

ACM Reference Format:
Eric Spero and Robert Biddle. 2025. Site Inspector: Improving Browser Communication of Website Security Information. *ACM Trans. Priv. Sec.* 28, 3, Article 34 (August 2025), 31 pages. https://doi.org/10.1145/3726867

1 Introduction

Phishing sites trick users into giving up sensitive data by mimicking trustworthy sites; they exploit users' limited understanding of the site's true identity. X.509 certificates can provide reliable information about a site's identity, but this is not communicated to users at present. With phishing attacks on the rise [2], and AI threatening to make these attacks more effective [14], the need for effective security indicators is greater than ever. This article presents a new approach in communicating this and other information pertinent to trust decisions to ordinary users, implemented as a Mozilla Firefox browser extension called **Site Inspector (SI)** (Figure 1).

Certificate Authorities (**CAs**) offer special certificates that include assurances of website identity: **Organization Validation** (**OV**) and **Extended Validation** (**EV**) certificates. The *Organization* field of these certificates provides the verified name of the organization responsible for the website's domain name. Only legitimate websites should ever be able to acquire this type of certificate, making this information valuable in protecting against fraudulent website attacks (we discuss this further below). This information is not made accessible to ordinary users through the browser UI: the 'lock' symbol next to the URL merely indicates that a website has a certificate issued by a CA, which ensures encrypted communication, but often does not have assurances of identity.

Browsers formerly featured indicators for EV certificates.[1] When a website had an EV certificate, the name of the verified organization was shown to the left of the URL. Google removed these indicators from Chrome in 2019, with their researchers citing a lack of effect on user behavior [30], and other browsers soon followed. However, browsers made no attempt to link the indicator to assurances of identity, and therefore its lack of influence on behavior is unsurprising. Google researchers leave open the possibility for UIs to educate users on identity indicators [30], thereby reclaiming their value. We see the present work as aligned with this perspective.

There are infrastructural issues surrounding identity verification,[2] as Thompson et al. [30] and others have noted, including the occasional granting of EV certs to malicious entities and a general lack of transparency about the verification processes. We acknowledge that these operational issues undermine the effectiveness of OV and EV certificates, and solving them must be part of a comprehensive solution that includes UI design as an integral part. However, the UI aspect of this problem can be addressed independently, and we believe that better communication of even the current implementation of certificates would benefit users.

Educating users about security matters through the user interface presents significant challenges, particularly those related to user knowledge, limited time and attention, and motivation for engaging with security issues. In previous work, we introduced a design paradigm called **Security Awareness Visibility and Evaluation** (**SAVE**) [25] to help overcome these challenges. We designed SI following SAVE principles, which distinguishes it from prior security indicators. We describe SAVE in Section 2.2.

[1]Browsers have never featured indicators for OV certificates.
[2]See Section 4.7.1 for more on these issues.

SI communicates three key pieces of information from X.509 certificates: domain name, identity (i.e., the owner of the site, if known) and confidentiality (i.e., whether the connection is encrypted). The first two can be used to distinguish fraudulent from real sites, and the latter guarantees that data sent by the user will not be intercepted in transit by a third party.

Identity and confidentiality are the outcome of specific processes that the user must have a basic understanding of in order to determine how this information should influence their trust decisions. SI includes brief, non-technical explanatory text organized using an abstraction hierarchy, informed by **Ecological Interface Design** (**EID**) [32]: In two levels of toggleable panels reveal further details about identity and confidentiality. This organization prevents first-time users from being overwhelmed, while allowing experienced users to skip information they already know.

SI also contains supplementary, experimental functionality that allows us to explore a distinct but related topic: the influence of visual appeal on trust decisions. SI can optionally 'uglify' the visited website if the website offers no assurance of identity (i.e., if it does not use an EV or OV certificate). Uglification involves inverting the website's color palette, and changing its font family to monospace. This aspect of SI relates to other research we have done, which we discuss, but it does not pertain to our main research question.

We conducted a user study with 30 participants to provide a preliminary evaluation of SI's effectiveness in helping users detect fraudulent websites and found it to be effective.

A full functional and technical description of SI is given in Section 3. Next, in Section 2, we provide a summary of background literature on topics relevant to SI, including mental models (Section 2.1), SAVE (Section 2.2), X.509 certificates (Section 2.3), and prior work on website security indicators (Section 2.4). In Section 4, we report on our user study, including its methodology, results, and a discussion.

2 Background

2.1 Mental Models

The concept of mental models suggests that people create simplified internal representations of reality to understand and interact with the world [12, 15, 22]. These representations have the same causal structure as their real-world counterparts, meaning they work in the same way. Mental models are essential in any cognitive activity that requires interaction with or reference to the external world, as they represent an individual's understanding of the world. These models enable people to interpret, predict, and respond to external events.

Mental models are simplified representations, and 'good' mental models are those that support successful interactions with the external world [15, 20]. For example, operating a car requires a much simpler model than repairing a car.

Mental models are constructed from the experience of particular objects and events [4]. Greater experience leads to more accurate models, and therefore to better-informed decisions.

The 'mental models' construct has helped shed light on user security behavior. It has been shown that non-experts have quite different mental models of security compared with experts [3]; users' models are more simplistic [8]. Studies have related weak mental models with insecure behavior. For example, users' limited understanding of end-to-end encryption and its benefits reduces their motivation to adopt secure tools [1], and users' limited understanding of wi-fi blinds them to the concept of a malicious access point, leading to risky behavior [16]. Wash and colleagues [33] found that mental models of malware correlated with following security best practices. Camp [10] was an early and strong advocate for using metaphors to influence users' mental models so that they better understood security risks.

2.2 Security Awareness Visibility and Evaluation (SAVE)

This section summarizes prior work by us published elsewhere [25].

Ideally, security matters are handled automatically by infrastructure and users are kept out of the loop. However, there are many security problems that infrastructure cannot cover, and users must protect themselves. For example, web certificates cannot tell a user whether they should trust a website. But very often users in this situation can be given supporting information to aid their decision-making. Web certificates can link a website to a legal entity, which may be familiar to the user; if unfamiliar, users can investigate it.

Users develop working models of the software they use mainly through interactions with the user interface [21]. We think UIs should provide information to support users when they are in an unprotected space like the one just mentioned. Unfortunately, system designers often hide security information for users, perhaps for pragmatic reasons (e.g., development costs, fear of overburdening users). There is then no opportunity to develop an adequate mental model of the situation. These systems put users at risk by requiring them to make uninformed decisions with major potential security consequences.

SAVE [25] is a high-level framework for introducing supporting security information in user interfaces to help users make informed decisions where there are gaps in security infrastructure. SAVE has four principles:

Preparedness: When systems may require users to make security decisions, help them develop a mental model that will support informed decision-making.

Visibility: Make the system's security state, and the consequences of security actions, easy for users to find in the UI.

Intelligibility: Information should be provided in terms that ordinary users will find familiar and understandable.

Veridicality: Supporting information must aim toward giving users mental models that accurately reflect reality (e.g., not metaphors or oversimplifications).[3]

Frugality: Demand as little of users limited attention as possible, without sacrificing intelligibility.

We think that EID [32], particularly its abstraction hierarchy, will be useful in designing SAVE applications. See Section 3.2 for a brief description of EID and how it relates to SI.

2.3 X.509 Certificates and Support for Detecting Fraudulent Websites

As mentioned in the previous section, infrastructure alone does not cover all potential security threats, and there are many situations where users must rely on their own understanding to keep themselves protected. The tool we present in this article is designed to help users in one such situation: deciding whether to share sensitive information with a website that could be real or fraudulent.

X.509 certificates [11] are integral to the HTTPS protocol and can provide information about the identity of the organization that controls a website's domain. These certificates are typically issued by CAs, who may conduct background checks during the signing process. Browsers come pre-installed with a set of trusted CAs.

There are three levels of organization verification. **Domain Validation** (**DV**) offers the least assurance of identity, confirming only that the certificate applicant controls the domain. OV verifies the existence of the organization (e.g., by referring to government records). EV involves additional steps to verify the organization's legitimacy, linking it to a legally recognized entity [9].

[3]This was implicit in our earlier description of SAVE, but not an explicit principle until now.

Unfortunately, at present browsers do not attempt to communicate this information to ordinary users, and this infrastructural support for owner verification is therefore mostly unknown. The certificate can be viewed after a complicated series of menu-item selections, but the presentation format is complex, and the content too arcane for ordinary users to understand.

Users should be told whether the website they are on has been verified by a CA, and the name of the organization if it is known. But in order for the 'verified' descriptor to have meaning one must also understand some basic facts about the verification process.

2.4 Prior Work on Web Security Indicators

Early Work. In a 2006 article, Wu et al. [34] evaluated several 'security toolbars'—browser extensions that augment the toolbar to communicate security information to users—in their capacity to help users detect phishing sites. None were found to be effecive. In 2007 Schechter et al. [24] found SSL indicators to be ineffective because users simply ignore them.

Communicating Identity. In 2009 Biddle et al. [6] created a novel design for certificate indicators that explained the identity and confidentiality assurances provided by SSL certificates using ordinary language. Compared with the existing certificate indicators in Internet Explorer 7.0, the alternative interface was found to be much more effective at communicating security-relevant information to users, and this resulted in better security decisions.

Mental Model Builder for Web Certificates. Stojmenović et al. [27] tested a visualization tool that provided a summary explanation of the process of fetching and checking X.509 certificates when visiting websites, and the assurances provided by each validation level (DV, OV, EV, and no certificate). An in-lab study showed that the interface changed participant attitudes to websites: there was a notable decrease in user's stated willingness to enter sensitive information into websites without certificates, and an increase for websites with EV certificates. Potential issues with this design include it's large size (it occupied an entire screen) and how much attention it required. It also includes a number of operational details that may not be strictly relevant to users' trust decisions.

Using Operating System Notifications. Stojmenović et al. [28] developed a browser extension that produces operating system notifications on page load that state whether the owning organization is known, and if so, who they are. It was found that the notifications were largely ignored; users seemed to distrust the notifications, suspecting that they might be malware. The authors speculated that users may have become habituated to OS notifiations through regular desktop OS use. The icons also lacked aesthetic polish, which may have contributed to the perception that they were malware.

EV Indicators Removed from Browsers. Until 2019 browsers provided a special indicator for websites that had EV certificates: the organization name was listed next to the URL. Google Chrome was the first to remove this feature. Google's reasoning is explained in a 2019 article by Thompson et al. [30] which includes a large-scale study on EV indicators. This study showed that EV indicators did not affect most user behavior, and they concluded that it likely did not help users defend against phishing attacks [30]. The authors also point out broader issues with EV certificates and indicators: that malicious entities have received EV certificates in the past, and users do not pay attention to UI indicators.

Browsers have never attempted to explain to users the meaning of EV indicators, making it unsurprising that they were ineffective. The presence of an organization name next to the URL provides several assurances: that the organization who controls the domain passed a multifactor background check to ensure they were a legitimate business, and they are linked to a known legal

entity who in theory can be pursued if the website causes users harm. Without any meaningful explanation of what the indicator indicates, the presence of an organization name is anomalous.

Thompson et al. did not give up on the idea of identity indicators entirely. They note that the presence of the EV indicator did cause users to click the "Page Info" bubble—a pop-up window that appears when the indicator is clicked—more often. The researchers state that having a prominent UI feature related to certificates could be an opportunity to educate users about identity indicators.

This work takes a step in the direction suggested by Thompson et al., aiming to develop a UI element that educates users about the key security assurances offered by X.509 certificates, with the goal of eventually reintroducing improved security indicators in browsers. We designed SI in accordance with SAVE framework (Section 2.2), which addresses challenges associated with introducing security 'mental model builders' in UIs. Through SI and the accompanying user study, we aim at assessing whether users can be taught a sufficient understanding of the security assurances in EV and OV certificates through a small, unobtrusive user interface element that could be feasibly integrated into existing browsers.

Our Approach. We attempt to integrate the lessons taught by the history of browser web certificate indicators and the related body of literature. Our approach to creating a SAVE-compliant security indicator can be outlined as follows:

— Provide a simple 'default' view that communicates key information minimally;
— Explain key concepts to help users understand the basic indicators;
— Make explanations brief, and available on-demand;
— Communicate only essential status items: domain name, identity status, and encryption status;
— Communicate identity and encryption status separately;
— Ensure a professional look and feel to avoid being perceived as malicious;
— Implement the interface as a fully functional prototype to increase realism.

2.4.1 Manipulating Visual Appeal to Convey Website Certificate Information. Stojmenović et al. [29] found a relationship between a website's perceived visual appeal and its perceived security, where a low visual appeal website was judged by users to be less secure than a high visual appeal website. The 'high' visual appeal website was a clone of a real website for the City of Gold Coast, Australia.[4] The 'low' visual appeal website featured the same content and layout, but used discordant color combinations, and featured color-inverted images. The websites were paired with the browser Chrome featuring standard web certificate indicators and a URL.

Participants reliably rated the 'high' visual appeal website as more secure than the 'low' visual appeal website regardless of the *actual* security status of the website indicated by the web certificate indicators. Remarkably, this was true even in cases where the certificate indicator of the 'high' visual appeal website read "Not secure". These findings demonstrate not only the influence of visual appeal on website security, but the failure of web certificate indicators to indicate security.

SI includes an experimental feature that reduces the visual appeal (i.e., 'uglifies') of DV and no-certificate websites on-the-fly in realistic web usage scenarios.

3 Site Inspector: Prototype for a SAVE-Compliant Security Indicator

SI's[5] primary purpose is to help users develop security-related mental models of a website by interpreting and explaining information from its X.509 certificate. The mental models should support user decisions about whether to share sensitive information with that website. SI's user interface

[4]In earlier work [26] users rated this website as *high in visusal appeal* and *unfamiliar*.
[5]The source code for SI is available at https://github.com/speroe/site-inspector

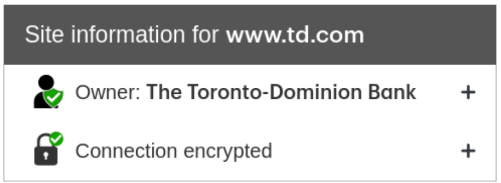

Fig. 2. Site inspector level 1: Basic status view.

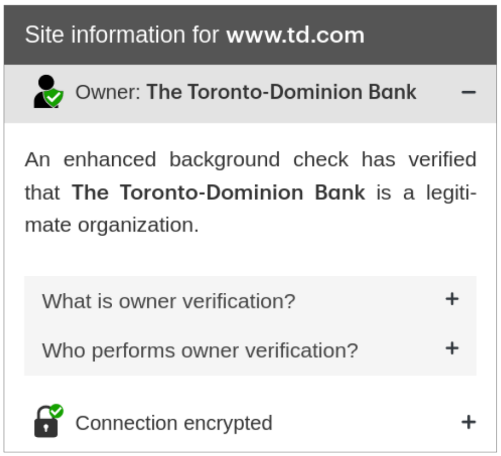

Fig. 3. SI Level 2: Status elaboration for owner identity.

is contained in a draggable floating window placed in the top-left of the website by default.[6] It features a 'basic' view that minimally communicates basic status information from the X.509 certificate pertinent to trust decisions through icons and a few words. Expandable panels provide explanations of the processes and mechanisms that underlie these statuses.

SI also has a supplementary, optional feature that reduces the visual appeal of (i.e., 'uglifies') that have a DV or no certificate.[7] The uglification process involves inverting the website's entire colour palette and changing the font family to monospace. Examples of this are shown in Figures 20–23 in Appendix B. We included this feature in hopes of expanding on findings from previous work [29], described in Section 2.4.1. We believe that on-the-fly uglification of websites might help induce a shift in users' attention toward potential security risks.

3.1 Functional Overview

3.1.1 Main Interface: Floating Window. The default view (Figure 2) shows three key pieces of information: the website's domain name; the identity of the organization that owns the website,[8] if it is known; and whether the connection to the website is encrypted.

Textual Content and Hierarchical Organization. SI's content is organized in a three-level hierarchy: *basic status* (e.g., Figure 2), *status elaboration* (e.g., Figures 3 and 5), *mechanism/process*

[6] A proper implementation of SI's UI would be built directly into browsers.
[7] A proper implementation of uglification would include a disclaimer message explaining that the changes in appearance are manipulations of the existing website, and are not properties of the site itself. This should help prevent unwanted side effects such as reinforcing the idea that visual appeal is related to website security.
[8] SI assumes that the organization named in the `Organization` field of the X.509 certificate also controls the website. In reality, the situation is a bit more complicated. This is discussed in Section 4.7.1.

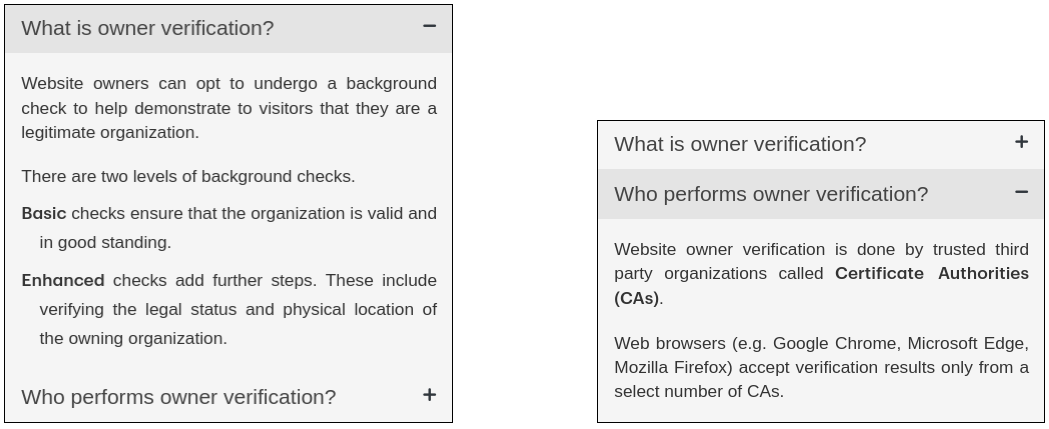

(a) What is owner verification?

(b) Who performs owner verification?

Fig. 4. SI Level 3: Mechanism/process description for identity verification.

Fig. 5. SI Level 2: Status elaboration for confidentiality.

description (e.g., Figures 4(a), 4(b), and 6). SI uses an *accordion menu*: deeper levels are toggled in and out of view by clicking headers featuring a '+' sign on the right.

Level 3 content (i.e., mechanism/process description) is always the same regardless of the website context. Levels 1 (basic status) and 2 (status elaboration) contain references to domains and organization names which vary from site to site. The rest of the content varies to reflect differences between the various X.509 certificate types and the respective assurances they offer. For example, when the connection is unencrypted, the Level 2 text says: *The data you exchange with [domain name] can be intercepted by attackers.* The figures used in this section show SI in the context of a website with an EV certificate; see Figures 13–15 in Appendix A to see how SI appears when the website has no certificate.

Icons. We created five icons to communicate three identity verification statuses: no validation or DV, OV, and EV, as well as two connection confidentiality statuses: not encrypted and encrypted. These are shown in Figure 7.

The icon indicating 'no identity verification' (DV certificate, no certificate) shows a blurred person shape with a question mark. This is meant to connote that the identity is *unknown* rather

Fig. 6. SI Level 3: Mechanism/process description for encryption.

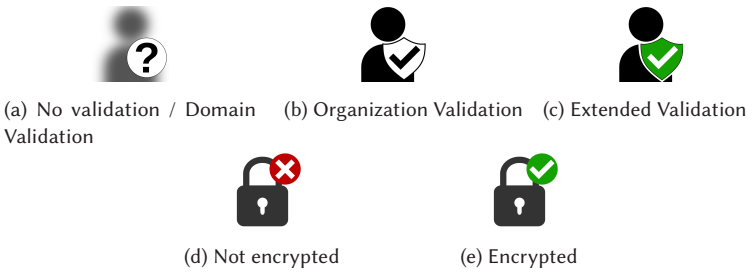

Fig. 7. Icons used in Site Inspector to communicate identity verification (top row) and confidentiality (bottom row) statuses.

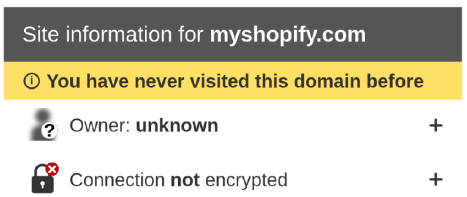

Fig. 8. The yellow alert banner is shown to users when visiting a domain for the first time.

than *dangerous*. This is consistent with our overall approach of imparting accurate understanding to the user, not merely modifying behavior. In real usage, many authentic websites that people visit use only DV certificates (e.g., *google.com*)—an unfortunate issue that we discuss in Section 4.7.1—so labeling DV sites with *danger* indicators might lead users to disregard the indicator.

Alert for First-time Visits. When visiting a domain for the first time, a yellow banner is added to SI featuring the message "You have never visited this domain before." (see Figure 8). This is to help users detect fraudulent versions of websites they are familiar with and visit frequently.

3.1.2 Browser Action Popup. User configuration options are provided in the "browser action popup". The "browser action" is a button in the browser's toolbar that is associated with an extension. When clicked, it can produce a "popup" displaying rendered HTML. Users can toggle the visibility of SI globally, or for just the current domain.

3.1.3 Developer Options. Pressing Ctrl+Alt+Y opens a set of developer options in a new tab, which helped us run the user study (presented in Section 4). They are hidden because we did not want our participants to access them. The options are "enable uglification", which applies the uglification changes to the website's CSS mentioned at the beginning of this section, and "clear local storage", which ensures users start the study with no SI-specific browsing history.

3.2 Site Inspector and Ecological Interface Design

EID [32] is a mental models–based approach to the design of human interfaces to complex engineering systems like those used in nuclear reactors. Interfaces in this domain are necessarily complex,[9] and EID helps manage this complexity by organizing interfaces in an abstraction hierarchy [23]—that is, in several encapsulated layers. Higher levels correspond to functional aspects of the system (e.g., the overall goals of the system), and lower levels to physical aspects (e.g., the equipment that realizes the function).

This design strategy enables operators to ignore aspects of system complexity that are not currently relevant, reducing cognitive workload and leading to better performance. During normal system functioning the operator need only consider the top-level 'functional' view; to diagnose system faults they can progressively descend through the layers to make sense of the issue.

Web users face different security challenges compared with operators of complex systems. Users do not need control over the underlying security mechanisms, as responsibility lies with system designers. Users simply need to learn enough about a security issue that the high level information we communicate makes sense to them. Consequently, security interfaces can be relatively simple when compared with canonical EID interfaces. However, unlike the operators of complex systems, ordinary web users lack formal security training and technical expertise, and they have limited time and attention to devote to their task since security is a secondary concern. While SI asks end-users to consider only a fraction of the information EID operators do, in both domains there is a critical need for managing cognitive load.

In presenting security information to end-users it would be easy to overwhelm by showing too much information at once. The abstraction hierarchy allows explanatory text and its associated complexity to be hidden until it is needed. Once this information is learned, it may never be needed again. Users who have no need for the explanatory text can rely exclusively on the simple high-level view, keeping the cognitive burden of understanding the current security-related system status to a minimum.

3.3 Site Inspector and X.509 Certificates

The diagram in Figure 9 summarizes the relationship between SI and X.509 certificates, using the EV certificate for `td.com` as an example. Level 1 presents several fields of the X.509 certificate that contain useful but arcane information. The `Organization` field in `Subject Name` (if present) indicates the name of the verified organization, while `Certificate Type` under `Certificate Policies` specifies the level of verification conducted. Additionally, if one of the listed `Purposes` is `Key Encipherment`, it signifies that the connection to the website is encrypted. Level 3 provides

[9]Unlike ordinary web use, complex systems require that operators have full control, necessitating complex interfaces (see Ref. [32] for more details).

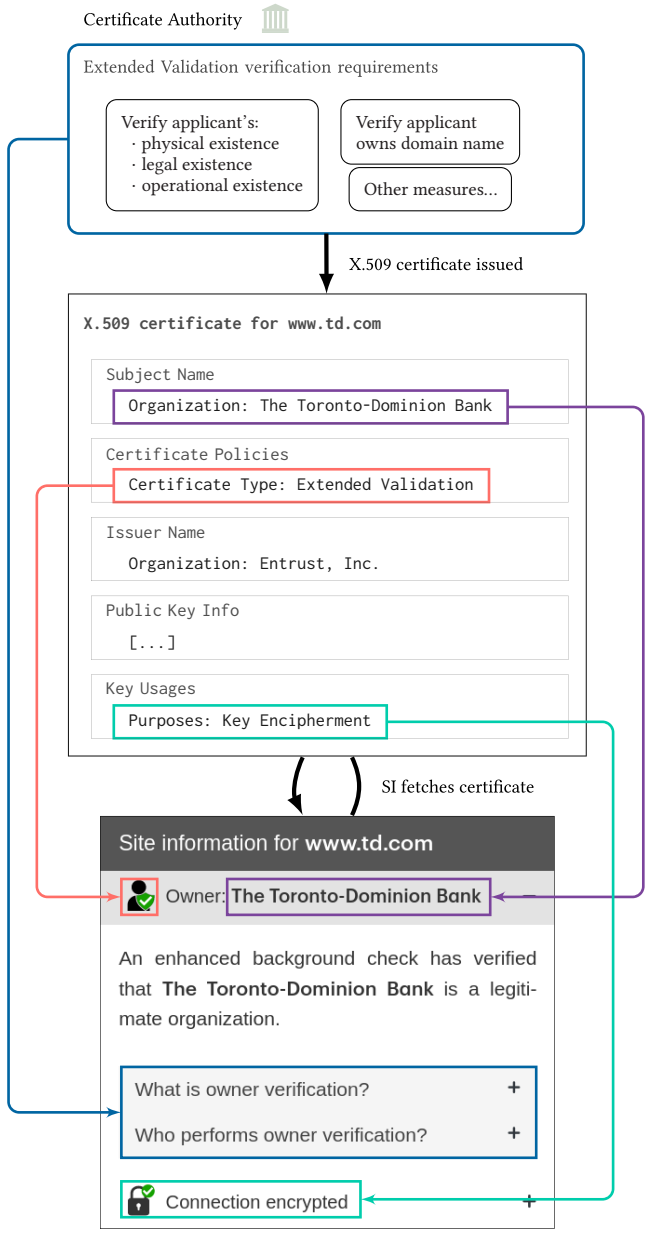

Fig. 9. SI makes obscure technical processes and artifacts visible.

context by explaining the real-world processes and mechanisms that give meaning to the X.509 fields. Level 2 serves as a succinct bridge between Levels 1 and 3.

3.4 Uglification

People become habituated to the user interface elements they encounter frequently [31]. In the case of more subtle security indicators like SI, there is a risk that users might neglect to notice the indicator in an important situation. Prior work has shown that users naturally associate inverted

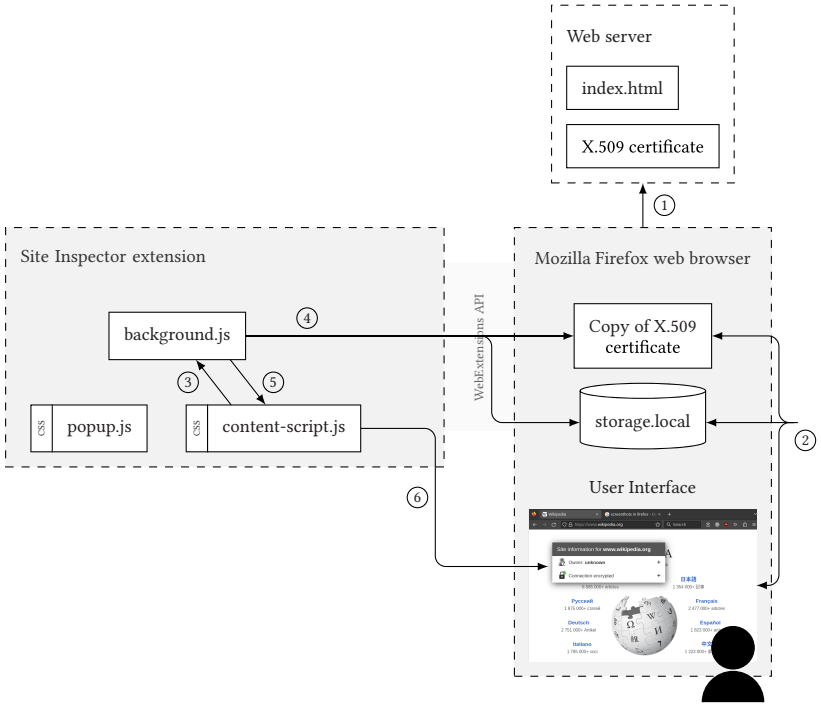

Fig. 10. Site Inspector architecture/basic process flow on page load. The numbers indicate the steps followed in execution.

website color palettes with insecurity [29]. This cue is difficult to ignore, and could serve as a valuable complement to a tool like SI. We are interested in investigating whether the experimental uglification feature amplifies SI's effectiveness.

3.5 Technical Overview

SI was implemented as a browser extension for Mozilla Firefox.[10] The diagram in Figure 10 shows the basic architecture of the extension, and an overview of the process flow that occurs each time a website is loaded.

The extension has three main components:

Content script (content-script.js) This is where modifications to the website's DOM are made. This script creates the SI interface and appends it to the DOM each time a page loads.

Background script (background.js) This script fetches the X.509 certificate and interacts with local storage. Local storage is used to log visits to websites (to determine whether to show the yellow 'first visit' banner) and to save some user-specific parameters (i.e., configuration options and floating window position).

Popup script (popup.js) This modifies the appearance of the 'browser action popup' (the window that appears when the extension's icon next to the URL bar is clicked). Here is where users can disable/enable SI globally, or for the current domain.

[10]We chose Firefox over the alternatives because it is the only major browser that provides WebExtensions support for inspecting X.509 certificates [19].

The components communicate with each other through messages. Every component listens for messages, and its response is conditional on the contents of the message it receives.

The process flow depicted by the numbered arrows in Figure 10 is as follows:

(1) The browser negotiates a TLS connection to a website.
(2) A copy of the X.509 certificate is retained and made available to extensions via the WebExtensions API. SI makes a record in the browser's local storage indicating that the site has been visited (if such a record does not already exist). The website is rendered.
(3) After the page load, a message is sent to the background script to initiate Step 4. Ensuring the page has loaded before proceeding guarantees that modifications made to the DOM by SI in Step 6 will not be overwritten.
(4) The background script fetches a copy of the X.509 certificate using `getSecurityInfo()`. It also checks the browser's local storage to see if the browser has visited this page before.
(5) A copy of the X.509 certificate and the 'visited before' status is passed to the content script in a message.
(6) The content script creates SI's floating window and appends it to the website's DOM.

3.6 Summary

In this section, we have presented SI: A browser extension to help users understand identity and confidentiality assurances offered by X.509 certificates. SI reads key information from X.509 certificates and presents it in an easy-to-understand format, while also explaining any necessary background concepts. SI conforms to the SAVE principles described in Section 2.2. The information in SI is organized according to an abstraction hierarchy allowing for the progressive disclosure of information as it is needed.

In the next section, we describe an effort to test the effectiveness of SI in a pseudo-realistic context.

4 User Study

As an initial proof of concept for the principled design of SI, we conducted a user study with 30 participants. Our primary goal was to determine whether participants using SI are better able to distinguish real sites from imposter sites compared with those who do not use it. Additionally, we have a secondary interest in testing an experimental feature of SI that reduces the visual appeal of (i.e.'uglifies') websites without a verified owner, to see if this feature could provide further assistance to participants.

Our research question is: Does SI help users differentiate real sites from imposters? To answer this, we formulate three hypotheses. Additionally, we also anticipate that uglifying websites will be even more effective than SI by providing a visceral cue for potential fraudulence. Since this feature is more experimental, we do not formulate specific hypotheses for it.

4.1 Method

In a user study, we asked participants to look at eight different websites—four 'real' and four 'imposter' sites—and answer Likert scale and semi-structured questions about them. The purpose was to see if the SI interventions helped users differentiate the two kinds of sites. This study was given ethical clearance by our Research Ethics Board (Clearance #118267).

4.1.1 Participants. We recruited 30 participants using recruitment posters posted (a) around our University campus and (b) on a Facebook group for recruiting HCI participants. Nine participated

Table 1. The URLs for Real Sites, along with the Corresponding Company Name

Company name	URL
Autozone	`www.autozone.com/signin`
T-Mobile	`account.t-mobile.com/signin/v2`
Silicon Valley Bank	`www.svbconnect.com/auth`
M&T Bank	`www3.mtb.com/log-in`

Table 2. The URLs for Imposter Sites, along with the Corresponding Company Names and Domain Names of the Real Sites They Impersonate

Company name	Imposter URL	Real domain
Cabela's	`cabelas.certlo.com`	`cabelas.ca`
Tiffany & Co.	`tiffany.secwww.com`	`tiffany.ca`
City National Bank	`cnb.secureuserlogin.com`	`cnb.com`
Chewy	`chewy.http-s.net`	`chewy.com`

in person, and 21 over Zoom. The vast majority, if not all, seemed to come from the Facebook group. The age of participants ranged from 18–53, with a mean of 27.5 and a median of 25. Gender counts were 25 female, 4 male, and 1 other.

All except one were university- or college-educated. Only three reported having an occupation or area of study in a computing-related field: one technical (Computer Engineer), and the other two more business-oriented (Technology Innovation Management).

4.1.2 *Materials.*

Websites. All websites used in the study were fully functional, and participants could interact with them if they wished. Figures 16–23 in Appendix B show screenshots of the websites used.

The four 'real' sites used in the study were the login pages for legitimate sites owned by legitimate organizations, and accessed through domains that these organizations control. Table 1 shows the company names associated with the 'real' websites used in the study, and the URLs used to access them.[11, 12]

The four imposter websites were created by us using the website downloader `wget2` [13] to copy the login pages of four legitimate websites. The resulting HTML file was hosted on a web server running in our research lab.

Participants accessed the imposter sites using domains controlled by us. The domain names were generic (e.g., `certlo.com`), and subdomains were used to reference the imitated company name (e.g., `cabelas.certlo.com`). This approach was designed to emulate real world practices: phishing sites are typically short-lived [5], and buying new domains for each site would be financially infeasible for attackers. Table 2 shows the subdomain and domain pairings used, along with the corresponding domain names and company names of the imitated websites.

[11]Technically, participants accessed the sites by clicking a hyperlink bearing the company name only. The URL was obscured by the link, which is expected in a phishing scenario.

[12]Our study was completed before the high-profile collapse of Silicon Valley Bank.

The imposter sites were accessible only from inside our network.

DV certificates for our imposter sites were obtained from Let's Encrypt. Our participants, therefore, saw the 'lock' icon next to the URL when visiting these sites, and the sites' URL bore the `https` scheme.

Cosmetic Issues with Websites. All four imposter sites were structurally identical to their real counterpart. There were two cosmetic differences, however, which affected the Tiffany's site: custom icons were replaced with placeholders, and the default sans-serif font was used. Both were noticed in only a small handful of cases, and it did not appear to signal to participants that the website was illegitimate.

Likert Questions and Rationale. The main source of quantitative data in our study comes from 5-point Likert scale questions. Participants are asked to answer the following five questions after viewing each website:

(1) I would log in to the linked-to website per the email's instructions;
(2) The linked-to website is visually appealing;
(3) The linked-to website looks professional;
(4) The linked-to website is easy to use;
(5) The linked-to website is of a legitimate business.

The first and fifth are the most important to us as they are intended to gauge perceptions of the website's security. We ask the second, third, and fourth because they are plausible alternative explanations to answers to security questions. For example, if we find that participants rate imposter sites as less secure, but also as less visually appealing, the visual appeal judgments may be driving the security ratings. These three questions also serve as checks for major flaws in the imposter websites we created. Our imposter sites should be as visually appealing, professional-looking, and easy to use as the others.

At the conclusion of the study, those who viewed the SI interface were asked the following questions to gauge their perceptions of its effectiveness:

(1) I found the information in the site information tool easy to understand;
(2) The information provided by the site information tool was relevant to my login decision.

4.1.3 Procedure. Participants were given the option of doing the study in-person in our lab, or over via Zoom teleconference. We wanted to provide a remote participation option as COVID-19 was still a concern during the course of data collection. In both cases, the study was delivered by a LimeSurvey [17] instance hosted in our lab.

When running the study over Zoom, participants were given remote access to the researcher's instance of Firefox which had SI installed.

Participants were assigned to one of three conditions: 1. Control condition (the extension is not active), 2. SI (participants see the SI floating window on each page load), and 3. SI plus uglification (the same as Condition 2, with additional CSS changes for DV websites: inverted colors and the font-family is changed to monospace). In the early stages of the study, conditions were assigned pragmatically: participants were assigned based on which condition we needed more data on at that time. After approximately 10 sessions, we began assigning participants in the following order: Control, SI, SI plus uglification.

Participants were asked to read the following written instructions at the beginning of each session:

> You are presented with 8 simple scenarios relating to online accounts. The scenarios all have the same basic form:
>
> — You receive an email from a company asking you to do something that requires logging in to their website. The email provides a link to the website to facilitate the login process.
>
> Some assumed context:
>
> — You have an existing online account with the company named in the email, for a website that **looks** like the one that is linked to
>
> — The email **seems** totally legitimate to you (but it might not be)
>
> Each scenario will feature a different company and linked website.
>
> Of the eight links we ask you to click on, **some will be legitimate** (i.e., link to a page that really belongs to the company named in the email) **but some will not**.
>
> Your (the participant's) task will be to click on the provided link to visit a website and answer a few quick questions about it.
>
> If anything about this is unclear, please let the researcher know.

Participants viewed eight websites each, in randomized order. Four sites were legitimate and accessed through a legitimate URL; the other four were 'imposter' (psuedo-phishing) websites created by us, and accessed through domains owned by us. Before visiting a new website participants were shown the following text to remind them of their task:

> **Scenario:**
>
> You receive a legitimate-looking email that appears to be from [company name] asking you to perform some action that requires logging in to their website. They provide a link to facilitate the login process.
>
> You have an existing account with a website belonging to [company name] that looks identical to the site they link to. However, it is always possible that the site they link to could be an imposter.
>
> This is the link provided in the email: [company name]

Participants in the SI and SI+U conditions would receive a verbal explanation closely resembling the following upon their first encounter with SI:

> The box you see in the top left is something we have added to the browser. It is meant to tell you some information about the website. The information is sourced from a third party trusted by the browser. There is no deception here; you can trust this information.

Methodological Adjustments. We allowed for the possibility of making methodological adjustments early in the study provided they improved our ability to evaluate SI, which was our main goal. We made two such changes, detailed below. No data were discarded after making these changes.

(1) *Replacing Vancity website.* One of our original 'real' websites was the Vancity Bank website (https://www.vancity.com). The login page for the Vancity bank used an OV cert, whereas the main site used an EV cert. One early participant noticed this discrepancy, which caused them to doubt the authenticity of the website. The participant's behavior was a reaction to an unusual situation in the real world pertaining to this particular website, not to identity

information in itself nor SI, so we decided to use another banking website (M&T Bank) that featured an EV on its login page and the main page.

(2) *Adding font manipulation to uglification process.* The uglification process initially only involved inverted the colors of websites. However, multiple early participants did not find the uglified sites ugly. In hopes of enhancing the uglification effect and making it more universally effective, we changed the uglified websites' font to the monospace family—a style typically used in terminal emulators and IDEs, and highly uncommon for modern professional websites.

4.2 Research Question and Hypotheses

Our study aims to provide an initial assessment of the potential benefits of SI, determining whether a more extensive study is warranted. The study is guided by the following research question and hypotheses:

RQ Does SI help users differentiate real sites from imposters?

H1 Participants who see the SI interface will consider imposter sites less secure than real sites

H2 Participants who see the SI interface will consider imposter sites less secure than those who do not see the SI interface.

H3 Participants who see the SI interface will consider real sites more secure than those who do not see the SI interface.

4.3 Quantitative Results

We use boxplots to provide a general overview of our results. In the boxplots presented in this section we aggregate responses by condition and *website legitimacy* (i.e., real or imposter): each participant's four responses about the four real websites are averaged, as are their responses about the imposter sites. This approach minimizes the influence of extreme responses and ensures the independence of individual data points. For a detailed breakdown of response distributions by website, see the boxplots in Figures 24–28 in Appendix C.

As we aggregated responses, we treat the data as continuous. After confirming the distributions were appropriate, we used *t*-tests to test our hypotheses.

After describing the boxplots, we report the results of *t*-tests on groups of data relevant to our hypotheses, focusing on the normal SI condition and the control condition. To assess how the experimental SI + Uglification condition influenced participant judgments, we later present a set of *t*-tests that parallel our hypothesis tests.

4.3.1 Security Questions. Figure 11 shows boxplots of the two questions directly related to security.[13] The overall pattern of responses seems similar across the two questions, which is expected as we think both are closely related to feelings of security and trustworthiness.

We see three noteworthy patterns which can be observed in the responses to both questions, with the first two directly relating to our hypotheses.

(1) Participants in both SI conditions report being less likely to log in to imposter sites than those in the control condition, and they were more likely to disagree that imposter sites were run by legitimate businesses.

(2) Participants in both SI conditions report being less likely to log in to imposter sites than real sites, and they were more likely to disagree that imposter sites were run by legitimate businesses than real sites.

[13]Condition 3 is labelled as "SI+Uglification" for brevity, but a more accurate label would be "SI+Uglification of DV sites". Only imposter sites would be uglified in this condition; real sites appear as normal.

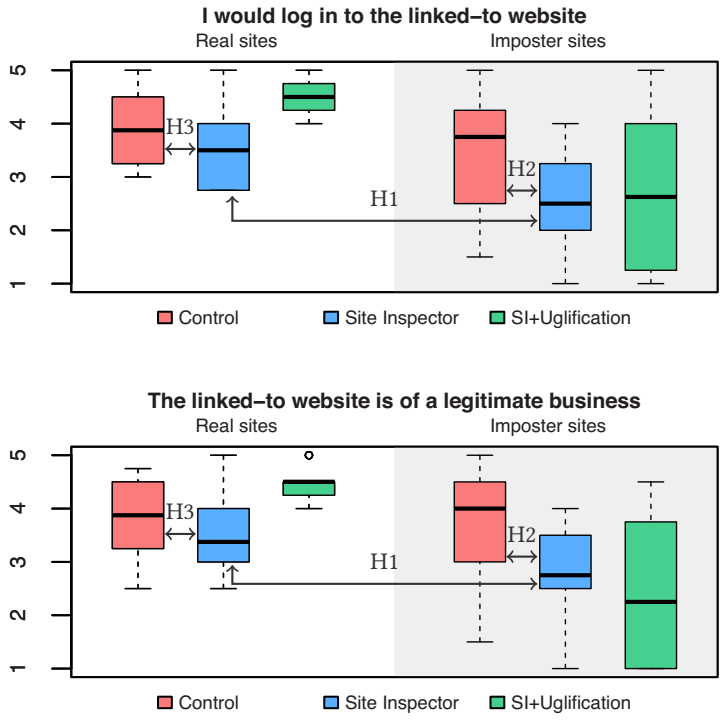

Fig. 11. Boxplots of security-related questions aggregated by website legitimacy (real or imposter).

(3) Participants in the SI+Uglification condition (SI+U) report being more likely to log in to real websites than those in the control condition, and they were more likely to agree that real websites are run by legitimate businesses.

Note that the second pattern is not observed in the control condition: participants are equally likely to log in to real and imposter sites, and they are equally likely to agree that real and fake sites are run by legitimate businesses. This reflects how phishing sites succeed.

Hypothesis Tests. We test the three hypotheses with one-sided t-tests comparing participant responses to the two security questions across different combinations of condition and website legitimacy (i.e., real or imposter websites). The t-test results are summarized in Table 3; means and standard deviation for each group are provided in the main text. Figure 11 shows boxplots of participants' response to the two security questions, and overlaid arrows indicate the group comparisons made in the t-tests.

To test H1: *Participants who see the SI interface will consider imposter sites less secure than real sites*, we compare how participants in the normal SI condition (i.e., without uglification) responded to security questions about real sites with fake sites.

Participants in the SI condition expressed a greater inclination to log in to real sites ($M = 3.58, sd = 0.86$) than imposter sites ($M = 2.58, sd = 0.98$). Participants in this condition also expressed a stronger belief that real sites ($M = 3.58, sd = 0.78$) were run by legitimate organizations than imposter sites ($M = 2.67, sd = 0.97$).

Incidentally, in the control condition, participants did not express a stronger inclination to log in to real sites than imposter sites, nor did they express a stronger belief that real sites were run by legitimate organizations than imposter sites.

Table 3. Summary of *T*-test Results for Hypothesis Tests

	Question	G1	Exp.	G2	df	t	p	Sig. lvl.
H1	Login	SI–Real	>	SI–Imposter	17.70	2.43	.013	*
	Legitimate	SI–Real	>	SI–Imposter	17.21	2.28	.018	*
H2	Login	SI–Imposter	>	Control–Imposter	17.75	1.87	.039	*
	Legitimate	SI–Imposter	>	Control–Imposter	17.65	2.24	.019	*
H3	Login	SI–Real	>	Control–Real	17.66	−0.90	> .05	ns
	Legitimate	SI–Real	>	Control–Real	17.76	−0.68	> .05	ns

The Exp. (expectation) column shows which group (G1/G2) we expect will have the greater means, and so the direction of the one-sided t-test. The remaining columns show degrees of freedom (df), t-test result (t), probability (p), and significance level (Sig. lvl.) with ns for 'not significant', and * for 'significant at .05'.

Table 4. Summary of *T*-test Results Involving SI+Uglification Condition

Question	G1	Exp.	G2	df	t	p	Sig. lvl.
Login	SI+U–Real	>	SI+U–Imposter	9.80	3.62	.002	**
Legitimate	SI+U–Real	>	SI+U–Imposter	10.17	4.53	< .001	***
Login	SI+U–Imposter	>	Control–Imposter	16.20	1.20	> .05	ns
Legitimate	SI+U–Imposter	>	Control–Imposter	17.32	2.32	.016	*
Login	SI+U–Real	>	Control–Real	12.37	2.52	.013	*
Legitimate	SI+U–Real	>	Control–Real	13.28	2.64	.01	*

These parallel the hypothesis tests, with SI+U replacing SI. The Exp. (expectation) column shows which group (G1/G2) we expect will have the greater means, and so the direction of the one-sided t-test. The remaining columns show degrees of freedom (df), t-test result (t), probability (p), and significance level (Sig. lvl.) with ns for 'not significant', * for 'significant at .05', and ** for 'significant at .01, and *** for 'significant at .001'.

To test H2: *Participants who see the SI interface will consider imposter sites less secure than those who do not see the SI interface.*, we compare how participants responded to security questions about imposter sites in the normal SI condition with the control condition.

Participants in the SI condition expressed a lesser inclination log in to imposter sites ($M = 2.58, sd = 0.98$) than those in the control condition ($M = 3.45, sd = 1.1$). Participants in the SI condition also expressed a weaker belief that imposter sites were run by legitimate businesses ($M = 2.67, sd = 0.97$) than those in the control condition ($M = 3.72, sd = 1.12$).

To test H3: *Participants who see the SI interface will consider real sites more secure than those who do not see the SI interface.*, we compare how participants responded to security questions about real sites in the normal SI condition with the control condition.

Participants in the SI condition did not express a greater inclination to log in to real sites ($M = 3.58, sd = 0.86$) than those in the control condition ($M = 3.9, sd = 0.75$). Participants in the SI condition also did not express a stronger belief that imposter sites were run by legitimate businesses ($M = 3.58, sd = 0.78$) than those in the control condition ($M = 3.8, sd = 0.7$).

Tests Involving Uglification Condition. Here, we present a set of t-tests that parallel those just presented, with the experimental SI + Uglification condition in place of the normal SI condition. The results are summarized in Table 4.

The first set of tests compare how SI+Uglification participants responded to security questions about real and imposter sites. In the SI+U condition, participants expressed a stronger inclination

to log in to real sites ($M = 4.55, sd = 0.33$) than imposter sites ($M = 2.72, sd = 1.56$). Participants in this condition also expressed a stronger belief that real sites ($M = 4.45, sd = 0.35$) were run by legitimate organizations than imposter sites ($M = 2.42, sd = 1.37$).

The second set of tests compare how SI+U participants responded to security questions about imposter sites with the control group. In the SI+U condition, participants did not express a lesser inclination to log in to imposter sites ($M = 2.72, sd = 1.56$) than those in the control condition ($M = 3.45, sd = 1.1$). However, participants in the SI+Uglification condition did express a weaker belief that imposter sites were run by legitimate businesses ($M = 2.42, sd = 1.37$) than those in the control condition ($M = 3.72, sd = 1.12$).

The third set of tests compare how SI+U participants responded to security questions about real sites with the control group. In the SI+U condition, participants expressed a greater inclination to log in to real sites ($M = 4.55, sd = 0.33$) than those in the control condition ($M = 3.9, sd = 0.75$). Participants in SI+U also expressed a stronger belief that imposter sites were run by legitimate businesses ($M = 4.45, sd = 0.35$) than those in the control condition ($M = 3.8, sd = 0.7$).

4.3.2 Non-Security Questions. Boxplots for the three non–security-questions are shown in Figure 12. The data for the Control and SI conditions look roughly equivalent in each plot.

The SI+Uglification condition deviates from this pattern. Participants in the SI+U condition rated the (uglified) imposter sites as less visually appealing than (non-uglified) real sites. They also rated (uglified) imposter sites less visually appealing than participants in the other conditions rated (non-uglified) imposter sites. SI+U participants' responses to all three questions about real sites also appear slightly higher than responses for real sites for participants in the non-uglified conditions.

4.3.3 Questions About Perceptions of SI's Effectiveness. Nineteen participants responded to the statement "The information provided by the site information tool was relevant to my login decision". One strongly disagreed, two responded "neutral", seven agreed, and nine strongly agreed.

Nineteen participants responded to the statement "I found the information in the site information tool easy to understand". Three disagreed, one responded "neutral", seven agreed, and eight strongly agreed.

4.4 Qualitative Results

Throughout the study, we collected qualitative data through our observations of participants as they interacted with the websites, as well as through informal dialogue with participants after each website and at the conclusion of the study.

Between websites, we would ask participants to verbally reflect on the overall experience, typically starting with the following prompt 'please give us your overall thoughts on the website, and provide a brief explanation of why you answered the questions the way you did'. Participants would generally begin by telling us the factors they considered when determining the website authenticity.

4.4.1 Determining Legitimacy. Content is a major factor in users' decisions of website legitimacy. When explaining the rationale for their decision, content was typically the first thing mentioned. Things like listing a contact number, showing products for sale, and having a familiar brand tended to make participants trust the website more.

Participants often '*poke and prod*'—interact with the site and compare the results with their expectations for legitimate sites. This included clicking on links and examining the page that opens up, or looking for pop-up elements when hovering over menus.

The website's URL was a non-factor in participant decisions: it was almost never mentioned, and only one participant understood how to parse the URL correctly.

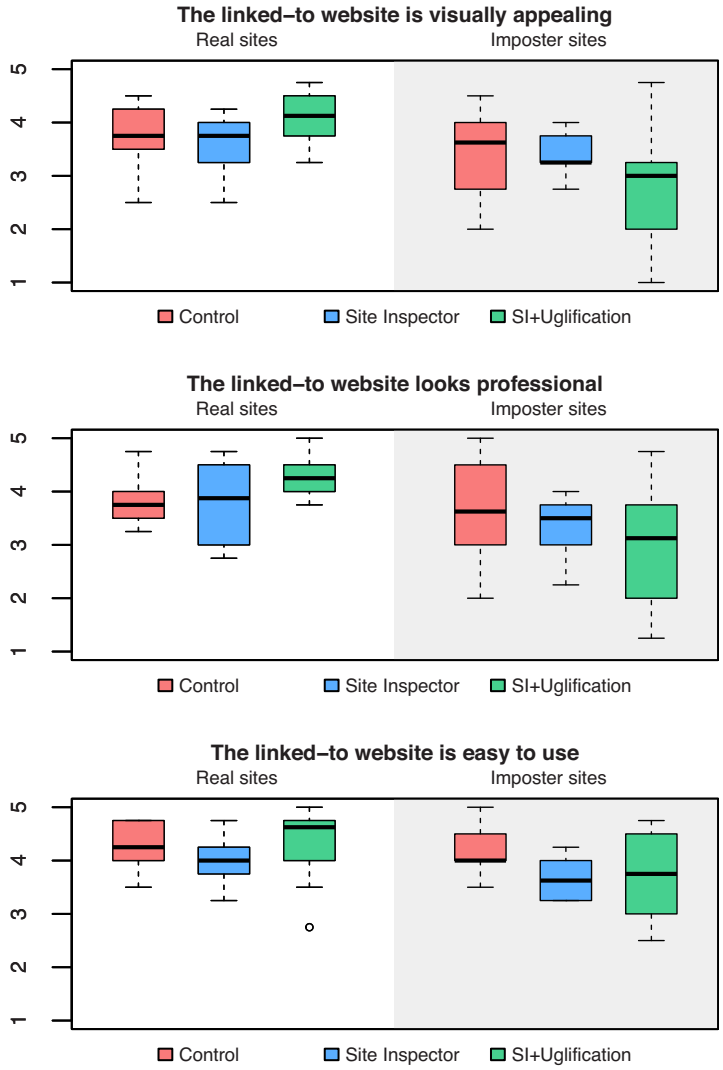

Fig. 12. Boxplots of questions relating to visual appeal, professionalism, and usability.

4.4.2 Observations on Site Inspector Usage. Owner identity information overtook *website content* as the most-often—and typically first-mentioned—reason for users judgments of website legitimacy. Content and aesthetics were still mentioned, but far less often.

SI did not appear to be overly intrusive. Participants would typically read all of SI's text once, and thereafter refer only to the basic level. Occasionally they might open one or more of the explanatory panels for a reminder. A minority of participants skimmed or skipped the explanatory text, and appeared willing to accept the basic level information without needing a deeper explanation.

A small minority of participants did not use information from SI in their judgments at all. Some of these expressed that a lack of a background check did not suggest anything dangerous. For example, the site owners may have forgotten to update the certificate, or perhaps they did not want to pay the fees.

4.4.3 Observations in the Uglification Condition. There appear to be individual differences with regard to aesthetic judgments. The 'uglified' websites were often not considered 'ugly' by our participants. This tended to happen when the sites had a simple color scheme and layout. For example, when the Tiffany & Co. website is inverted it could come across as a dark theme rather than an ugly look. Even when the sites had many colors and a busy layout, some participants did not find it ugly—much to our surprise.

The uglification effect appears to be amplified if the uglification is applied after the website fully loads (with its original color scheme) and then a brief delay, as is what happens with some larger websites.

4.5 User Study Discussion

4.5.1 Summary of Hypothesis Tests. We found preliminary support for each of our first two hypotheses. In support of H1: *Participants who see the SI interface will consider imposter sites less secure than real sites*, we found that participants in the normal SI condition (i.e., without uglification) reported being less likely to log in to fake sites than real sites, and they were also less likely to think that real sites were run by legitimate businesses. Participants who saw SI were able to distinguish real sites from imposter sites. The same could not be said for those who did not see SI.

In support of H2: *Participants who see the SI interface will consider imposter sites less secure than those who do not see the SI interface.*, we found that participants in the normal SI condition were less likely to log in to fake sites than those in the control condition, and they were less likely to think that imposter sites were run by legitimate businesses. Participants who saw SI were better at classifying imposter sites as imposters than those who did not see SI.

We did *not* find support for H3: *Participants who see the SI interface will consider real sites more secure than those who do not see the SI interface.*. Participants in the normal SI condition were not more likely to log in to real sites than those who did not see SI, nor were they more likely to think that real sites were run by legitimate businesses. Participants who saw SI were not better at classifying real sites as real than those who did not see SI.

4.5.2 Reconsidering Security Indicators. As mentioned in the introduction and Section 2.4, browsers have removed all identity indicators without ever having explained them to users. One objective of the present work was to see if users could be taught the meaning of security indicators for identity and confidentiality during normal web usage.

Our user study provides preliminary evidence that this is indeed possible. When each session began, we directed users' attention toward the new UI element via a verbal prompt. Users typically explored the explanatory content once, and then subsequently glanced at the indicators when making trust decisions. Users who saw the indicators in SI were better able to detect fraudulent websites than those who did not.

We believe similar results could be achieved by integrated security indicators directly into to the browser's UI. The current 'basic status view' is too large to fit in the browser Chrome, but the interface could be further reduced to just the identity and confidentiality icons, which could be put in place of the current 'lock' symbol next to the URL. Displaying the verified owner name in a manner similar to the legacy EV indicator might also be helpful. When the user clicks these, a window like the current SI interface could be shown, with hierarchically-organized explanatory texts. We believe the user should be prompted to this new UI feature when it is first installed, perhaps through a minimally obtrusive 'demo' that directs users' attention to it and notifies them of its basic purpose.

As mentioned in the introduction, we acknowledge there are operational issues surrounding identity verification, including those mentioned in Section 4.7.1, which undermines the

effectiveness of OV and EV certificates. A proper comprehensive solution to this problem must involve addressing these issues in addition to any improved UI features.

4.5.3 Skimming/Skipping and the Abstraction Hierarchy. One qualitative finding was that a significant minority of participants who saw SI either skimmed through or skipped the explanatory text entirely. Yet the majority of these participants still seemed to incorporate the information in SI in their decision-making. For those who did read the explanatory text, it was rarely revisited, and they relied almost exclusively on the information provided in the basic status view.

This skimming/skipping behavior is actually intended, and is made possible through SI's use of the abstraction hierarchy (borrowed from EID [23, 32]). We believe there are many users who want reliable information about a website's security but are currently under-served. We also recognize there is likely significant variability in how much attention users are willing to devote to this information. Some users might already have the necessary knowledge, perhaps from previous use of SI, while others may have low motivation to learn more and prefer to accept the basic status information at face value, provided it comes from a trustworthy source like the browser.

Our goal is not to turn users into security experts or require users to spend a significant amount of time thinking about security. Instead, our goal is to provide those who seek insight into a website's security status with the means to do so, requiring minimal expenditure of attentional resources. The abstraction hierarchy used by SI enables this goal by making information available on demand, eliding unwanted information.

4.5.4 Uglification Results.

Variability in Effectiveness. Participant responses in the SI+U condition were somewhat surprising. The manipulation of visual appeal should provide an additional, visceral signal of website security or insecurity in addition to those given by SI; this is what we found in prior work [29]. If uglification provides another source of help to users, one would expect participants in the SI+U condition to be a little better at identifying imposter sites than those in the normal SI condition. However, this was not the case: for both security questions about imposter sites, the central tendency for the two conditions were approximately the same.

There is much more variability in the SI+U condition for security questions about imposter sites (see Figure 11). This is consistent with our observation mentioned in Section 4.4 that the 'uglified' sites were often not considered ugly by our participants. Some of this we attribute to difference in aesthetic tastes between individuals, and perhaps cultures as well: 'ugliness' is not a purely objective phenomenon. Some may even be due to differences in perceptual ability: some of our participants may have been color blind, for example.

We think factors relating to the source websites also played a role. The method of uglification we used also worked better on some websites than others. Ratings of visual appeal for the Tiffany website (shown in Figure 21(a)), for example, were about the same whether it was 'uglified' or not (see Figure 26). The uglified Cabela's website (Figure 23) was judged most ugly relative to its non-uglified variants, perhaps due to its choice of colors and use of many images.

We think it would be worthwhile to explore additional CSS changes to the uglification method to help increase the perceived ugliness of uglified websites, and to help ensure the websites are considered ugly by all users.

Tests Paralleling Third Hypothesis. While participants in the normal SI condition were not better able to classify real sites as real compared with those in the control condition, participants in the SI+U condition *were* (see Table 4 and Figure 11). There is a remarkably *low* amount of variability in SI+U participant responses to security questions about real sites—the exact opposite of what was

found for their responses about imposter sites—with SI+U participants being very likely to log in to real sites, and very likely to say real sites are run by legitimate businesses.

The most obvious explanation for the high security ratings for real sites relative to the other conditions is that uglification made it more obvious which were the imposter sites, which had the side effect of reducing uncertainty about the legitimacy of real sites. But the high variability of SI+U security judgments about fake sites, along with the high variability of SI+U judgments about the visual appeal of uglified sites, and our qualitative observations of participants, contradict this explanation.

This matter needs more research.

4.6 User Interface vs. Infrastructure Issues in OV and EV Certificates

The effectiveness of web certificates in helping users understand website identity hinges on addressing two distinct categories of challenges: (1) the adoption and credibility of web certificates and the identity verification processes of CAs, and (2) and the usability and comprehendability of interfaces to web certificates. These categories are fundamentally different, each requiring different kinds of solutions and potentially the expertise of different types of specialists.

The present research pertains to the second of these categories. Our goal is to make progress toward showing that identity information from X.509 certificates can be communicated to end-users in a pseudo-realistic setting. The infrastructural issues surrounding OV and EV certificates, as discussed in Section 4.7.1, pertains to the first category. Our study sidesteps the infrastructural issues because they would act as confounding factors in our evaluation of SI as a communication tool.

We are optimistic that these infrastructural issues can be fixed with the commitment of relevant stakeholders, including CAs and perhaps regulatory bodies. Carefully considering how to address these issues would, we think, be a valuable research endeavor. Moreover, advancements in the user interface aspects of the problem could provide CAs or regulators the motivation to act.

4.7 Limitations

4.7.1 Idealized Assumptions. Our study adopts several idealized assumptions regarding the implementation of owner verification certificates, deliberately avoiding the existing issues in the current system. This approach allows us to focus on assessing whether SI can effectively communicate identity and other information from X.509 certificates to end-users in realistic usage scenarios. Striving for complete realism in our methodology would undermine our research goals, as these systemic issues would act as confounds. We discuss this more in Section 4.6.

Idealization 1: All and Only Legitimate Sites use OV or EV Certificates. Our study assumes that only legitimate websites have OV certificates or higher, and only fraudulent sites use DV certs or lower. This is clearly not the case in the real world, where the use of DV certificates by legitimate websites is widespread: even some major websites like *google.com* use DV certificates. If users commonly see popular, trusted websites being labeled as 'unverified' by the UI, the label's effectiveness in signaling caution to users would be significantly undermined.

Our study also assumes that verified identities are always non-malicious. However, in 2018 a researcher received an EV certificate for "Stripe, Inc." [7], a company with the same name as the popular payment processor but registered in a different state. This highlights the imperfections in the EV process.

Idealization 2: The Verification Process used by CAs is known and Reliable. The explanatory text in SI implies that the process CA's use to verify organizations is known to us. In reality, however, we only know what CAs *should* do: for example, the CA/Browser provides guidelines for issuing

and managing EV certificates [9]. While we assume that CAs follow the relevant guidelines for issuing and managing certificates, it is unclear what their procedures actually are. It is therefore unclear what value this information has for users in their trust decisions.

Idealization 3: The Verified Organization in the X.509 Certificate Controls the Website. Our study also assumes that the entity listed under `Organization` in the `Subject Name` field of the X.509 certificate is the entity that has direct control over the website. Unfortunately, this is often not the case. For example, in our search for websites to use in our study, we noticed that Cloudflare, Inc. was frequently listed as the 'organization' on OV certificates. In a real-world application of SI this would be confusing to users.

4.7.2 Study Method and Threats to Validity. Our main contribution is the principled design of SI, with the study intended as an initial proof of concept, and to inform iterative design improvements. We acknowledge several threats to validity in the present study that should be amended in more confirmatory studies.

Population Validity. We recruited for the study using a Facebook group for recruiting HCI participants. It is possible that this group is not fully representative of the target population of end-users (e.g., they may be more educated). We note that this issue is common in studies involving human participants.

Ecological Validity. Our study takes place in a controlled setting where the researcher observes the participants as they complete their tasks. It is possible that participants interactions with SI were influenced by the knowledge that they were being observed by the researcher, otherwise known as the Hawthorne Effect [18]. The goal of this study was not to be fully ecologically valid, but to better understand how the tool is used and understood. Once the tool has reached maturity, confirmatory studies striving for maximal ecological validity should be conducted to minimize this risk. Ultimately, the tool could be integrated into web browsers and deployed to real users.

Small Sample Size. Our study includes 30 participant, and 10 per condition. We believed this number was suitable for an initial examination of the tool where one researcher observed all participants as they completed their tasks. Future studies of more refined iterations of the tool should use a larger participants pool.

4.7.3 Influence of Known Brands. The real websites and the source websites for our imposter sites were chosen because we did not think users would be familiar with the *websites*. Excepting a few cases, it seemed that were indeed not familiar with the websites in the study. However, it was common for a participant to be familiar with at least one of the *brands* associated with the websites (particularly T-Mobile and Tiffany & Co.), which seemed to have a subtle but real influence on their responses (e.g., If they knew the brand they seemed to trust the site more). An effective strategy for a follow-up study might be to focus on businesses with a regional as opposed to national or international reach.

4.7.4 Methodological Adjustments. As mentioned in Section 4.1.3, we made two methodological adjustments mid-study to address errors we identified. These adjustments helped us better meet our main goals of evaluating SI and determining whether more rigorous studies would be warranted. It is important to note that the decision to retain the study data prior to the adjustments (as opposed to destroying this data and recruiting new participants) actually reduces the measured effectiveness of SI, so we do not consider it a threat to validity.

5 Conclusions

Currently, browser UIs lack indicators of website identity assurances provided by X.509 certificates, leaving users vulnerable to various attacks involving fraudulent websites. Browsers formerly featured EV indicators, but they were never explained to users, making them useless, and they were subsequently removed in 2019. We believe the industry has abandoned identity indicators prematurely.

We designed and implemented SI in part to explore this possibility. SI implements the SAVE paradigm [25] (Section 2.2), which offers guidelines for creating UIs that users develop security mental models without unduly interfering with their everyday tasks.

SI features a basic view showing the website's domain along with two security statuses: identity and confidentiality. It includes text explaining the meaning of these indicators, organized within abstraction hierarchy, which users can reveal by clicking buttons. In accordance with SAVE, SI is highly visible, uses familiar language, respects users' limited attentional resources, and directly describes real-world phenomena.

We conducted a user study to provide a preliminary evaluation of SI—to see if more rigorous evaluations would be worthwhile, to inform design improvements, and to explore the effectiveness of SAVE. Thirty participants were placed into one of three conditions: 1. SI is inactive (control condition); 2. SI is active; 3. SI is active, plus websites with DV/no certificates are 'uglified'. Participants viewed eight websites (four real, and four imposters created by us) and answered Likert scale questions. We were most interested in seeing if the SI interface helped users distinguish real sites from imposter sites.

We found that SI did in fact help users distinguish real sites from imposter sites (H1), which appears to be mostly driven by an improvement in the ability to detect fraudulent sites (H2), and not by an improved ability to detect real sites (H3). Most participants found the interface useful in their decision whether to log in to the site, and they found the information in the interface intelligible. SI seemed to help users overcome a strong bias to use a site's appearance and content when judging its legitimacy.

The uglification technique was an experimental attempt at applying findings from previous work [29], which showed that users make security decisions based on a website's visual appeal. It seems that the technique we used worked for some users (who found the sites ugly and subsequently rated the sites as less secure than those in the normal SI condition) better than others (who did not find the uglified sites ugly and who did not rate them as less secure). Curiously, participants in the uglification condition did seem to have an improved ability to detect real sites (H3).

Limitations of the study include a relatively small sample size (30 in total; 10 per condition), using websites with recognized brands (which appeared to increase trust), and minor methodological adjustments made early in the study.

If browsers adopted an approach similar to SI, the usefulness of OV and EV certificates for ordinary users could potentially be reclaimed. Along with improvements to the validation process, this would allow users to better protect themselves against attacks involving fraudulent websites.

References

[1] Ruba Abu-Salma, Angela M. Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina, and Matthew Smith. 2017. Obstacles to the adoption of secure communication tools. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP'17)*. IEEE, New York, NY, USA, 137–153. DOI: https://doi.org/10.1109/SP.2017.65

[2] Anti-Phishing Working Group. 2023. APWG Phishing Activity Trends Report. (11 2023). Retrieved August 18, 2024 from https://docs.apwg.org/reports/apwg_trends_report_q4_2023.pdf

[3] Farzaneh Asgharpour, Debin Liu, and L. Jean Camp. 2007. Mental models of security risks. In *Proceedings of the Financial Cryptography and Data Security (FC'07)*, Sven Dietrich and Rachna Dhamija (Eds.). Springer, Berlin, Germany, 367–377.

[4] Lawrence W. Barsalou. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22, 4 (1999), 577–660. DOI: https://doi.org/10.1017/S0140525X99002149
[5] Simon Bell and Peter Komisarczuk. 2020. An analysis of phishing blacklists: Google safe browsing, OpenPhish, and PhishTank. In *Proceedings of the Australasian Computer Science Week Multiconference*. ACM, New York, NY, USA, 1–11.
[6] Robert Biddle, P. C. van Oorschot, Andrew S. Patrick, Jennifer Sobey, and Tara Whalen. 2009. Browser interfaces and extended validation SSL certificates: An empirical study.
[7] Bramus. 2018. Extended Validation Is Broken. (2 2018). Retrieved August 20, 2024 from https://www.bram.us/2018/02/05/extended-validation-is-broken/
[8] Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, and Saranga Komanduri. 2011. Bridging the gap in computer security warnings: A mental model approach. In *Proceedings of the IEEE Symposium on Security and Privacy (SP'11)*. IEEE, New York, NY, USA, 18–26. DOI: https://doi.org/10.1109/MSP.2010.198
[9] CA/Browser Forum. 2019. Guidelines For The Issuance And Management Of Extended Validation Certificates. (2019). Retrieved March 9, 2023 from https://cabforum.org/wp-content/uploads/CA-Browser-Forum-EV-Guidelines-v1.7.0.pdf
[10] L. Jean Camp. 2009. Mental models of privacy and security. *IEEE Technology and Society Magazine* 28, 3 (2009), 37–46. DOI: https://doi.org/10.1109/MTS.2009.934142
[11] D. Cooper, S. Santesson, S. Farell, S. Boeyen, R. Housley, and W. Polk. 2008. *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. RFC 5280. RFC Editor. Retrieved from https://www.rfc-editor.org/info/rfc5280
[12] Kenneth Craik. 1943. *The Nature of Explanation*. Cambridge University Press, Cambridge, United Kingdom.
[13] GNU Wget2. 2023. Wget2 (Version 2.0.1) [Computer software]. GitLab. https://gitlab.com/gnuwget/wget2
[14] Bruce Heiding, Fredrik Schneier and Arun Vishwanath. 2024. AI Will Increase the Quantity—and Quality—of Phishing Scams. (5 2024). Retrieved August 18, 2024 from https://hbr.org/2024/05/ai-will-increase-the-quantity-and-quality-of-phishing-scams
[15] Philip N. Johnson-Laird. 1983. *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Harvard University Press, Cambridge, MA, USA.
[16] Predrag Klasnja, Sunny Consolvo, Jaeyeon Jung, Benjamin M. Greenstein, Louis LeGrand, Pauline Powledge, and David Wetherall. 2009. When I am on Wi-Fi, I am fearless: Privacy concerns & practices in everyday Wi-Fi use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'09)*. ACM, New York, NY, USA, 1993–2002.
[17] LimeSurvey. LimeSurvey Homepage. (n.d.). Retrieved February 4, 2023 from https://www.limesurvey.org/
[18] Elton Mayo. 1933. *The Human Problems of an Industrial Civilization*. The MacMillan Company.
[19] Mozilla Corportation. webRequest.getSecurityInfo(). (n.d.). Retrieved January 9, 2023 from https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions/API/webRequest/getSecurityInfo
[20] Don Norman. 2013. *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books, Inc., Hachette, NY, USA.
[21] Donald A. Norman. 1986. Cognitive engineering. In *Proceedings of the User Centered System Design: New Perspectives on Human-computer Interaction*, Donald A. Norman and Stephen W. Draper (Eds.). CRC Press, Boca Raton, FL, USA, 266–290.
[22] Donald A. Norman. 2014. Some observations on mental models. In *Proceedings of the Mental Models*. Psychology Press, 15–22.
[23] Jens Rasmussen. 1985. The role of hierarchical knowledge representation in decisionmaking and system management. *IEEE Transactions on Systems, Man, and Cybernetics* 15, 2 (1985), 234–243.
[24] Stuart E. Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. 2007. The Emperor's new security indicators. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy (SP'07)*. IEEE, New York, NY, USA, 51–65.
[25] Eric Spero and Robert Biddle. 2020. Out of sight, out of mind: UI design and the inhibition of mental models of security.
[26] Milica Stojmenović, John Grundy, Vivienne Farrell, Robert Biddle, and Leonard Hoon. 2016. Does textual word-of-mouth affect look and feel?. In *Proceedings of the 28th Australian Conference on Computer–Human Interaction*. 257–265.
[27] Milica Stojmenović, Temitayo Oyelowo, Alisa Tkaczyk, and Robert Biddle. 2018. Building website certificate mental models. In *Proceedings of the International Conference on Persuasive Technology*. Springer, 242–254.
[28] Milica Stojmenović, Eric Spero, Temitayo Oyelowo, and Robert Biddle. 2019. Website identity notification: Testing the simplest thing that could possibly work. In *Proceedings of the 17th Annual Conference on Privacy, Security and Trust (PST'19)*. IEEE, New York, NY, USA, 310–316. DOI: https://doi.org/10.1109/PST47121.2019.8949048
[29] Milica Stojmenović, Eric Spero, Miloš Stojmenović, and Robert Biddle. 2022. What is beautiful is secure. *ACM Transactions on Privacy and Security* 25, 4, Article 30 (7 2022), 30 pages. DOI: https://doi.org/10.1145/3533047

[30] Christopher Thompson, Martin Shelton, Emily Stark, Maximilian Walker, Emily Schechter, and Adrienne Porter Felt. 2019. The web's identity crisis: Understanding the effectiveness of website identity indicators. In *Proceedings of the 28th USENIX Security Symposium (USENIX Security'19)*. USENIX Association, Berkeley, CA, USA, 1715–1732.
[31] Richard F. Thompson. 2009. Habituation: A history. *Neurobiology of Learning and Memory* 92, 2 (2009), 127–134.
[32] Kim J. Vicente and Jens Rasmussen. 1992. Ecological interface design: Theoretical foundations. *IEEE Transactions on Systems, Man, and Cybernetics* 22, 4 (7 1992), 589–606. DOI: https://doi.org/10.1109/21.156574
[33] Rick Wash. 2010. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*.
[34] Min Wu, Robert C. Miller, and Simson L. Garfinkel. 2006. Do security toolbars actually prevent phishing attacks?

Appendices

A Site Inspector: No Certificate

The following figures show SI when the user has visited a domain without an X.509 certificate. In this case the user has never visited this domain before.

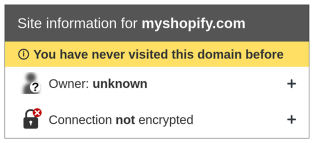

Fig. 13. Site Inspector Level 1: basic status view for `myshopify.com`.

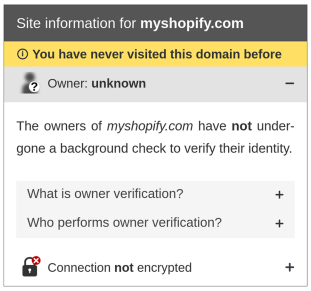

Fig. 14. Site Inspector Level 2: status elaboration view for `myshopify.com`.

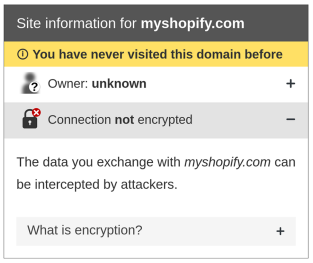

Fig. 15. Site Inspector Level 3: mechanism/process description view for `myshopify.com`.

B Websites Used in Study

B.1 Real Sites

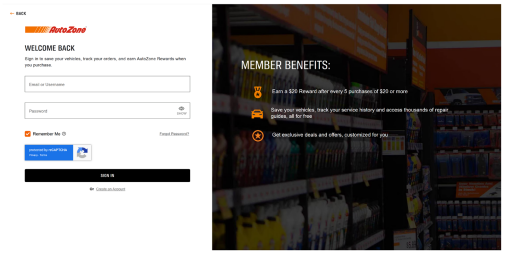

Fig. 16. Autozone: `www.autozone.com/signin`.

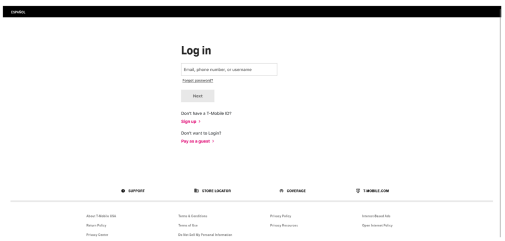

Fig. 17. T-Mobile: `account.t-mobile.com/signin/v2`.

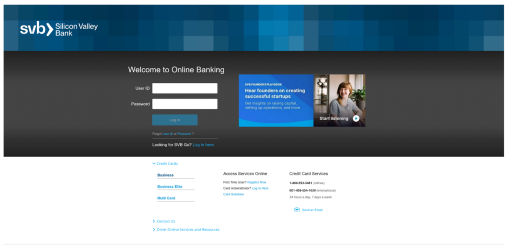

Fig. 18. Silicon Valley Bank: `www.svbconnect.com/auth`.

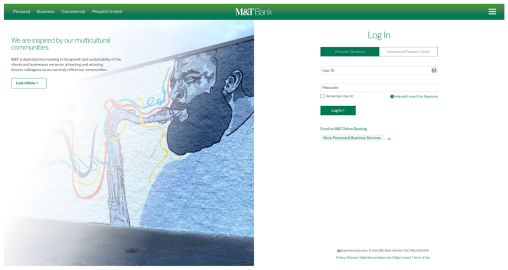

Fig. 19. M&T Bank: `www3.mtb.com/log-in`.

B.2 Imposter Sites

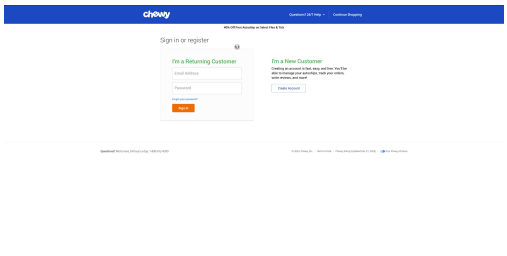

(a) Normal

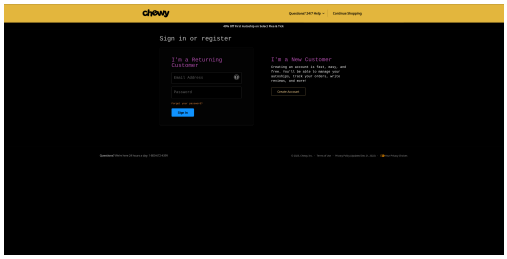

(b) Uglified

Fig. 20. Chewy: `chewy.http-s.net`.

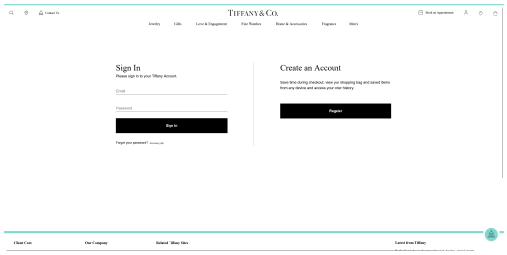

(a) Normal

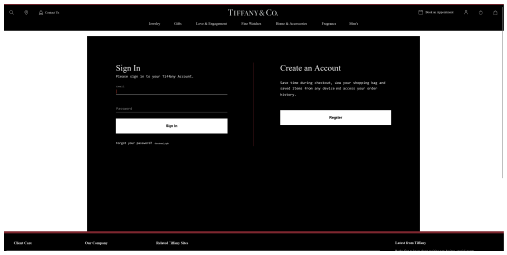

(b) Uglified

Fig. 21. Tiffany & Co.: `tiffany.sec-www.com`.

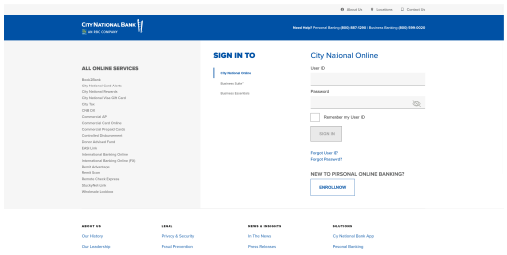

(a) Normal

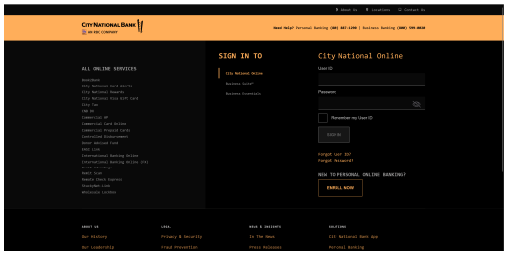

(b) Uglified

Fig. 22. City National Bank: `cnb.secureuserlogin.com`.

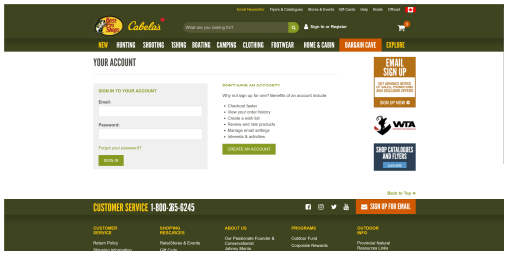

(a) Normal

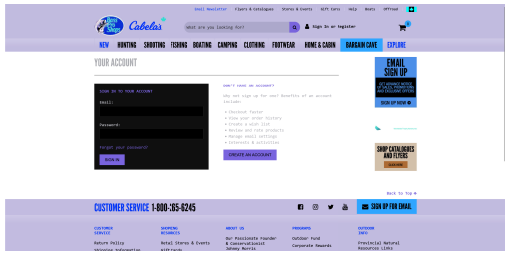

(b) Uglified

Fig. 23. Cabela's: `cabelas.certlo.com`.

C Per-Website Boxplots

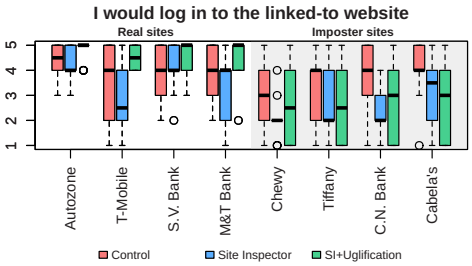

Fig. 24. Individual websites boxplots: 'Login' question.

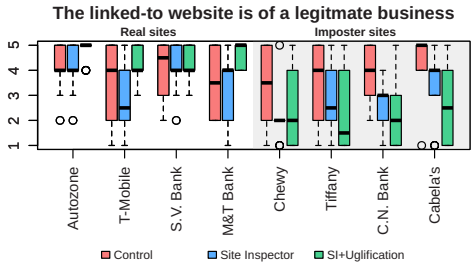

Fig. 25. Individual websites boxplots: 'Legitimate' question.

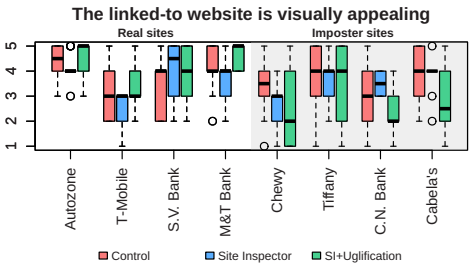

Fig. 26. Individual websites boxplots: 'Visual appeal' question.

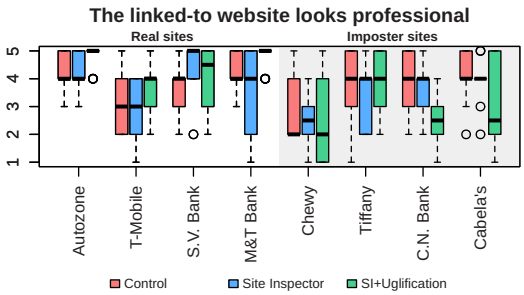

Fig. 27. Individual websites boxplots: 'Professionalism' question.

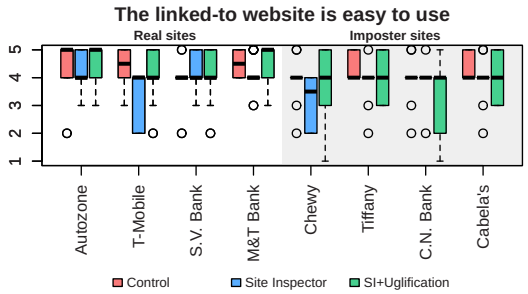

Fig. 28. Individual websites boxplots: 'Easy to use' question.

Received 2 September 2024; revised 23 December 2024; accepted 23 March 2025