What Do We Really Know about How Habituation to Warnings Occurs Over Time? A Longitudinal fMRI Study of Habituation and Polymorphic Warnings

Anthony Vance, Brock Kirwan, Daniel Bjornn, Jeff Jenkins, Bonnie Brinton Anderson

Brigham Young University

{anthony.vance, kirwan, dbjornn, jeffrey_jenkins, bonnie_anderson}@byu.edu

ABSTRACT

A major inhibitor of the effectiveness of security warnings is habituation: decreased response to a repeated warning. Although habituation develops over time, previous studies have examined habituation and possible solutions to its effects only within a single experimental session, providing an incomplete view of the problem. To address this gap, we conducted a longitudinal experiment that examines how habituation develops over the course of a five-day workweek and how polymorphic warnings decrease habituation. We measured habituation using two complementary methods simultaneously: functional magnetic resonance imaging (fMRI) and eye tracking.

Our results show a dramatic drop in attention throughout the workweek despite partial recovery between workdays. We also found that the polymorphic warning design was substantially more resistant to habituation compared to conventional warnings, and it sustained this advantage throughout the five-day experiment. Our findings add credibility to prior studies by showing that the pattern of habituation holds across a workweek, and indicate that cross-sectional habituation studies are valid proxies for longitudinal studies. Our findings also show that eye tracking is a valid measure of the mental process of habituation to warnings.

Author Keywords

Security warnings; habituation; polymorphic warnings; functional magnetic resonance imaging (fMRI); eye tracking; longitudinal experiment.

ACM Classification Keywords

H.5.2. Information interfaces and presentation: User Interfaces; J.4. Social and Behavioral Sciences.

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

Copyright is held by the owner/author(s).
CHI 2017, May 06–11, 2017, Denver, CO, USA
ACM 978-1-4503-4655-9/17/05.
<http://dx.doi.org/10.1145/3025453.3025896>

INTRODUCTION

Users often represent the last line of defense between attackers and organizations. User responses to security warnings is thus a critical aspect of behavioral security [5, 26, 34]. A major inhibitor of the effectiveness of security warnings is habituation: diminished attention due to frequent exposure to warnings [36]. Through this process—also known as warning blindness [53] or fatigue [1]—users’ attention to warnings can attenuate to the point where they hardly see the warning any longer. Although this problem is widely recognized [e.g., 3, 6, 21, 32, 38, 46], a major limitation of past studies that examine habituation is that they used cross-sectional (i.e., single point in time) experimental designs. However, habituation is fundamentally a neurobiological phenomenon that develops over time [40]. Thus, past research on habituation to security warnings has provided only a snapshot of a dynamic problem. Our first research question is therefore:

RQ1. How does habituation develop in the brain in response to security warnings over time?

We address this question by extending the work of Anderson et al. [4, 6] in two key respects. First, we performed a longitudinal experiment that examined user habituation to security warnings over the course of a five-day workweek. This allowed us to measure not only the attenuation of users’ responses over the course of the workweek but also another core characteristic of habituation: *response recovery*—the increase in user response after a rest period during which the stimulus is absent [40]. Given that past work has been based exclusively on cross-sectional experimental designs, this paper is the first to explore how users recover from habituation effects between exposures to warnings.

Second, Anderson et al. [4, 6] used fMRI and mouse cursor tracking to measure habituation, neither of which directly measures visual processing. Because security warnings are mainly graphical in nature, it is important to understand how visual attention to the warnings changes over time. In this paper, we measured habituation using two neurophysiological methods simultaneously, fMRI and eye tracking, using an fMRI-compatible, long-range eye tracker (see Figure 1). This allowed us to measure how both

Figure 1. EyeLink 1000 Plus long-range eye tracker, mounted under the MRI viewing monitor.

cognitive processing and visual inspection of a warning habituate over time.

Third, Anderson et al. [4, 6] showed that polymorphic warnings that repeatedly change their appearance can be effective in maintaining attention during a single experimental session, but they left unresolved whether this novelty fades with time. We extend their study by testing their polymorphic design in our longitudinal experiment, hypothesizing that the polymorphic warning will exhibit less attenuation and greater recovery across the five-day workweek as compared to conventional warnings.

Our second research question therefore is:

RQ2. Are polymorphic warnings resistant to habituation over time?

Our results showed a dramatic drop in attention in terms of neural activity and eye fixations after only the second warning exposure, with further decreases throughout the workweek. We found that participants' attention partially recovered between workdays when the stimulus was absent. The polymorphic warning design was substantially more resistant to habituation as compared to conventional static warnings, and this advantage persisted throughout the five-day experiment. Further, our findings add credibility to prior studies by showing that the pattern of habituation holds across a workweek, and indicate that cross-sectional habituation studies are valid proxies for longitudinal studies. Our findings also show that eye tracking is a valid measure of the mental process of habituation to warnings.

Why fMRI

We chose fMRI over traditional behavioral methods with superior ecological validity because it provides unique neural insights to better explain *why* people habituate to warnings. First, fMRI can capture automatic or unconscious mental processes that are difficult or impossible to measure with traditional tools [24]. We show in this paper that habituation to repeated warnings occurs automatically at the neural level. This explains that the root cause of habituation

is an obligatory neurobiological response, rather than willful carelessness on the part of users. Second, fMRI allows us to show how the brain recovers to habituation effects between exposures to warnings. None of these insights would be possible using traditional methods.

Finally, the use of fMRI aids in the development of more habituation-resistant security warnings. As noted in Dimoka et al. [24], “rather than relying on perceptual evaluations of IT artifacts, the brain areas associated with the desired effects can be used as an objective dependent variable in which the IT artifacts will be designed to affect” [24, pg. 700]. In this study, fMRI allows us to evaluate our polymorphic design using neural activity as our dependent variable. This enabled us to more directly address the problem of habituation.

LITERATURE REVIEW

Habituation is widely recognized as “the simplest and most basic form of learning” [39, p. 125]. It is believed to be ubiquitous in the animal kingdom, having been found “in every organism studied, from single-celled protozoa, to insects, fish, rats, and people” [16, 40, p. 125]. Habituation is an important survival mechanism allowing organisms to filter out irrelevant stimuli in the environment and thus to conserve energy for responding to stimuli that are relevant for survival [49]. Not surprisingly, humans also exhibit habituation to a wide variety of stimuli—visual, auditory, and others—and this is evident as early as infancy [17].

Given its strong security implications, habituation is frequently cited as a key contributor to users' failure to heed warnings. However, many studies infer the presence of habituation rather than empirically examining it [13, 27, 48, 52]. For example, Egelman et al. [26] found a correlation between user disregard for warnings and user recognition of warnings as previously viewed, attributing this correlation to habituation. Sunshine et al. [54] observed that participants remembered their responses to previous interactive security warnings and applied them to new warnings—even if the level of risk or context had changed—and likewise pointed to habituation as the probable cause. Akhawe and Felt found that the most common browser SSL error had the lowest adherence rate and the shortest response time, noting that this result was “indicative of warning fatigue” [1, p. 268].

Bravo-Lillo et al. [10, 11] empirically measured habituation, albeit indirectly. They measured habituation in terms of the percentage of users who immediately recognized that the contents of a dialog message had changed after a rapid habituation period. Only 14% of the users in their study immediately recognized the change in the dialog message [11]. A follow-up study examined four different levels of warning exposure frequency. They found that increasing the frequency with which a dialog box was displayed caused a threefold decrease in the proportion of users who immediately recognized a change in the dialog message [10].

In contrast to the above studies, Anderson et al. [4] used fMRI to measure habituation in the brain in response to warnings. Their results showed a large drop in activity in the visual processing centers of the brain after only the second exposure to a warning, and they found further decreases with additional exposures.

However, all these prior studies share a major limitation: they are based on single experimental sessions. Cross-sectional experiments that observe habituation at a single moment in time are unable to capture how habituation develops over several days. They are also unable to measure response recovery after a warning has been withheld, which is a key characteristic of habituation [40]. As a result, our understanding of how habituation changes over time as well as of how to address the problem is limited. We address this research gap in this study.

HYPOTHESIS

We develop our hypotheses around the two most prevalent characteristics of habituation: (1) response decay—an attenuation of a response with multiple exposures—and (2) response recovery—the increase in response after a rest period during which the stimulus is absent [40]. Hypothesis 1 explores how user response to security warnings weakens over the course of several repeated viewings and how polymorphic warnings (described below) can deter this effect. Hypothesis 2 explores how user response to warnings recovers after the warning is withheld and how polymorphic warnings enhance this recovery. Our hypotheses rest on a prominent theory from neurobiology: the dual-process theory of habituation (DPT) [31].

Response Decay

DPT [31] explains that the brain creates a mental model when exposed to a stimulus (e.g., when seeing a warning). When people see the stimulus again, they automatically and unconsciously compare the stimulus to this model. If the model and stimulus are similar, people pay less attention to it. This automatic, unconscious mechanism becomes more ingrained with each successive repetition of the warning.

We predict that this habituation will occur both when viewing repeated warnings within a single computing session and when viewing repeated warnings in computing sessions over consecutive days [30]. When viewing repeated warnings within a single computing session, the brain creates a robust mental model of the security message, resulting in habituation during that session. However, these mental models can also persist across several days and even for much longer periods. Over successive days, users will thus rely on their mental models rather than actively process the warning [40]. In summary, we hypothesize:

H1a: Users habituate to warnings in computing sessions over consecutive days.

We hypothesize that users will habituate more slowly to polymorphic warnings—warnings that change their appearance with each repetition [4]—than to static warnings. Wogalter states that, “habituation can occur even with well-designed warnings. . . . Where feasible, changing the warning’s appearance may be useful in reinvigorating attention switch previously lost because of habituation” [58, p. 55]. Changing the appearance of a warning creates novelty. The *orienting reflex*, described by DPT as the primary reaction of the body to a novel stimulus, is influenced by a comparison of the current stimulus with a mental model of the stimulus as it was previously experienced. If a new or changed stimulus is experienced that does not match the mental model, then the response strength will recover (e.g., people will pay more attention to the warning) [51]. This process is called *sensitization* [31]. Sensitization counterbalances habituation [40]. Consequently, by changing the appearance of a warning, users will be unconsciously sensitized and therefore less habituated to polymorphic warnings on both the neural and behavioral levels [7].

We predict that polymorphic warnings will engender sensitization, reducing habituation within a single computing session as well as between computing sessions over multiple days. When users encounter a polymorphic warning in a future computing session, it may contradict a weaker mental model and be perceived as novel (i.e., cause an orienting reflex) [18, 55]. In summary, we hypothesize:

H1b: Users habituate less to polymorphic warnings than to static warnings in computing sessions over consecutive days.

Recovery

Although users will habituate to warnings, we predict that they will partially recover from the habituation after a day’s rest period without seeing warnings. Decay theory [8] explains that memory becomes weaker due to the passage of time. When a warning is withheld for a day, the mental model of the warning will become weaker. Therefore, when users see this warning in the future, it will be less likely to match the mental model and will appear novel. In response to this novelty, the response strength will recover and the sensitization process will increase a person’s attention to the warning, thus counteracting habituation [14].

Although the mental model diminishes with time, it is unlikely to fade completely within a single day. The brain will still inhibit the behavioral response to the stimulus, and habituation will occur. However, this response inhibition or habituation is likely to be weaker when users see a warning after it has been withheld for a day as compared to when they see it repeatedly within a single computing session [40]. In summary, we hypothesize:

H2a: If warnings are withheld after habituation occurs, the response recovers at least partially the next day.

Message Content: Pictorial symbols (e.g., an exclamation point) [35, 50]

Warning Appearance: Color [9, 44]

Animation: Jiggle, scale/zoom [11, 28, 37]

Figure 2. Symbol, background color, zoom and jiggle variations.

We predict that the amount of recovery from day to day will be greater for polymorphic warnings than for static warnings. As previously discussed, the mental models of polymorphic warnings are weaker and less stable than the models of static warnings. Less stable mental models (i.e., mental models with less reinforcement) fade more quickly than stable models [40]. After users do not see a warning for a day, they are more likely to perceive the polymorphic warning as novel. As a result, user response to polymorphic warnings will recover to a greater degree than the user response to static warnings.

Furthermore, if the polymorphic warning continues to change its appearance from day to day, it is even more likely to differ from the existing mental model, weakening behavioral inhibition, increasing sensitization, and enhancing response recovery [40]. Conversely, with static warnings, response recovery will be weaker because the

mental model is more robust, reinforced by repetitive exposures to the same warning over time [29, 31]. The behavioral response will be inhibited to a greater degree, and habituation will be more pronounced [40]. In summary:

H2b: If warnings are withheld after habituation occurs, response recovery is stronger for polymorphic warnings than for static warnings on the next day.

POLYMORPHIC WARNING DESIGN

Anderson et al. [4] found that four of their twelve graphical variations maintained attention better than the others. We used these four variations for our study: (1) including a pictorial symbol, (2) changing the warning's background color to red, (3) using a "jiggle" animation when the warning appears, and (4) using a zoom animation to make the warning increase in size. Figure 2 shows each variation for one sample warning with its supporting sources. Given this support, we used these four variations of the polymorphic warning to test our hypotheses.

Neurophysiological tools can be used to evaluate UI designs. Riedl et al. explained that neurophysiological measures are beneficial "to the design of ICT artifacts" [41, p. ii] and that "researchers could use the theory of controlled and automatic brain processes to . . . allow for a better design of IT artifacts and other interventions" [42, p.250]. Further, Dimoka et al. [25] argued that these measures should be used as dependent variables in evaluating IT-artifact designs:

"Rather than relying on perceptual evaluations of IT artifacts, the brain areas associated with the desired effects can be used as an objective dependent variable in which the IT artifacts will be designed to affect (p. 700)."

We use precisely this approach to evaluate the polymorphic warning design.

METHODS

To test our hypotheses, we conducted a multimethod study, simultaneously collecting both fMRI and eye-tracking data. This allowed us to capitalize on the strengths of each method while mitigating their limitations [57]. The use of these methods is preferable over simple behavioral measures such as response time or response choice because they are able to more directly measure the effect of the UI artifact on the underlying cognitive processes. fMRI is useful in measuring neural activity by tracking changes in blood oxygenation levels (the blood oxygen level-dependent or BOLD response) in specific areas of the brain. This allows researchers to identify distinct regions of the brain where activity is correlated with cognitive processes. fMRI identifies regions in terms of voxels or small 3 mm cubes, which makes it ideal when high spatial resolution is required [23]. A neural manifestation of habituation to visual stimuli in the brain is called repetition suppression (RS): the reduction of neural responses to stimuli that are repeatedly viewed [29]. In our case, high spatial resolution was

important because it allowed us to disentangle RS effects from sensory adaptation or fatigue effects [40].

We used fMRI to capture evidence of the RS effect, which is a reduction in the degree of fMRI activation (as measured by the BOLD response) that occurs as a participant is exposed to multiple repetitions of a stimulus—a robust indicator of habituation [29]. We utilized the differential RS effect in various brain regions to map sensitivity to repetitive security warning stimuli.

Concurrent with the fMRI scan, we used an eye tracker to measure the eye-movement memory (EMM) effect—another robust indicator of habituation [45]. The EMM effect manifests in fewer eye-gaze fixations and less visual sampling of the regions of interest within the visual stimulus. Memory researchers have discovered that the EMM effect is a pervasive phenomenon in which people unconsciously pay less attention to images they have viewed previously. With repeated exposure, the memories become increasingly available, thus requiring less visual sampling of an image [33].

One strength of eye tracking is its temporal resolution, which allows researchers to measure with millisecond precision the attentional process of participants' responses to repeated stimuli. Thus, fMRI (with high spatial resolution) and eye tracking (with high temporal resolution) complement each other, measuring both a behavioral manifestation of attention (i.e., eye movements) as well as the neural activity that drives attention.

Participants

We recruited 16 participants from a large US university. This number of participants is consistent with other fMRI studies [22]. Participants were between 19 and 29 years of age (the mean age was 23.3 years), right-handed, native English speakers, had normal or corrected-normal visual acuity, and were primarily PC users. One subject was excluded from the study due to scanner malfunction, resulting in 15 total participants (eight male, seven female).¹ Each participant engaged in five fMRI scans: one at the same time each day for five consecutive days. Upon arrival, the participants were screened to ensure MRI compatibility. They were then given instructions about the task and placed in the scanner. Each scan lasted 30 minutes, beginning with a structural scan and followed by two functional scans that displayed the warnings and images.

Ethics

The university Institutional Review Board (IRB) approved the protocols used. Participants were verbally briefed about MRI procedures as well as the task and purpose of the experiment before entering the scanner.

Experiment Design

For each participant, warning stimuli were randomly assigned to conditions that remained the same across the five-day experiment. In addition, the order of the presentation of the warning stimuli was randomized per participant, per day, as follows. First, 40 images of various computer-security warnings (e.g., browser malware and SSL warnings, anti-virus software warnings, and software signing errors) were randomly split into two pools: one for the static condition and the other for the polymorphic condition. The 20 warnings in the static condition would be repeated four times at random times each day for the course of the study (20 static warnings $\times$ 4 repetitions = 80 static warning images to display).

The four polymorphic variations were then applied to the 20 warnings in the polymorphic pool so that each of the polymorphic variations randomly appeared once a day during the experiment. That is, for each polymorphic warning, all four polymorphic variations were shown each day (20 polymorphic warnings $\times$ 4 variations = 80 polymorphic images to display). The order of these variations was randomized so that each day had a different presentation order for the variations of each warning.

Next, 120 images of general software (e.g., Windows Explorer, Control Panel, Microsoft Outlook) were also randomly split into two sets. The first set consisted of 20 images that would be randomly repeated four times each day (20 general software images $\times$ 4 repetitions = 80 general software images to display). The remaining 100 images would be divided evenly across days so that a different set of 20 images would be randomly displayed once on each day. These images were used to create a baseline of unique presentations throughout the task. By comparing the responses for each repeated image to the unique baseline images, we were able to distinguish the habituation effect from attention decay attributable to participants' fatigue over time.

Upon completion of the randomization there were 260 images to be presented each day (80 static images + 80 polymorphic images + 80 general software images + 20 unique general software images = 260). These 260 images were randomly displayed across two blocks of 7.7 minutes each, with a 2-minute break in between blocks. Images were displayed for 3 seconds each, with a 0.5-second inter-stimulus interval. While participants saw the warnings, they were required to rate the severity of the content of the warning on a 4-point scale. They did this using a fMRI compatible button-input device. The purpose of this task was to help keep the participant engaged in a context relevant to warnings. The technical details of the fMRI scans and procedures are documented in the appendix.

¹ We conducted a pilot study that revealed a large estimated effect size for the repetition effect (partial $\eta^2 = .7$). Using this estimated effect size, an a priori power analysis indicated that we would need four subjects to achieve

power greater than .8, indicating that a sample size of 15 is more than adequate.

ANALYSIS

We analyzed each hypothesis separately for the fMRI and eye-tracking data. Our analyses are described below, followed by tests of our hypotheses.

fMRI Analysis

MRI data was analyzed using the Analysis of Functional NeuroImages (AFNI) suite of programs [20] (see appendix for details). Briefly, following standard data pre-processing, we conducted regression analysis on the individual subject fMRI data wherein behavioral vectors coding for stimulus type (e.g., security warnings, general software images) were modeled using a stick function convolved with the canonical hemodynamic response. Beta values for the conditions of interest were then entered into group-level analyses as whole-brain, multivariate model analyses to identify significant clusters of activation, or regions of interest (ROIs), consistent with the hypothesized pattern. Our prior work using fMRI to investigate habituation examined habituation within a single day [4]. Consequently, we were unsure if brain regions that demonstrate habituation effects within a day would also exhibit habituation effects over multiple days or if additional brain regions would also demonstrate long-term habituation effects. Accordingly, we used a data-driven approach to identify functional ROIs. Group comparisons were corrected for multiple comparisons using a voxel-wise threshold of $p < .02$ and a spatial-extent threshold of 40 contiguous voxels (1080 mm^3) for an overall corrected p -value $< .05$, as determined through Monte Carlo simulations [59].

All our hypothesis tests utilized the same ROIs. Graphs of brain activity in response to polymorphic and static warnings over consecutive days are presented for two brain regions in Figures 3 and 4.

Eye-Tracking Analysis

Eye-tracking data was collected using an MRI-compatible SR Research EyeLink 1000 Plus (see Figure 1). Fixations were defined as periods of time between eye movements that were not also part of blinks. Fixation count was used as the dependent variable in each analysis.²

The number of fixations for polymorphic and static warnings per warning repetition per day is shown in Figure 5. The mean and standard deviations of fixation count and fixation duration per day are shown in Table 1. Some of the polymorphic warnings were animated, which prevented participants from fixating upon the warning during the animation. To control for this, we normalized all intercepts to zero and controlled for warning type in the analysis, allowing for individual warning intercepts. This control allowed us to focus on and accurately analyze how fixations changed over time as an indicator of habituation.

² We chose fixation count as a more appropriate measure of habituation than fixation duration because the warning stimuli were displayed to subjects for the same duration. However, we replicated all analyses using fixation

Figure 3. Activity in the right inferior temporal gyrus in response to each presentation of static and polymorphic warnings. Beta values were extracted from a whole-brain analysis for each subject and then averaged across subjects according to stimulus condition.

Figure 4. Activity in the right ventral visual pathway in response to each presentation of static and polymorphic warnings. Beta values were extracted from a whole-brain analysis for each subject and then averaged across subjects according to stimulus condition.

Figure 5. Change in eye gaze fixations across viewings

duration as the dependent variable, and the results were the same as those obtained using fixation count as the dependent variable.

	Day 1	Day 2	Day 3	Day 4	Day 5
Fixation count mean	9.1	8.08	8.09	7.71	7.35
Fixation count SD	2.65	2.18	2.27	2.48	2.32
Fixation duration mean (ms)	2349	2204	2135	2113	2081
Fixation duration SD (ms)	450	325	384	441	444

Table 1. Absolute fixation count and fixation duration by day.

HYPOTHESES RESULTS

H1a Analysis: Users habituate to warnings over consecutive days.

fMRI Analysis: We conducted a whole-brain, multivariate model analysis [15] on the fMRI data holding gender, day, repetition number, and stimulus type (static warning and polymorphic warning) fixed, to find areas that responded to a linear trend on day number, collapsing across repetitions and stimulus types. In this analysis, two main ROIs were identified: the right and left insula.

To quantify the extent of the decrease in these ROIs, beta values were extracted for these regions and tested using a within-subjects, repeated measures ANOVA. Both the right [F (1, 597) = 67.87, $p < .001$] and left insula [F (1, 597) = 86.19, $p < .001$] exhibited a significant habituation effect across days (Table 2). Thus, the fMRI analysis supported H1a.

Eye-Tracking Analysis: In a linear mixed-effects model, we included fixation count as the dependent variable and the subject ID and warning ID as random factors. The presentation number (across days) was treated as a fixed factor, and visual complexity⁴ was included as a covariate. The eye-tracking analysis supported H1a; the beta of presentation number across days was significantly negative [χ^2 (1, N = 11,976) = 212.89, $p < .001$, $\beta = -0.1031$], indicating habituation. Visual complexity was also significant [χ^2 (1, N = 11,976) = 34.85, $p < .001$, $\beta = 0.3815$]. The R^2 of the model was 0.13.

6.3.1 H1b Analysis: Users habituate less to polymorphic warnings than to static warnings over consecutive days.

fMRI Analysis: We conducted a whole-brain analysis for a day by stimulus-type interaction. Two ROIs, the left middle frontal gyrus [F (1, 595) = 5.188, $p < .05$] and left middle

ROIs for Main Effect of Day

Region	# Voxels	Peak x	Peak y	Peak z	F Value	p Value
R. insula	160	-43	-16	3	67.87	< .001
L. insula	158	40	-16	0	86.19	< .001

ROIs for Day by Stimulus-Type Interaction

Region	# Voxels	Peak x	Peak y	Peak z	F Value	p Value
L. middle frontal gyrus	190	49	-31	18	5.19	.02
L. middle occipital gyrus	118	25	76	39	4.70	.03

Table 2. ROIs for habituation across days.

occipital gyrus [F (1, 595) = 4.697, $p < .05$], displayed a significant habituation interaction across days and between stimulus types (Table 2).

Eye-Tracking Analysis: We specified the same mixed-effects model as in H1a, except that we included an interaction term between the presentation number (across days) and a polymorphic dummy variable (coded as 1 for polymorphic and 0 for static).

The eye-tracking analysis supported H1b; the interaction between the presentation number and polymorphic dummy was significantly positive [χ^2 (1, N = 11,976) = 10.70, $p < .001$, $\beta = 0.024$], indicating that participants habituate less to polymorphic warnings than to static warnings over the course of several days. The main effects for both presentation number [χ^2 (1, N = 11,976) = 493.42, $p < .001$, $\beta = -0.115$] and polymorphism [χ^2 (1, N = 11,976) = 64.71, $p < .001$, $\beta = -0.725$] were also significant. Visual complexity, however, was not significant: χ^2 (1, N = 11,976) = 0.17, $p > .05$, $\beta = 0.026$. The R^2 of the model was .137.

6.3.2 H2a Analysis: If warnings are withheld after habituation occurs, user response recovers at least partially the next day.

fMRI Analysis: We first calculated recovery scores by subtracting the mean beta value of the last display of each stimulus type from the first display of that stimulus type on the following day (i.e., Day 2 Display 1 – Day 1 Display 4; etc.). A whole-brain, multivariate model analysis was then conducted to test for regions that displayed changes from baseline activation, which, collapsing across days, revealed four ROIs where there was significant recovery.

⁴ A MATLAB script was used to calculate visual complexity 43. Ruth Rosenholtz, Yuanzhen Li and Lisa Nakano. 2007. Measuring visual clutter. *Journal of vision*, 7 (2). 17. .

Hypothesis	Eye Tracking	fMRI
H1a: Users habituate to warnings over consecutive days.	Supported	Supported
H1b: Users habituate less to polymorphic warnings than to static warnings over consecutive days.	Supported	Supported
H2a: If warnings are withheld after habituation occurs, user response recovers at least partially the next day.	Supported	Supported
H2b: If warnings are withheld after habituation occurs, user response recovery is stronger for polymorphic warnings than for static warnings the next day.	Not supported	Not supported

Table 4. Summary of results

Post hoc analysis comparing specific days showed significant recovery for Days 2–4 in nearly every area, with no significant recovery on Day 5 (Table 3). Thus, H2a was supported by the fMRI data.

Eye-Tracking Analysis: We subtracted the fixation count for the first viewing of a warning on a given day from the fixation count of the last viewing of the warning on the previous day. We then tested this hypothesis using a t-test. The test results supported H2a: participants experienced a significantly positive recovery ($m = 0.369$, $sd = 3.171$) from day to day [$t(2377) = 5.672$, $p < .001$, $d = 0.233$].

6.3.3 H2b Analysis: *If warnings are withheld after habituation occurs, response recovery is stronger for polymorphic warnings than for static warnings the next day.*

fMRI Analysis: We analyzed the same ROIs found for H2a but augmented the model by including stimulus type (polymorphic or static) as a factor. None of the regions

displayed a significant recovery by stimulus-type interaction (Table 3). Thus H2b was not supported.

Eye-Tracking Analysis: We subtracted the fixation count for the first viewing of a warning on a given day from the fixation count for the last viewing of the warning on the previous day. We specified a linear mixed-effects model that tested whether warning type (polymorphic vs. static) predicted this difference. The subject ID, day interval (e.g., the difference between Day 1 and Day 2 was coded as 1), and warning ID were included as random factors. Polymorphism was included as a fixed factor, and visual complexity was included as a covariate. The eye-tracking analysis did not support H3b. Neither the warning type [$\chi^2(1, N = 2,400) = 1.92$, $p > .05$, $\beta = -0.166$] nor the visual complexity [$\chi^2(1, N = 2,400) = 1.16$, $p > .05$, $\beta = 0.072$] significantly predicted recovery between days. Table 4 summarizes the results for all hypotheses.

DISCUSSION

This paper makes several foundational contributions. Although habituation to security warnings has been studied before, it was previously unknown whether prior findings were valid outside of short experimental sessions [e.g., 1, 4, 11, 12, 32, 47, 54]. This is because habituation develops over time. Our results (1) add credibility to prior studies by showing that the pattern of habituation holds across a workweek, and indicate that cross-sectional habituation studies are valid proxies for longitudinal studies.

Similarly, Anderson et al. [4] showed that polymorphic warnings can slow habituation, but it wasn't clear whether they are effective over time. This study extends their work and (2) shows that polymorphic warnings sustain their novelty and thereby decrease habituation over time. We therefore recommend eliminating repetitive, routine security alerts and using polymorphic warnings for important security messages or to draw users' attention to new information on familiar-looking warnings to improve users' security behaviors.

					Recovery > 0		Recovery by Stimulus-Type Interaction		Recovery across Days		Day by Stimulus-Type Interaction	
Region	# Voxels	Peak x	Peak y	Peak z	t Value	p Value	F Value	p Value	F Value	p Value	F Value	p Value
R. ventral visual stream	334	-31	73	-9	4.52	< .001	0.015	.90	2.79	.10	1.02	.31
L. ventral visual stream	206	40	52	-12	4.00	< .001	0.02	.89	1.33	.25	.004	.95
L. inferior frontal gyrus	188	40	-1	27	5.31	< .001	0.083	.77	.01	.91	.06	.81
R. inferior frontal gyrus	54	-37	-1	30	4.35	< .001	1.115	.29	2.06	.15	.66	.42

Table 3. Regions of interest (ROIs) for recovery.

Further, our results showed that polymorphic warning designs need only to iterate through a small set of visual variations to sustain novelty over time. In our study, the polymorphic warning iterated through four design variations: including (a) a pictorial symbol, (b) changing the warning's background color to red, (c) using a "jiggle" animation when the warning appears, and (d) using a zoom animation to make the warning increase in size. Although participants saw hundreds of warnings each day for an entire workweek (and each variation therefore hundreds of times across the entire study), simply iterating through these four visual variations resulted in substantially less habituation over a five-day period as compared to static warnings. Thus, our results suggest that polymorphic warning designs need not be elaborate nor need they have unique variations every appearance. Rather, iterating through a simple set of polymorphic designs may be a cost-effective, easily-deployable solution to decrease habituation. However, we add the caveat that we only tested the effectiveness of polymorphic warnings in terms of neural activation. Further research is needed to examine the effect of polymorphic warnings on actual behavior.

We also (3) show using fMRI recovery of habituation after a warning is withheld for a time. Response recovery is a key characteristic of habituation [40], but it has not yet been examined in the context of security warnings. In our study, participants did not see the warnings for a 24-hour period between computing sessions. After this rest period, we found that participants experienced greater activation in the brain and fixated more on the warnings than they did at the end of the computing session on the previous day. Unfortunately, this wasn't enough to offset continued warning exposure. These results support the idea that habituation can be effectively mitigated by allowing time between displaying warnings. We recommend that future research examine if longer periods of time between warnings (greater than 24 hours) is more effective in offsetting habituation due to continued warning exposure.

Methodologically, we (4) triangulate eye tracking with fMRI. We found that the eye-tracking results closely mimic the fMRI results, suggesting that eye tracking is a valid and cost-effective alternative to fMRI for studying habituation to warnings as a mental process (as opposed to behavior that is influenced by this mental process) using a similar design as used in this study. This finding enables future researchers to conduct more ecologically valid habituation studies that use eye tracking in a normal computing environment.

Finally, we observe that the above contributions would not have been possible by observing behavior in traditional laboratory experiments. This is because warning disregard behavior does not explain why habituation occurs. In this paper, we were able to show that habituation of attention to warnings occurs automatically at the neurobiological level, and that the extent of habituation increased throughout the

week. Observing behavior would also not have provided a direct examination of recovery of habituation, nor allowed us to validate that eye tracking is a valid measure of the mental process of habituation.

LIMITATIONS AND FUTURE RESEARCH

Our research is subject to a number of limitations. First, the fMRI methodology requires users to lie down in a supine position while being scanned. Participants are limited in their movement during the scan. Thus, the physical setting of our experiment is different than how people naturally interact with computers and security warnings. However, only through the use of fMRI are we able to demonstrate that habituation is a natural result of how the brain works, suggesting that attempts to train users to try harder to pay attention to warnings are incomplete on their own. This insight is difficult to establish through behavioral HCI studies.

Second, participants viewed 260 warnings in each scan session. It is unlikely that a typical user would ever be exposed to so many warnings in a real work setting. However, this artificiality arguably made our test of habituation and recovery more robust [2, 11]. That is, if polymorphic warnings can slow habituation when users receive an unrealistically high number of warnings, they will likely be even more effective when users receive a realistic amount. The same logic applies to recovery from habituation effects [2, 11].

Third, participants saw images of warnings, rather than interacting with warnings in the context of normal computer use. This was necessary for precision in the fMRI methodology. However, in normal computer use, several factors may influence the generalizability of the results, such as dimensions of the computing experience that compete for attention, task demands, etc. Future research should corroborate our results with a study that has participants interact with warnings in their natural setting.

Fourth, our tests measured the habituation of attention to warnings rather than actual security behavior [56]. Although paying attention to a warning does not guarantee that the person will respond securely, it is a necessary precondition. In addition, this work extends that of Anderson et al. [4], who demonstrated that their polymorphic warning design was effective in reducing habituation behavior, as measured by mouse cursor tracking.

Fifth, our window of analysis was a five-day workweek. It is possible that the pattern of habituation may change after a longer period of time. However, this work is a significant improvement over past examinations of habituation that only examined habituation within a single point in time. Further, to our knowledge, no other examination of human habituation in the fields of neuroscience or neurobiology has used a window of analysis this long. Regardless, our five-

day window is sufficient to demonstrate (1) how the effects of habituation and recovery develop over time in response to security warnings, and (2) that the polymorphic warning design is substantially more resistant to habituation than are conventional warnings.

Finally, we address only visual habituation to visual security warnings. Future research should investigate other cognitive processes, such as increased semantic fluency, which may affect the repetition effect for security warnings.

CONCLUSION

Our results suggest a fundamental explanation of why users habituate to security warnings: human biology. Drawing on neuroscience and neurobiology, we observed how habituation develops over the course of a five-day workweek. When users viewed polymorphic warnings, they habituated less than when viewing static warnings. This suggests that polymorphic warnings are a cost-effective solution that can improve users' responses to important security messages. This study also shows that neurophysiological tools are useful to understand how the biology of the user relates to security UI design.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grant CNS-1422831 and a Google Faculty Research Award.

APPENDIX: FMRI AND EYE-TRACKING DETAILS

Equipment: MRI scanning took place at a university MRI research facility with the use of a Siemens 3T Tim-Trio scanner. For each scanned participant, we collected a high-resolution structural MRI scan for functional localization in addition to a series of functional scans to track brain activity during the performance of the various tasks. Structural images for spatial normalization and overlay of functional data were acquired with a T1-weighted magnetization-prepared rapid gradient-echo (MP-RAGE) sequence with the following parameters: matrix size = 224×256 ; TR = 1900 ms; TE = 2.26 ms; field of view = 219×250 mm; NEX = 1; slice thickness = 1.0 mm; voxel size = $1 \times .977 \times .977$ mm³; flip angle = 9°; number of slices = 176. Functional scans were acquired with a T2*-weighted gradient-echo echoplanar pulse sequence with the following parameters: matrix size = 64×64 ; field of view = 192 mm; slice thickness = 3 mm; TR = 2000 ms; 229 TRs; TE = 28 ms; number of slices = 39; voxel size = $3 \times 3 \times 3$ mm; flip angle = 90°. Slices were aligned parallel with the rostrum and the splenium of the corpus callosum. The first three volumes acquired were discarded to allow for T1 stabilization.

Eye-tracking data were collected on each scan using an MRI-compatible SR Research EyeLink 1000 Plus long-range eye tracker with a spatial resolution of 0.01° and sampling at 1,000 Hz. Eye movements were recorded for the right eye. A nine-point calibration routine was used to map

eye position in order to screen coordinates prior to each scanning block. Eye-fixation data was processed with DataViewer software (SR Research Ltd., version 1.11.900) to identify fixations and saccades. Saccades were defined as eye movements that met three different parameters: eye movement of at least .1°, velocity of at least 30°/second, and acceleration of at least 8,000°/second. Fixations were defined as periods of time between the saccades that were not also part of blinks.

Protocol: Scan sessions occurred at the same time each day, over a period of five days for each participant, resulting in five scans per participant. Upon arrival at the facility, participants completed a screening form to ensure MRI compatibility. Participants were verbally briefed about the MRI procedures and the task and were then placed supine in the scanner. Visual stimuli were viewed using a mirror attached to the head coil; this reflected a large monitor outside the scanner that was configured to display images in reverse so that they appeared normal when viewed through the mirror. Participants responded to stimuli using an MRI-compatible button box.

fMRI Data Analysis Details: Functional data were slice-time corrected to account for differences in acquisition time for different slices of each volume; then, each volume was registered with the middle volume of each run to account for low-frequency motion. A three-dimensional automated image registration routine, 3dVolreg [19], which uses Fourier interpolation, was applied to the volumes to realign them with the first volume of the first series used as a spatial reference. Data from each run was aligned to the run nearest in time to the acquisition of the structural scan. The structural scan was then co-registered to the functional scans. Spatial normalization was accomplished by calculating a transformation from each subject's structural scan to a template brain with advanced neuroimaging tools (ANTs) and then applying the transformation to the structural and functional data for each subject.

Behavioral vectors were created that coded for stimulus type (e.g., security warnings, general software images) and repetition number. These were then entered separately into single-participant regression analyses for each day. Stimulus events were modeled using a stick function convolved with the canonical hemodynamic response. Regressors that coded for motion and scanner drift were also entered into the model as nuisance variables. Spatial smoothing was conducted by blurring the resulting beta values with a 5-mm FWHM Gaussian kernel to increase the signal-to-noise ratio. Beta values for the conditions of interest were then entered into group-level analyses as we tested each hypothesis (below). Group comparisons were corrected for multiple comparisons using a voxel-wise threshold of $p < .02$ and a spatial-extent threshold of 40 contiguous voxels (1080 mm³) for an overall corrected p -value $< .05$, as determined through Monte Carlo simulations [59].

REFERENCES

1. Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness *Proceedings of the 22nd USENIX Conference on Security*, USENIX Association, Washington, D.C., 257-272.
2. T. S. Amer and Jo-Mae B. Maris. 2007. Signal words and signal icons in application control and information technology exception messages—Hazard matching and habituation effects. *Journal of Information Systems*, 21 (2). 1-25.
3. Bonnie Brinton Anderson, Jeffrey L Jenkins, Anthony Vance, C. Brock Kirwan and David Eargle. 2016. Your memory is working against you: How eye tracking and memory explain habituation to security warnings. *Decision Support Systems*, 92. 3-13.
4. Bonnie Brinton Anderson, C. Brock Kirwan, Jeffrey L. Jenkins, David Eargle, Seth Howard and Anthony Vance. 2015. How Polymorphic Warnings Reduce Habituation in the Brain: Insights from an fMRI Study *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, Seoul, Republic of Korea, 2883-2892.
5. Bonnie Brinton Anderson, Anthony Vance, C Brock Kirwan, David Eargle and Jeffrey L Jenkins. 2016. How users perceive and respond to security messages: a NeuroIS research agenda and empirical study. *European Journal of Information Systems*, 25 (4). 364-390. doi:10.1057/ejis.2015.21
6. Bonnie Brinton Anderson, Anthony Vance, C. Brock Kirwan, Jeffrey L Jenkins and David Eargle. 2016. From warnings to wallpaper: Why the brain habituates to security warnings and what can be done about it. *Journal of Management Information Systems*, 33 (3). 713-743.
7. Robert J. Barry. 2009. Habituation of the Orienting Reflex and the Development of Preliminary Process Theory. *Neurobiology of Learning and Memory*, 92 (2). 235-242. <http://dx.doi.org/10.1016/j.nlm.2008.07.007>
8. Marc G Berman, John Jonides and Richard L Lewis. 2009. In Search of Decay in Verbal Short-Term Memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35 (2). 317-333.
9. Curt C Braun and N Clayton Silver. 1995. Interaction of Signal Word and Colour on Warning Labels: Differences in Perceived Hazard and Behavioural Compliance. *Ergonomics*, 38 (11). 2207-2220.
10. Cristian Bravo-Lillo, Lorrie Cranor, Saranga Komanduri, Stuart Schechter and Manya Sleeper. 2014. Harder to Ignore? Revisiting Pop-Up Fatigue and Approaches to Prevent It. in, *USENIX Association*, 105-111.
11. Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W. Reeder, Manya Sleeper, Julie Downs and Stuart Schechter. 2013. Your Attention Please: Designing Security-decision UIs to Make Genuine Risks Harder to Ignore *Proceedings of the Ninth Symposium on Usable Privacy and Security* ACM, Newcastle, United Kingdom, 1-12.
12. José Carlos Brustoloni and Ricardo Villamarín-Salomón. 2007. Improving Security Decisions with Polymorphic and Audited Dialogs, ACM, 76-85.
13. José Carlos Brustoloni and Ricardo Villamarín-Salomón. 2007. Improving security decisions with polymorphic and audited dialogs. in *Proceedings of the 3rd symposium on Usable privacy and security*, ACM, 76-85.
14. Münire Özlem Çevik. 2014. Habituation, Sensitization, and Pavlovian Conditioning. *Frontiers in Integrative Neuroscience*, 8. 13. 10.3389/fnint.2014.00013
15. G. Chen, N. E. Adleman, Z.S Saad, E. Leibenluft and R.W. Cox. 2014. Applications of Multivariate Modeling to Neuroimaging Group Analysis: A Comprehensive Alternative to Univariate General Linear Model. *NeuroImage*, 99. 571-588.
16. G. R. J. Christoffersen. 1997. Habituation: Events in the History of Its Characterization and Linkage to Synaptic Depression. A New Proposed Kinetic Criterion for Its Identification. *Progress in Neurobiology*, 53 (1). 45-66. 10.1016/S0301-0082(97)00031-2
17. John Colombo and D. Wayne Mitchell. 2009. Infant Visual Habituation. *Neurobiology of Learning and Memory*, 92 (2). 225-234. <http://dx.doi.org/10.1016/j.nlm.2008.06.002>
18. Sam F. Cooke, Robert W. Komorowski, Eitan S. Kaplan, Jeffrey P. Gavornik and Mark F. Bear. 2015. Visual recognition memory, manifested as long-term habituation, requires synaptic plasticity in V1. *Nature Neuroscience*, 18 (2). 262-271. 10.1038/nn.3920
19. R. W. Cox and A Jesmanowicz. 1999. Real-time 3D image registration for functional MRI. *Magnetic Resonance in Medicine*, 42. 1014-1018.
20. Robert W Cox. 1996. AFNI: Software for Analysis and Visualization of Functional Magnetic Resonance Neuroimages. *Computers and Biomedical Research*, 29 (3). 162-173.
21. Lorrie Faith Cranor. 2008. A Framework for Reasoning about the Human in the Loop, *USENIX Association* 1:1-1:15.
22. A. Dimoka, R. D. Banker, I. Benbasat, F. D. Davis, A. R. Dennis, D. Gefen, A. Gupta, A. Lschebeck, P. H. Kenning, P. A. Pavlou, G. Muller-Putz, R. Riedl, J. vom Brocke and B. Weber. 2012. On the Use of Neurophysiological Tools in IS Research: Developing a Research Agenda for NeuroIS. *MIS Quarterly*, 36 (3). 679-702.
23. Angelika Dimoka. 2012. How to Conduct a Functional Magnetic Resonance (fMRI) Study in Social Science Research. *MIS Quarterly*, 36 (3). 811-840.

24. Angelika Dimoka, Paul A Pavlou and Fred D Davis. 2011. Research Commentary—NeuroIS: The Potential of Cognitive Neuroscience for Information Systems Research. *Information Systems Research*, 22 (4). 687-702.
25. Angelika Dimoka, Paul A Pavlou and Fred D Davis. 2011. Research Commentary—NeuroIS: The Potential of Cognitive Neuroscience for Information Systems Research. *Information Systems Research*, 22 (4). 687-702.
26. Serge Egelman, Lorrie Faith Cranor and Jason Hong. 2008. You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Florence, Italy, ACM, 1065-1074. 10.1145/1357054.1357219
27. Serge Egelman, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov and Cormac Herley. 2013. Does My Password Go Up to Eleven?: The Impact of Password Meters on Password Selection. in, ACM, 2379-2388. 10.1145/2470654.2481329
28. Steven Furnell. 2010. Usability Versus Complexity—Striking the Balance in End-user Security. *Network Security*, 2010 (12). 13-17.
29. Kalanit Grill-Spector, Richard Henson and Alex Martin. 2006. Repetition and the Brain: Neural Models of Stimulus-specific Effects. *Trends in Cognitive Sciences*, 10 (1). 14-23.
30. Matthew Grizzard, Ron Tamborini, John L. Sherry, René Weber, Sujay Prabhu, Lindsay Hahn and Patrick Idzik. 2015. The thrill is gone, but you might not know: Habituation and generalization of biophysiological and self-reported arousal responses to video games. *Communication Monographs*, 82 (1). 64-87. 10.1080/03637751.2014.971418
31. Phillip M. Groves and Richard F. Thompson. 1970. Habituation: A Dual-Process Theory. *Psychological Review*, 77. 419-450.
32. Marian Harbach, Markus Hettig, Susanne Weber and Matthew Smith. 2014. Using personal examples to improve risk communication for security & privacy decisions *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, ACM, Toronto, Ontario, Canada, 2647-2656.
33. Jennifer J. Heisz and David I. Shore. 2008. More efficient scanning for familiar faces. *Journal of Vision*, 8 (1).
34. Jeffrey L. Jenkins, Bonnie Brinton Anderson, Anthony Vance, C. Brock Kirwan and David Eargle. 2016. More Harm Than Good? How Messages That Interrupt Can Make Us Vulnerable. *Information Systems Research*, 27 (4). 880-896. 10.1287/isre.2016.0644
35. Michael J Kalsher, Michael S Wogalter and Bernadette M Racicot. 1996. Pharmaceutical Container Labels: Enhancing Preference Perceptions with Alternative Designs and Pictorials. *International Journal of Industrial Ergonomics*, 18 (1). 83-90.
36. MJ Kalsher and KJ Williams. 2006. Behavioral Compliance: Theory, Methodology, and Result. in Wogalter, M. ed. *Handbook of Warnings*, Lawrence Erlbaum Associates, Mahwah NJ, 313-331.
37. Chun-Ming Leung. 2009. Depress Phishing by CAPTCHA with OTP. in *3rd International Conference on Anti-Counterfeiting, Security, and Identification in Communication (ASID)*, Hong Kong, IEEE, 187-192.
38. Ajaya Neupane, Nitesh Saxena, Keya Kuruvilla, Michael Georgescu and Rajesh Kana. 2014. Neural Signatures of User-Centered Security: An fMRI Study of Phishing, and Malware Warnings *The Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, 1-16.
39. Catharine H. Rankin. 2009. Introduction to Special Issue of Neurobiology of Learning and Memory on Habituation. *Neurobiology of Learning and Memory*, 92 (2). 125-126.
<http://dx.doi.org/10.1016/j.nlm.2008.09.010>
40. Catharine H. Rankin, Thomas Abrams, Robert J. Barry, Seema Bhatnagar, David F. Clayton, John Colombo, Gianluca Coppola, Mark A. Geyer, David L. Glanzman, Stephen Marsland, Frances K. McSweeney, Donald A. Wilson, Chun-Fang Wu and Richard F. Thompson. 2009. Habituation revisited: An updated and revised description of the behavioral characteristics of habituation. *Neurobiology of Learning and Memory*, 92 (2). 135-138.
<http://dx.doi.org/10.1016/j.nlm.2008.09.012>
41. René Riedl, Fred D. Davis and Alan R. Hevner. 2014. Towards a NeuroIS Research Methodology: Intensifying the Discussion on Methods, Tools, and Measurement. *Journal of the Association for Information Systems*, 15 (10). i-xxxv.
42. René Riedl, Adriane B Randolph, Jan vom Brocke, Pierre-Majorique Léger and Angelika Dimoka. 2010. The Potential of Neuroscience for Human-Computer Interaction Research. in *Proceedings of the 9th Annual Workshop on HCI Research in MIS*, 1-5.
43. Ruth Rosenholtz, Yanzhen Li and Lisa Nakano. 2007. Measuring visual clutter. *Journal of vision*, 7 (2). 17.
44. Christina M Rudin-Brown, Mike P Greenley, Andrea Barone, Joe Armstrong, Alice F Salway and Beverley J Norris. 2004. The Design of Child Restraint System (CRS) Labels and Warnings Affects Overall CRS Usability. *Traffic Injury Prevention*, 5 (1). 8-17.
45. Jennifer D. Ryan, Robert R. Althoff, Stephen Whitlow and Neal J. Cohen. 2000. Amnesia is a deficit in relational memory. *Psychological Science*, 11 (6). 454-461.
46. Florian Schaub, Rebecca Balebako, Adam L Durity and Lorrie Faith Cranor. 2015. A Design Space for Effective Privacy Notices. *Symposium on Usable Privacy and Security*, Ottawa.

47. S. E. Schechter, R. Dhamija, A. Ozment and I. Fischer. 2007. The Emperor's New Security Indicators *IEEE Symposium on Security and Privacy, 2007. SP '07*, 51-65.
48. Stuart E Schechter, Rachna Dhamija, Andy Ozment and Ian Fischer. 2007. The Emperor's New Security Indicators. in *Security and Privacy, 2007. SP'07. IEEE Symposium on*, Berkeley, CA, IEEE, 51-65.
49. Susanne Schmid, Donald A. Wilson and Catharine H. Rankin. 2014. Habituation Mechanisms and Their Importance for Cognitive Function. *Frontiers in Integrative Neuroscience*, 8. 97. 10.3389/fnint.2014.00097
50. Russell J Sojourner and Michael S Wogalter. 1997. The Influence of Pictorials on Evaluations of Prescription Medication Instructions. *Drug Information Journal*, 31 (3). 963-972.
51. Evgeny Sokolov. 1963. Higher Nervous Functions: The Orienting Reflex. *Annual Review of Physiology*, 25. 545-580.
52. Andreas Sotirakopoulos, Kirstie Hawkey and Konstantin Beznosov. 2011. On the challenges in usable security lab studies: lessons learned from replicating a study on SSL warnings *Proceedings of the Seventh Symposium on Usable Privacy and Security*, ACM, Pittsburgh, Pennsylvania, 1-18.
53. Nathalie Stembert, Arne Padmos, Mortaza S Bargh, Sunil Choenni and Frans Jansen. 2015. A Study of Preventing Email (Spear) Phishing by Enabling Human Intelligence. in *Intelligence and Security Informatics Conference (EISIC), 2015 European*, IEEE, 113-120.
54. Joshua Sunshine, Serge Egelman, Hazim Almuhiemedi, Neha Atri and Lorrie Faith Cranor. 2009. Crying Wolf: An Empirical Study of SSL Warning Effectiveness, *USENIX Association*, 399-416.
55. Richard F. Thompson. 2009. Habituation: A History. *Neurobiology of Learning and Memory*, 92 (2). 127-134. <http://dx.doi.org/10.1016/j.nlm.2008.07.011>
54. Joshua Sunshine, Serge Egelman, Hazim Almuhiemedi, Neha Atri and Lorrie Faith Cranor. 2009. Crying Wolf: An Empirical Study of SSL Warning Effectiveness, *USENIX Association*, 399-416.
56. Anthony Vance, Bonnie Brinton Anderson, C. Brock Kirwan and David Eargle. 2014. Using Measures of Risk Perception to Predict Information Security Behavior: Insights from Electroencephalography (EEG). *Journal of the Association for Information Systems*, 15 (10). 679-722.
57. Vinod Venkatraman, Angelika Dimoka, Paul A. Pavlou, Khoi Vo, William Hampton, Bryan Bollinger, Hal E. Hershfield, Masakazu Ishihara and Russell S. Winer. 2015. Predicting Advertising Success Beyond Traditional Measures: New Insights from Neurophysiological Methods and Market Response Modeling. *Journal of Marketing Research*, 52 (4). 436-452. doi:10.1509/jmr.13.0593
58. Michael S Wogalter. 2006. Communication-Human Information Processing (C-HIP) Model. in Wogalter, M.S. ed. *Handbook of Warnings*, Lawrence Erlbaum Associates, Mahwah, NJ, 51-61.
59. J. H. Xiong, J. H. Gao, J. L. Lancaster and Peter T Fox. 1995. Clustered pixels analysis for functional MRI activation studies of the human brain. *Human Brain Mapping*, 3 (4). 287-301.