Gathering Evidence: Use of Visual Security Cues in Web Browsers

Tara Whalen Kori M. Inkpen

Faculty Of Computer Science
Dalhousie University
Halifax, NS, Canada, B3H 1W5
{whalen, inkpen}@cs.dal.ca

Abstract

Web browsers support secure online transactions, and provide visual feedback mechanisms to inform the user about security. These mechanisms have had little evaluation to determine how easily they are noticed and how effectively they are used. This paper describes a preliminary study conducted to determine which elements are noted, which are ignored, and how easily they are found. We collected eyetracker data to study user's attention to browser security, and gathered additional subjective data through questionnaires. Our results demonstrated that while the lock icon is commonly viewed, its interactive capability is essentially ignored. We also found that certificate information is rarely used, and that people stop looking for security information after they have signed into a site. These initial results provide insights into how browser security cues might be improved.

Key words: usable security, web browsing, visual feedback, secure web-based transactions.

1 Introduction

Web tasks requiring security are common. They include accessing email, shopping, and banking. Browsers provide feedback about security to the user, to reassure them that a secure connection has been established. For example, many browsers display a lock icon when Secure HTTP (`https`) is in place, and remove it when the connection is no longer encrypted. Although visual feedback is in widespread use, there has been little evaluation of how effectively it works.

The challenge of aligning usability and security has been a subject of recent research in both the HCI and security communities. This work highlights the need to provide appropriate feedback about security, so that users can make informed decisions about their sensitive data. Given the potential consequences of exposing banking passwords and credit cards, users are understandably concerned about the risks of online transactions. People must be given the ability to discover and understand security information when using the web. The overall goal of this research is to develop feedback that clearly informs users about security without overburdening them with distractions. The first step towards this goal is to understand how people currently use browser security information.

We have completed a preliminary investigation into how people use visual cues currently provided in browsers. Eyetracking and questionnaire data were collected and analyzed to ascertain whether people noticed visual security indicators and whether they were able to find security information. The results of this analysis identified areas requiring design improvements, in order to better support awareness of web security.

In presenting our work, we first discuss related research in the area of usable security, and then describe our study design and methodology. Finally, we present our results and discussion, including future research directions.

2 Related Work

The ideas underlying usable security were first discussed in 1975 by Saltzer and Schroeder, who included "psychological acceptability" in their list of essential principles for information protection systems [7]. This topic went dormant for decades, during which time HCI evaluation techniques evolved but were not often applied to security applications. The field was revitalized in 1999, with Whitten and Tygar's paper [9] about the usability problems of PGP 5.0, a (supposedly) easy-to-use encrypted email program. Since that time, research has been carried out on such diverse topics as password usability [1] and network security visualization [10].

Security offers difficult challenges to the usability community [9]. It is often complex: hard to understand and easy to misconfigure. In addition, users are not motivated to spend time and energy on security when completing tasks. The human-factors aspects of security are the subject of much recent interest, spawning workshops [3] and special publications [6].

To date, there has been little work specifically in the area of usable browser security. One of the few pieces of work in this area is Friedman et al.'s research on value-sensitive design in browsers [5]. In 2002, they conducted semi-structured interviews about browser security with 72 people from three diverse communities (rural, suburban professional, and high-tech). They found that around half of the respondents, across all communities, could not recognize a secure connection.

In addition, they created a list of the types of evidence people use in identifying a secure connection, such as the lock icon or the type of website. They used this information to develop an approach for handling cookies [4], but did not continue work on the recognition aspects of web security.

3 Security Cues Study

The goal of this study was to gather details about people's attention to visual security details, and to gather subjective data about how people interpret this information. We investigated which security cues were ignored, which were easily noticed, and which were hard to find. This phase of our research was designed to provide the groundwork for explorations into improved designs for expressing browser security information.

3.1 Study Design

An eyetracker was used to reveal gaze and fixation on aspects of the screen. Eyetracking data was gathered in two different phases: first, when the user was conducting secure transactions as part of normal browsing, and second, when the user was asked to look specifically for security information when performing a task. Both phases took place within a one hour-long session.

After each phase, participants completed a questionnaire that asked whether participants checked for security on a subset of the pages they viewed (those that involved a confidential task, such as logging into a bank site). If they stated they *had* checked for security information, they were asked whether they concluded that the page was secure or insecure. If they *hadn't* checked the security information they were asked why not. The participants were also asked what sort of evidence they use when determining whether a web page or web connection is secure. A list of possible sources was provided based on items gathered from Friedman et al.'s study [5]: `https` in location bar/URL of web page; lock or key icon displayed by browser; certificate information; type of site (e.g., bank site is expected to have security); type of information being submitted (e.g., expect passwords to be treated securely); and statements on page that say the site is secure. We also provided space to list additional sources.

Participants were presented with this list of evidence sources twice: at the end of Phase 1, participants were asked what evidence they use *in general* when evaluating security, and again at the end of the study, when they were asked what evidence they had *actually* used when carrying out the tasks in Phase 2. Participants could select as many list items as they wished. In addition, the eyetracking data were used to verify whether people's reported behaviour matched their actual behaviour.

Phase 1: Normal Browsing

For the first half of the study, we wanted to observe how users attend to, and interact with, visual security cues when they are browsing naturally. Ideally, we would have observed participants in a non-laboratory setting. However, that ideal environment would not have enabled us to use the eyetracker to gather data about attention focus. More importantly, it is highly invasive for an observer to see details of private, secure transactions. This presents a challenge: providing participants with confidential data that is not their own and motivating them to protect it without unnaturally drawing their attention **too** strongly towards security cues.

We attempted to address this problem through a carefully-designed script. At the beginning of the session, participants were told that we were studying how people looked at browsers and web pages; security was not explicitly mentioned. They were told to browse as they would at home or at work, and that there would be a number of tasks that included reading news, email, and banking. For tasks that required logging in to accounts and using credit cards, they were told that they would use accounts and a card belonging to the lab. They were explicitly asked to treat this information "as if it were their own," and were instructed to keep it confidential. We hoped that these instructions would make the study data (passwords and credit cards) appear valuable, and thus the participants might make some effort to protect the information. At the same time, we felt that these instructions were not so strong as to make people overly conscious about security. (Unfortunately, as described in our results, this attempt to simulate normal browsing was not successful.)

A short set of tasks was created, both with and without secure transactions. Websites both familiar and unfamiliar to our participants were chosen, with the expectation that their attention to security cues may vary depending on their past experience with a particular website. This phase had five tasks in all: three secure tasks interleaved with two non-secure tasks. The first task was to use web-based email (Gmail), followed by a second task of reading a news article. The third task was to log into their bank website using a lab account rather than their own. The fourth task was to perform a simple Google search, and the last task was to purchase a cable from a small obscure retailer, using a research account credit card.

After completion of the five tasks, participants completed a questionnaire. This questionnaire focused on only those pages that had security tasks associated with them (such as webmail login page) and inquired whether the users checked for security and what evidence was used in making decisions about security on

those pages. They also told us what security evidence they used in general.

Phase 2: Focus on Security

For the second half of the study, we wanted to observe how people used security cues when they were **explicitly** instructed to pay close attention to security. In this case, we were interested in how well security cues worked when they were deliberately being sought, in effect providing a "lower bound": if experienced participants who were motivated to seek out security information could not easily find it, then we could conclude that the existing cues are ineffective.

We told participants that they would perform similar tasks, but this time they should be extra-cautious, taking the time to look for information about security as they completed the tasks. They were told to take whatever actions they required in order to decide whether it was safe to complete the task. We informed them that if they concluded that a page was insecure, they should report that fact on the questionnaire, but they should still proceed to enter the confidential data.

Participants completed three tasks in this phase, all of which involved entering confidential data. The first task was to access an online banking account from a bank website that they had never used. The second task was to access email through Microsoft's Hotmail web service. The last task was to buy a book from a popular online bookstore, using the research lab's credit card.

After completion of the three tasks, participants completed a questionnaire similar to the one at the end of Phase 1, which inquired whether they checked for security on each page, and what evidence they used in making decisions about security on those pages.

3.2 Study Infrastructure

Gathering eyetracking data and building the web infrastructure required a careful technical arrangement. The experimental setup consisted of three main components: the eyetracker, the web client, and the web server with study-specific web content. These components simulated normal web browsing, including secure websites, without requiring private online accounts.

A Tobii x50 eyetracker was used in this study, a small device which sat directly beneath the monitor used for web browsing. This eyetracker does not require the user to wear any special equipment; they need only stay within range of the device. Users sat approximately 50 cm away from the eyetracker; at this distance, the average error is approximately 1 cm between the intended and actual gaze point [8]. During a session, all screen activity was captured for analysis, including cursor movements and scrolling. We performed two calibrations for each user, one before each of the two phases.

A snapshot from the Tobii analysis software is shown in Figure 1. This still was taken from the video playback of the captured gaze data, at a moment when the user was looking at the lock icon. Note the bottom right-hand corner: the line shows the path of recent eye motion, while the circle represents fixation on one spot.

The web client ran on a PC with Windows XP. All participants used Internet Explorer (IE). The study was limited to one browser to ensure that the types of visible security cues were identical for all users, and IE was chosen because it is the most popular web browser. The browser security settings were configured specifically for this study. In particular, pop-up alert messages were turned off. Although browser alert messages are a significant aspect of a user's security experience, we felt that alerts were a complex issue best left for a future study. We also imported the proxy server's root certificate as a "trusted root certificate", which would prevent any messages about untrusted certificates. Finally, to make navigating consistent, all of the websites required for the study were placed into the IE Favorites list.

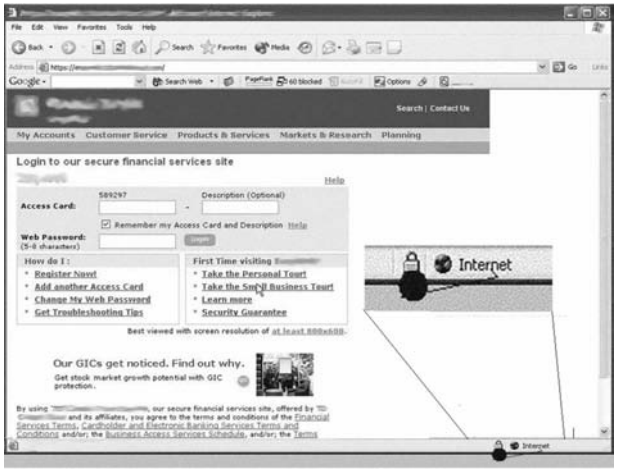

Figure 1: A still frame taken from video playback of eyetracking data. The enlargement (added for this figure) shows the user's gaze in the lock area; the dot shows fixation at that spot.

3.3 Simulating Secure Transactions

The browser in our study was configured to send requests through a lab proxy server. Users were able to access the web normally, with the exception of those pages requiring a real bank account or submission of a real credit card. For these types of connections, a "shadow" web was created, composed of pages mirrored from the real websites and stored on the local server. The pages were changed to allow logins with our lab usernames and passwords, and banking pages

were created to look exactly like real account pages, but with dummy financial information. These pages resided on our web server, and the proxy server returned these pages rather than the real sites. Participants were not made aware of the proxy server.

All the secure sites for our study (running on our local web server) were given domain names to match the real sites. IP-based virtual hosting in Apache 2.0 was used for this, with Secure Sockets Layer (SSL) support for all the secure sites. This approach allowed us to show `https` in the Location Bar, with the correct domain names, and to have the lock icon appear in IE.

For example, when a participant requested the Scotiabank Online Banking Login page, the page request would first go to the proxy server. The server would determine that the Scotiabank site was hosted locally, and would return our specialized content. If a request for any "regular" page was sent (such as Google), the server would determine that the site was hosted remotely, and the real page would be returned. The only secure connections that went to the actual web were those to web-based email sites, for which we created real email accounts (under the lab's name).

Our shadowed sites were very similar to the real sites. A few links were removed due to the difficulty of mirroring all content stored on secure and/or complex sites. We used self-signed (unvalidated) certificates for all the secure sites; this was sufficient for providing the appearance of security for most users, but those with knowledge of certificates could determine that they were not connected to the real institution's website.

3.4 Participants

Participants were recruited from within our local university, including faculty, staff, and students. Email lists were used as the primary method of solicitation; participants were offered $10 in compensation. A prescreening questionnaire was used to select participants who were experienced web users: all subjects used the web daily, and had used the web for at least 5 years. As well, only people who were familiar with IE and who had carried out banking, email, and shopping tasks on the web in the past six months were selected. Experienced web users were selected into order to determine whether browser security cues were useful for *any* population. Security cues are intended to provide information to those who carry out secure transactions; if this target group has difficulties with them, this would demonstrate a severe lack of effectiveness with the current approach.

Sixteen participants (10 female and 6 male) took part in the study. Nine participants worked for the university (faculty or staff), and seven were students. Based on background information provided by the participants, it was estimated that just over half (nine) of the participants had a technical background.

4 Results

Although this was only a preliminary study, our early results yielded several interesting observations, both in terms of current usage and potential browser improvements. Due to experimental difficulties with Phase 1, the eyetracking observations were drawn only from Phase 2. The only data used from Phase 1 was questionnaire data on the types of security evidence used in general.

4.1 Phase 1 Results: Normal Browsing

Although the first phase was designed to simulate natural web browsing, we found no evidence that security was checked at all, for any page. Based on our observations and questionnaire comments, we concluded that we were unsuccessful in reproducing normal browsing behaviour during our study. One reason participants did not check whether a page was secure was that it was not their own data and thus took no care to protect the information. Despite being asked to treat the data as if it was their own, most participants were completely unmotivated to take any precautions. (Some participants concluded the data could not be genuinely confidential, or we would not have let them use it.)

Another reason participants did not check whether a page was secure was the distraction of the testing environment. Carrying out a series of tasks in a lab setting tends to make people focus on following instructions; and leaves little attention for anything peripheral. In this case, security was not seen as part of the explicit task completion, and thus was ignored in favour of completing the required task. Based on these difficulties, we were unable to draw any conclusions about how people use security cues when they are carrying out normal web browsing behaviour.

4.2 Phase 2 Results: Focus on Security

In the second phase of our study, participants were asked to pay particular attention to security information. We wanted to motivate them to look for this data to see how successfully they found the necessary information in the browser. This phase provided a number of observations, which are described below.

Overall usage of web security information

In both post-condition questionnaires, people self-reported on what security information they used. In Phase 1, they stated what they used in *general* when browsing, and in Phase 2, what they used specifically *in this study*. (For Phase 2, they were asked what cues they

had used overall across all four pages, rather than what was used specifically for each page.) Table 1 shows the combined results from both phases. The numbers in the first two columns represent the number of participants (out of 16) that reported using that evidence. The third column shows the number of participants we *confirmed* as having used that cue in the study, as verified by the eyetracking data from Phase 2. "Not applicable" indicates that the type of evidence could not be verified from eyetracking data (e.g., "type of site").

type of evidence	in general	in study	verified
`https`	4	7	7
lock or key icon	12	12	11
certificate	7	3	2
site statements	8	7	9
type of site	14	14	n/a
type of information	7	11	n/a

Table 1: The number of participants who used each security cue (/16). The first column represents self-reported usage in general, the second column self-reported usage for this study, and the third column indicates the number verified. n/a indicates information that cannot be verified from the type of data gathered.

Recollection of the items *actually* checked in the study was good, with few contradictions between the eyetracking data and the self-reporting data. Note that for "site statements", we observed more people actually checking this information than was self-reported. (Included in the category of "site statements" were items such as security and privacy policies, and assertions about the use of SSL.) It is possible that some people read these items but did not include them in their list of reliable evidence. People also relied on additional information, such as the type of site or information; for example, people may expect that security features have been provided on certain sites (such as online banks) and for entering confidential data (such as passwords). This information cannot be verified by eye-tracking data, as it is contextual rather than visual; therefore, we omitted it from our analysis.

Given that a list of possible security cues was provided shortly at the end of Phase 1 (right before starting Phase 2), it is possible that the questionnaires influenced participant behaviour, by alerting them to cues that they would have otherwise ignored. Participant responses suggest that this bias is minor, in that there were few differences between self-reported sources of information in Phases One and Two; for example, few people stated that they did **not** use a specific cue, and then proceeded to start using it. The greatest change was observed for the https cue, where three people did not select it in Phase One but then used in it Phase Two. However, for the lock icon, there were no changes between phases; the 12 people who said they usually used it were the same ones who reported they actually did so in Phase Two. Although the bias appears minor, it should be avoided in future studies.

Some participants provided additional sources that they used in making web security assessments, which are listed in Table 2. The table also indicates at which point in the study the item was provided (i.e., in general or specifically in this study), and whether this use was verified, where appropriate.

type of evidence: other	in general	in study	verified
reputation of host company	✓		
browser warnings	✓		
popularity of site; known by my friends	✓		
for retailers, comments made by others	✓		
amount of encryption used	✓		
corporation size: bigger ones are more trustworthy	✓		
my browser preferences; enabling and disabling options to reduce risks	✓		
reputation of site		✓	n/a
when they direct you to sign out rather than close the window		✓	n/a
privacy report, page source info		✓	✓

Table 2: Types of evidence that people listed in addition to the items in Table 1. Checkmarks indicate whether that item was reported as generally used or used in the study, and whether we verified this (in the study only).

Four of our participants did not use browser cues at all, relying instead on elements such as statements about site security or the type of site. This was a small group, but these were not novice users. As with our other participants, these four people had been using the web for over 5 years and frequently engaged in secure transactions. However, to this group, browser security cues were practically invisible. It is unclear if this group requires education about browser security, a more obvious set of cues, or both; what is clear is that they could benefit from greater awareness of security.

Usage of browser-specific security cues

Our work focuses mainly on those cues provided by the browser itself, so we will discuss those elements in more detail. The lock icon was by far the most-used item, with 11 out of 16 participants using it (verifiably) at some point in the study. Interestingly, although `https` was used somewhat (7 out of 16), it was used only by people who also used the lock. In fact, the eyetracking data showed that most of those who

checked `https` also checked the lock directly before or afterwards. These two elements seem to be linked in people's minds. This makes sense, given that both items refer to SSL connections, but it also suggests that the `https` and the lock provide redundant information. In fact, the lock provides more data than appears at first glance, but the user must interact with the lock in order to reveal this data. When the cursor is rolled over the lock icon, a tool-tip appears with the message "SSL Secured" and the number of bits in the encryption key. In addition, when double-clicked, the lock icon pops up a window with certificate details. This information is crucial for verifying that the remote party is the one to whom you wish to send your private data.

If the lock is not clicked, it is not being used to its full potential. It merely presents the same information as `https` in the URL, with two possible advantages: 1) having an additional SSL cue might help if the `https` was overlooked, and 2) the lock metaphor may suggest a secure connection more so than the word `https` does. In our study, only two people clicked on the lock.

Because the lock icon is the portal to the certificate data, the two people who clicked on the lock were the only two people who checked the certificate. A third person investigated some certificate data, but only viewed the certificates that had been imported to the local host, rather than the remote certificates. Overall, the certificate was an unpopular source for security information. This is understandable, in light of the complexity of certificate information. Although two people examined the certificates, neither one noticed that they were completely false—they were self-signed and did not make any reference to the real site that they supposedly belonged to (such as the bank). Thus, even when the information is found, it may not be helpful.

For the most part, those who were familiar with the lock and `https` were able to find them with little difficulty. One interesting phenomenon we observed was that two people looked in the wrong corner for the lock—the lower left rather than the lower right. Because there is an IE logo in that corner, they concluded they had seen a lock and believed that the page was secure. (In neither case was the real lock found, so it could not have been used as evidence.) One of these people was familiar only with IE, so could not be confused by the layout of another browser; it is unclear why they thought they had seen a lock. The second person used Netscape and Mozilla, both of which place the security icon in the lower left corner. (This person did eventually find the lock, but only in the final of 3 tasks.) If the icons are small, and located in the periphery, they may be misidentified, or confused with icons from other browsers. Those who use a variety of browsers may not realize that they have not looked closely enough at icon to recognize it. This suggests that standard layouts may be required for security icons across browsers, or that the icons themselves need to be more prominent and clear.

Recognition of secure connections

In addition to investigating how people used specific cues, it was also important to know what people had concluded about the security state of their connections, given those cues. We asked people which pages they thought were secure, and which were not secure. In reality, only the Gmail and Hotmail pages were secure, although the Hotmail page did not show this clearly. All the other pages—banks and retailers—used encrypted channels, but the data was sent to our local webserver, not to the real site; this was revealed only through the certificate. (None of our participants recognized this at any point in the study.)

Our webserver provided SSL wherever the real site did; for example, we did not use unsecured HTTP on the bank sites, because all of those sites used secure HTTP in practice. We wanted to simulate reality, not fool our participants. After we completed the second phase, in which we asked people to pay attention to security, we asked them to give us their conclusions. We asked for assessments of the security of the bank sign-in page, the bank account information page, the Hotmail sign-in page, and the bookstore credit card payment page. We asked: based on the security information you saw, what did you conclude? Overall, people's recognition of an encryption channel was good. However, their use of certificates was weak, as they were rarely used and poorly interpreted.

Bank sign-in: Fifteen participants (out of 16) thought that the bank sign-in page was secure. The one person who thought it was insecure based their decision on lack of clear security statements on the bank's information page. None of the participants used the certificate data to conclude the connection was insecure.

Bank account page: We expected that very few people would continue to check for security after they signed in, so participants were asked about the account information page, which loaded after the login page. Twelve out of 16 participants stated that they *had* checked. However, the eyetracking data failed to verify that most participants actually performed the check: there was some evidence that three people had checked, and very weak evidence for two more. (This is less than half of those who claimed to have performed a check.) Of the 12 who claimed they checked, only one of these participants believed that the page was not secure (the same person as in the previous question, again using only website statements).

The lack of evidence checking after sign-in indicates that people feel a secure "phase" has been entered, which suggests that the user can relax and stop looking for cues. In this particular case, the data would be protected, as all of the real sites maintain SSL security after the sign-in page. However, if a security problem occurred after sign-in, a strong alert would be required to get the user's attention.

Hotmail sign-in: We chose to use Hotmail because of its familiarity, and because the site does not look secure at first glance. The sign-in page is not hosted on an SSL server, but an SSL connection is opened up to transmit the password when the "sign in" button is pressed. Users would have to examine the page's source code to discover this, and they could very briefly see the lock and `https` flash by, if they paid close attention. There were a variety of conclusions about the state of Hotmail security. Nine people checked for security, three did not, and four did not find security information.

Out of the nine people who stated that they checked for security, four concluded the page was secure, three concluded it was not, and one simply stated "I didn't understand." Overall, the cues provided during the sign-in process were confusing. Only three people saw the lock flash by, and one of these people used this transitory cue as evidence that the sign-in page was *insecure*.

Of the four people who could not find security information, two concluded that the page was *not* secure. The eyetracking data revealed how most of these people visually searched for security information: they all looked in the lower right during the search, where the lock would have been located. One person also used the MSN search tool to look for "encryption", in an unsuccessful attempt to find out what kind of security was being used for Hotmail.

Credit card page: Finally, we asked participants about the credit card payment page. All web pages in the book purchasing process were secured with SSL, including the single page that collected the credit card number. There was no sign-in process for our task (as there would be for a real customer). Thirteen of the 16 participants concluded that the page was secure. Of the three that did not, one based their decision only on site information ("that page didn't seem familiar"), a second could not find security information, and the third person did not give a reason. No participants used certificate information to conclude that the connection was forged. Two people checked the certificate (the same people as before), but neither noticed a problem with it, stating that the connection was secure.

4.3 Follow-up Survey: The Lock

Because so few people clicked on the lock, we wanted to know the reasons why. It was possible that they did not know that the lock was interactive. Nothing about the icon that suggests it performs an action, and in browsers such as Apple's Safari, it is *not* interactive. Another possibility is that people do not find its information to be useful, because of the expertise required to interpret certificate data. In order to find out more about this issue, we asked our participants to complete a post-study questionnaire. We emailed questions to our participants one week after the original study. They were asked what the lock told them about security, how often they used it (and why), and whether they ever clicked on it. For those who *had* clicked on it, we asked: what does it tell you, and is this information helpful to you? (Participants were asked to not look up the answer before responding.)

Responses were received from 14 of our 16 original participants, including the two who had been observed using the certificate function. Six of these 14 reported that they *had* clicked on the lock in the past; this suggests that although at least six participants knew of the additional functionality, only two took advantage of it in our study. All those who said they have clicked on the lock were able to correctly identify that the certificate information appears. This group was divided on how useful this certificate information is. Two people thought the information was very useful, although neither person used it in our study. The other four felt it had limited usefulness. In general, respondents had trouble interpreting the certificate information. These results suggest three challenges: how to make users aware of this certificate information, how to easily get this information to those who want it, and how to describe certificate information in a meaningful way.

5 Design Implications

Our research in visual security cues discovered information that can be applied to browser design and evaluation. In summary, we found that

- the lock icon is the browser security cue that is most often looked at, but few interact with it;
- some experienced web users do not take any notice of browser security cues;
- small browser icons can be easily misidentified or confused, especially given the non-standard layouts among browsers;
- certificates as sources of information are seldom used and rarely understood; and
- people tend to stop looking for security information after they have signed into a site.

These findings suggest a number of potential improvements for security support. For example, the lock has become a standard symbol for a secure connection. Making major modifications to this symbol, such as using a different object, may be disorienting: users now

expect to find a lock in a browser window. At the same time, the lock's small size and peripheral location can cause it to be overlooked by the casual observer; this icon could be made more prominent. In addition, the lock does not clearly indicate that it is interactive, which means that few click on it to get certificate data.

Another consideration is standard layout of icons across browsers. For example, IE has the lock on the lower right, and older versions of Firefox had it on the lower left. Because IE puts its small logo in the lower left, Firefox users may mistake this symbol for the lock, unless they look carefully. This suggests that standardized locations for security icons would be beneficial.

Finally, certificate data is poorly presented; it is both difficult to locate and to interpret. It is possible that it should have its own icon, separate from the lock, which makes it easier to find. As well, certificate data must be presented in a clearer manner; even those with technical backgrounds find it difficult to interpret.

6 Conclusions and Future Work

Although secure web transactions are common, the effectiveness of browser security cues is not well understood. We have addressed this problem by completing an assessment of current usage, which identified aspects of browser security that could be improved and raised significant questions for further research.

One component of our work that must be completed is the study of browser alerts. Because this is the other main source of browser-specific security information — besides the visual elements already considered—it represents a crucial piece of our overall research goals. As part of this research, we must incorporate personalized browser security settings, and consider how these settings change over time. For example, people may turn off or ignore certain warnings if they interfere with the task at hand. One possible approach to reducing overload would be to create translucent pop-ups that automatically close after a short period; a study would be required to determine the most effective means of dealing with these pop-up alert messages.

In addition, we were unable to determine how people use cues when they are not in a state of heightened security awareness. We intend to pursue this question further, although it presents methodological difficulties; it may be possible to design a minimally-invasive study that would be acceptable to our participants. If we are able to observe natural browsing behaviour, then we can develop and test designs to making the cues more obvious without making them overly intrusive. HCI research in the area of notification may be helpful here, such as Bartram's work on moving icons [2].

We would also like to evaluate the visual cues in Mozilla Firefox: in version 1.0, the location bar turns yellow for `https` connections, and remains white for regular http connections. In addition, lock icons appear in *both* the location bar and a lower corner of the browser; both can be clicked on to reveal certificate information. These cues do not appear to have been formally evaluated for effectiveness, and it is unclear how these multiple cues can be used effectively in combination. For example, is it confusing to have two identical icons, with identical functionality, in different locations? Does the yellow location bar draw attention away from the lock icons (and therefore away from the certificate data)? Do people recognize that they can interact with these icons? A formal study would help clarify the effectiveness of these design choices.

Acknowledgements

The authors wish to thank NSERC and the Killam Trust for financial support, and the EDGELab members and GI reviewers for their valuable insights and feedback.

References

[1] Adams, A. and M. A. Sasse (1999). "Users are not the enemy: Why users compromise security mechanisms and how to take remedial measures." Communications of the ACM 42(12): 40-46.

[2] Bartram, L., C. Ware, et al. (2001). Moving Icons: Detection and Distraction. Interact 2001, Tokyo, Japan.

[3] DIMACS (2004). Program of the DIMACS Workshop on Usable Privacy and Security Software.

[4] Friedman, B., D. C. Howe, et al. (2002). Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design. Thirty-Fifth Annual Hawaii International Conference on System Sciences.

[5] Friedman, B., D. Hurley, et al. (2002). Users' conceptions of Web security: A comparative study. CHI 2002, Minneapolis, MN.

[6] IEEE (2004). IEEE Security and Privacy Magazine Special Issue on Security and Usability, IEEE Computer Society.

[7] Saltzer, J. H. and M. D. Schroeder (1975). "The Protection of Information in Computer Systems." Proceedings of the IEEE 63(9): 1278–1308.

[8] Tobii Technology AB (2003). User manual, Tobii eye-tracker.

[9] Whitten, A. and J. D. Tygar (1999). Why Johnny can't encrypt: A usability evaluation of PGP 5.0. 8th Usenix Security Symposium, Washington, D.C.

[10] Yurcik, W., J. Barlow, et al. (2003). Two Visual Computer Security Network Monitoring Tools Incorporating Operator Interface Requirements. CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale, Florida.