Eyes Glazed Over: Using Eye Tracking and fMRI to Measure Habituation to Warnings over a Workweek

Abstract

A major inhibitor of the effectiveness of security warnings is habituation: diminished attention due to frequent exposure to warnings. Although this problem is widely recognized, previous security studies have largely inferred or indirectly measured the occurrence of habituation. Moreover, although habituation develops over time, previous studies have examined habituation only within a single experimental session. It therefore remains unclear how habituation to security warnings evolves over longer periods of time.

We address this gap by conducting a longitudinal experiment that examines how habituation to security warnings develops over the course of a five-day workweek. In addition, we measure the occurrence of habituation using two neurophysiological methods simultaneously: fMRI and eye tracking. Our results show a dramatic drop in attention in terms of neural activity and eye fixations after only the second exposure to a warning, with further decreases throughout the workweek. We also find that participants' attention partially recovers between workdays when there was no exposure to the warning stimulus. Finally, as a potential cost-effective measure to mitigate habituation, we test a polymorphic warning that updates its appearance with each repetition. We find that such warnings are substantially more resistant to habituation across the workweek as compared to conventional warnings.

1 Introduction

Users often represent the last line of defense between attackers and organizations. User response to security warnings is thus a critical aspect of behavioral security [23]. A major inhibitor of the effectiveness of security warnings is habituation: diminished attention due to frequent exposure to warnings [31]. Through this process—also known as warning blindness [50] or fatigue [2]—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., 18; 29; 35; 44], few studies have examined habituation empirically. Moreover, the few empirical studies that do exist either infer habituation or measure it indirectly through a behavioral proxy [2; 9; 10; 29; 45; 51]. An exception is [4], which used fMRI to examine how habituation to warnings develops in the brain.

Another major limitation of prior research is that it is based on cross-sectional experimental designs. However, habituation is fundamentally a neurobiological phenomenon that evolves over time [37]. Therefore, past research on habituation to security warnings has provided only a static snapshot of a dynamic problem. Our first research question is therefore:

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

We addressed this question by extending the work of [4] 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 experimental design allowed us to measure not only the attenuation of user warning response over the course of the workweek, but also another core characteristic of habituation: *response recovery*, that is, the increase in user response after a rest period in which the stimulus is absent [37]. Given that past work has been based exclusively on cross-section experimental designs, this paper is the first to explore how users recover from habituation effects between exposures to warnings.

Secondly, [4] used fMRI and mouse cursor tracking to measure habituation. However, neither of these methods directly measures visual processing. fMRI measures cognitive activity that lags visual processing by 3–5 seconds. Mouse cursor tracking provides a surrogate measure of attention as the mouse cursor hovers over UI elements, but this too follows the eyes' inspection of visual elements, and only provides an incomplete view of attention. 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. This allowed us to measure how both cognitive processing and visual inspection of a warning habituate over time.

Our second research question is:

RQ2. *How can security warnings be designed to be more resistant to habituation over time?*

Previously, [4] showed that polymorphic warnings that repeatedly change their appearance can be effective in maintaining attention during a single experimental session, but 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 results showed a dramatic drop in attention in terms of neural activity and eye fixations after only the second exposure to a warning, with further decreases throughout the workweek. We also found that participants' attention partially recovered between workdays when the stimulus was absent. Interestingly, we found that the polymorphic warning design was substantially more resistant to habituation as compared to conventional static warnings, and that this advantage persisted throughout the five-day experiment. The polymorphic design may thus be a cost-effective solution which can easily be put into practice.

2 Literature Review

Habituation is widely recognized as "the simplest and most basic form of learning" [36, 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" [13; 37, p. 125]. Habituation is an important survival mechanism because it allows organisms to filter out irrelevant stimuli in the environment, and to thus conserve energy for response to stimuli which are relevant for survival [46]. Not surprisingly, humans also exhibit habituation to a wide variety of stimuli—visual, auditory, and others—and this is evident as early as infancy [14].

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 examine it. For example, Egelman et al. [23] found a correlation between user disregard for warnings and user recognition of warnings as previously viewed, and attributed this correlation to habituation. Sunshine et al. [51] 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, and noted that this result was "indicative of warning fatigue" [2, p. 268].

Bravo-Lillo et al. [8; 9] empirically measured habituation, albeit indirectly. For example, 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 [9]. A follow-up study examined four different levels of warning exposure frequency. They found that increasing the frequency with which a dialog was displayed caused a threefold decrease in the proportion of users who immediately recognized a change in the dialog message [8].

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 found further decreases with additional exposures.

However, all of these prior studies share a major limitation in common: they are based on single experimental sessions. This is a problem because habituation is a fundamentally neurobiological process that occurs over time [37]. Consequently, cross-sectional experiments that observe habituation at a single moment in time are unable to capture how habituation evolves over several days. Furthermore, they are unable to measure response recovery after a warning stimulus has been withheld. As a result, our understanding of how habituation evolves and of how to address the problem is limited. This is the primary research gap addressed in this study.

3 Hypotheses

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 in which the stimulus is absent [37]. 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 upon two prominent theories of habituation from neurobiology: the stimulus-model comparator theory (SMCT) [49] and the dual-process theory (DPT) [28]. Although their mechanisms differ, both models describe a consistent process of habituation (see Figure 1).

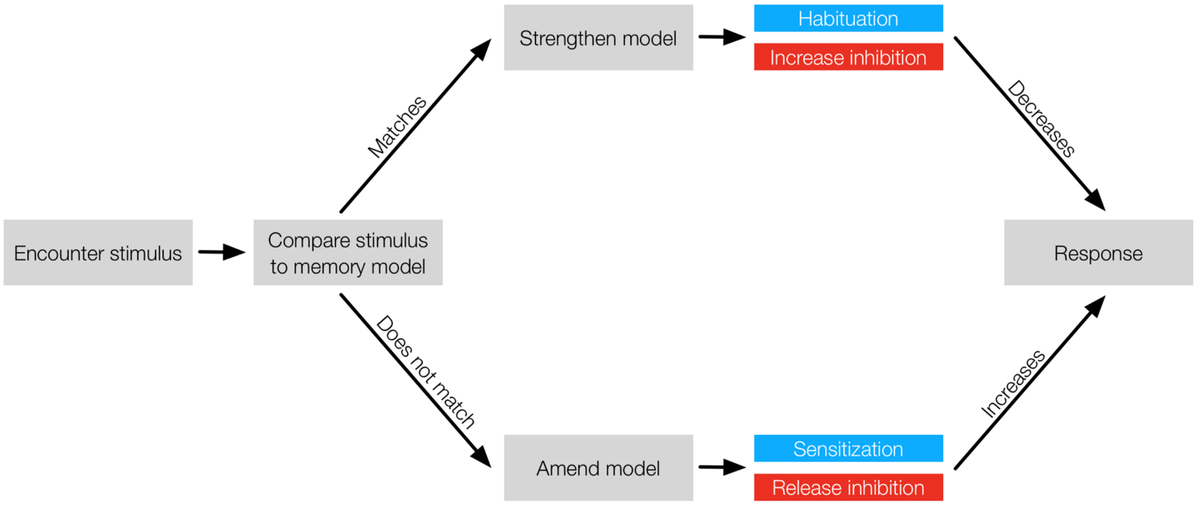

Figure 1: Graphical depiction of the stimulus-model comparator theory (unique terminology in red) and dual-process theory (unique terminology in blue).

3.1 Response Decay

SMCT [49] explains that the brain creates a mental model when exposed to a stimulus (e.g., when seeing a warning). When people see the same stimulus again, they automatically and unconsciously compare the stimulus to this model. If the model and stimulus are similar, an blocking system in the brain inhibits behavioral responses to the stimulus—e.g., people pay less attention to the stimulus [52].

DPT [28] describes this reduced response to stimuli as habituation. In the context of security warnings, users unconsciously compare subsequent warnings to the mental model of warnings they have seen previously. If users unconsciously determine that a warning is similar to others they have seen before, they pay less attention to it. This automatic, subconscious 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 [27]. When viewing repeated warnings within a single computing session, the brain creates a robust mental model of the security message, which results 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 [37]. 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" [56, p. 55]. Changing the appearance of a warning creates novelty. The *orienting reflex,* described by SMCT 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 response strength will recover—(e.g., people will pay more attention to the warning) [49]. DPT describes this process as *sensitization,* an energizing process that strengthens the orienting reflex and thereby the attention span [28]. Sensitization counterbalances habituation [37]. Consequently, by changing the appearance of a warning, users' orienting reflexes are unconsciously sharpened, and thus users will habituate less to polymorphic warnings on both the neural and behavioral levels [5].

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). [15; 52]. In summary, we hypothesize:

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

3.2 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 [6] explains that memory becomes weaker due to the mere 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 [11].

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 [37]. In summary, we hypothesize:

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

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 that have not received as much reinforcement) fade more quickly than stable models [37]. Thus, 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 one day to the next, it is even more likely to differ from the existing mental model, thus weakening behavioral inhibition, increasing sensitization, and enhancing response recovery [37]. Conversely, with static warnings, response recovery will be weaker because the mental model is more robust, reinforced by repetitive exposures to the same warning on previous days [26; 28]. The behavioral response will be inhibited to a greater degree, and habituation will be more pronounced [37]. 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.

4 Polymorphic Warning Design

Anderson et al. [4], developed a polymorphic warning artifact based on an extensive review of the warning-science literature. They created 12 graphical variations of a warning dialog that was expected to sustain attention. Using fMRI data, they tested the different polymorphic variations and found that, in terms of maintaining attention, four of the variations performed better than the rest: (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 for 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" [38, 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" [40, p.250]. Further, Dimoka et al. [22] 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.

5 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 [55]. 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.

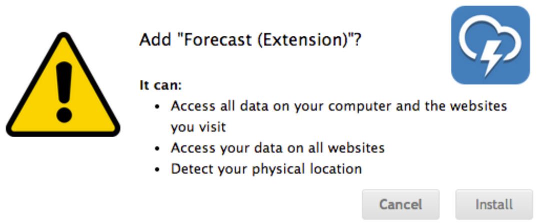

Add "Forecast (Extension)"?

It can:

- Access all data on your computer and the websites you visit
- Access your data on all websites
- Detect your physical location

Message Content: Pictorial symbols (e.g., an exclamation point) [32; 48]

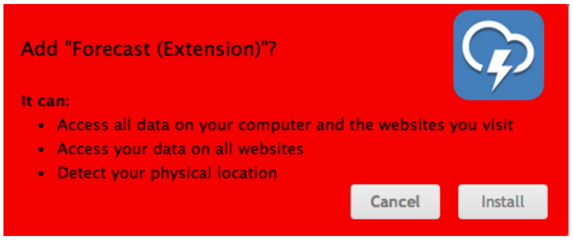

Add "Forecast (Extension)"?

It can:

- Access all data on your computer and the websites you visit
- Access your data on all websites
- Detect your physical location

Warning Appearance: Color [7; 42]

Animation: Jiggle, scale/zoom [9; 24; 33]

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

fMRI identifies regions in terms of voxels or small 3 mm cubes, which makes it ideal when high spatial resolution is required [20]. 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 [26]. In our case, high spatial resolution was important because it allowed us to disentangle RS effects from sensory adaptation or fatigue effects [37].

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 [26]. 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 [43]. 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 before. With repeated exposure, the memories become increasingly available, thus requiring less visual sampling of an image [30].

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.

5.1 Participants

We recruited 16 participants from a large US university (eight male, eight female). This number of participants is consistent with other fMRI studies [21]. 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, 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.

5.2 Ethics

The university Institutional Review Board (IRB) approved the protocols used. Upon arrival at the facility, participants completed a screening form to ensure MRI compatibility. Participants were verbally briefed about MRI procedures as well as the task and purpose of the experiment before entering the scanner.

5.3 Experiment Design

Our experimental design (Figure 3) consisted of five steps. In Step 1, computer-security warning images

---

[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.

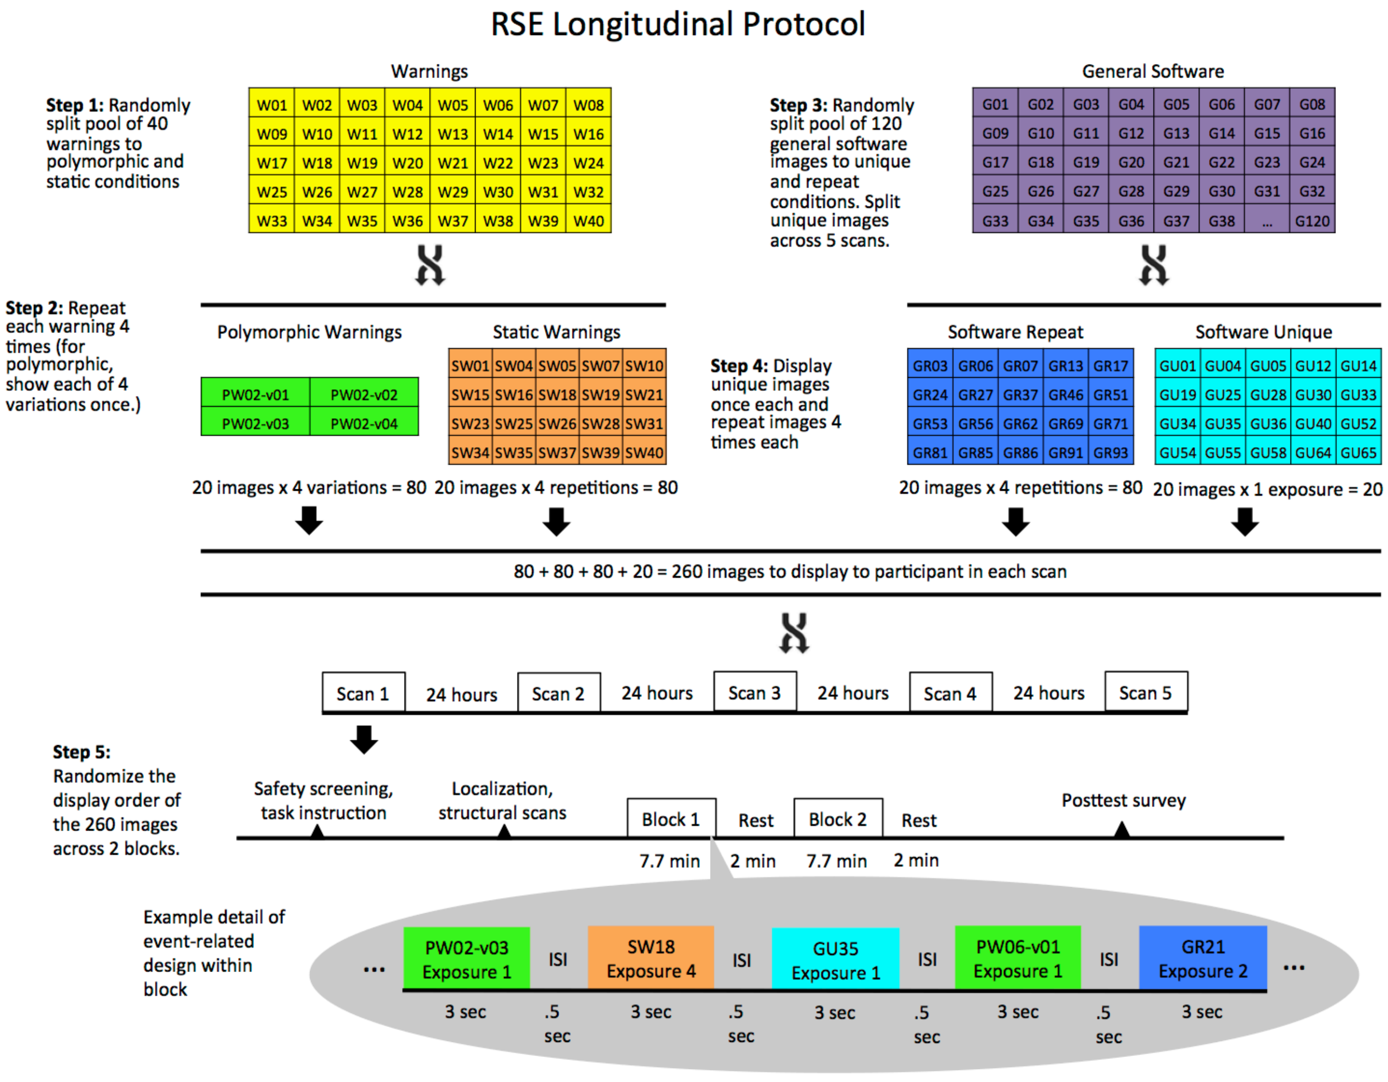

Figure 3. fMRI repetition-suppression-effect (RSE) longitudinal protocol.

were randomly split into two pools: one for the static condition and the other for the polymorphic condition. In Step 2, warnings in the polymorphic pool were randomly assigned to one of the four variations depicted in Figure 2, with the order of polymorphic variations also randomized. In Step 3, general software images were randomly split into two sets of images. In Step 4, 20 of the images in the first set were shown four times each, whereas the other 20 images were displayed only once. These unique general software 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.

Overall, there were 260 images, randomized for each participant, across two blocks of 7.7 minutes each, with a two-minute break in between blocks. Images were displayed for 3 seconds each, with a 0.5-second interstimulus interval. The technical details of the fMRI scans and procedures are documented in the appendix.

6 Analysis

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

6.1 fMRI Analysis

MRI data was analyzed using the Analysis of Functional NeuroImages (AFNI) suite of programs [16] (see appendix for details). Whole-brain, multivariate model analyses were conducted on the fMRI data to identify significant clusters of activation, or regions of interest (ROIs), consistent with the hypothesized pattern. All of 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 4 and 5.

6.2 Eye-Tracking Analysis

Eye-tracking data was collected using an MRI-compatible SR Research EyeLink 1000 Plus (see Figure 6). 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]

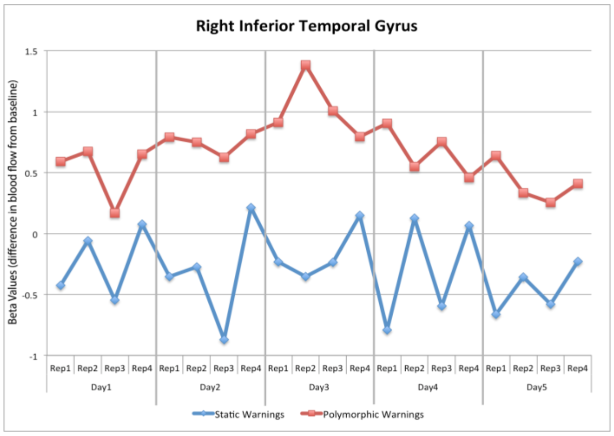

Figure 4. Activity in the right inferior temporal gyrus in response to each presentation of static and polymorphic warnings. Beta values were extracted from a whole-brain analysis for each subject and then averaged across subjects according to stimulus condition.

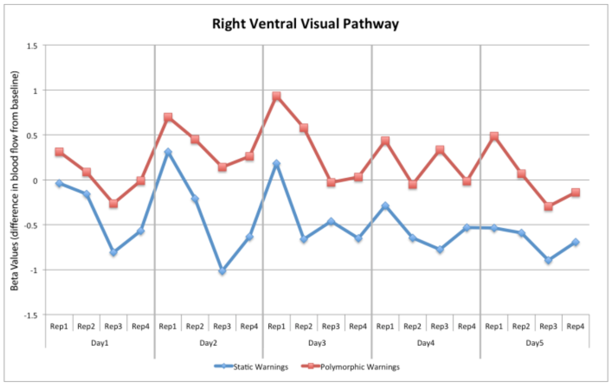

Figure 5. Activity in the right ventral visual pathway in response to each presentation of static and polymorphic warnings. Beta values were extracted from a whole-brain analysis for each subject and then averaged across subjects according to stimulus condition.

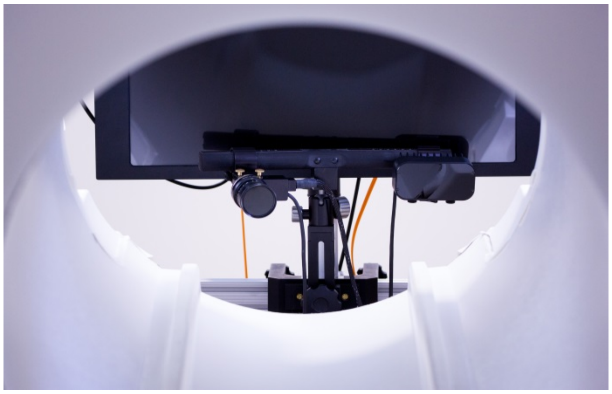

Figure 6. EyeLink 1000 Plus long-range eye tracker, mounted under the MRI viewing monitor.

---

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

The number of fixations for polymorphic and static warnings per warning repetition per day is shown in Figure 7. 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.

6.3 Hypotheses Results

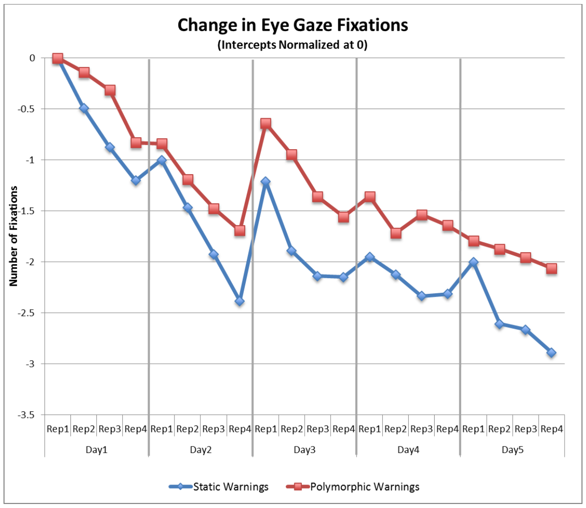

Figure 7. Change in eye gaze fixations across viewings

	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.

6.3.1 H1a Analysis: Users habituate to warnings over consecutive days.

fMRI Analysis: We conducted a whole-brain, multivariate model analysis [12] on the fMRI data holding gender,[3] 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 complexity[4] 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, β = -0.1031], indicating habituation. Visual complexity was also significant [$\chi 2$ (1, N = 11,976) = 34.85, p < .001, β = 0.3815]. The R^2 of the model was 0.13.

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.

---

[3] We controlled for sex because it has been shown to have a significant effect on behavior relating to technology [e.g., 25; 39; 54].

[4] A MATLAB script was used to calculate visual complexity [41].

6.3.2 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 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, β = 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, β = -0.115] and polymorphism [$\chi2$ (1, N = 11,976) = 64.71, p < .001, β = -0.725] were also significant. Visual complexity, however, was not significant: χ^2 (1, N = 11,976) = 0.17, p > .05, β = 0.026. The R^2 of the model was .137.

6.3.3 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. 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 significantly positive recovery (m = 0.369, sd = 3.171) from day to day [t(2377) = 5.672, p < .001, d = 0.233].

6.3.4 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, the 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 [χ^2 (1, N = 2,400) = 1.92, p > .05, β = -0.166] nor the visual complexity [χ^2 (1, N = 2,400) = 1.16, p > .05, β = 0.072] significantly predicted recovery between days. Table 4 summarizes the results for all hypotheses.

					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.

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

7 Discussion

Our results caution against the overuse of warnings. We found that users rapidly habituate to warnings—that is, after only a few exposures. In our study, user response decay was the result not of carelessness or inattentiveness, but of the neural process of habituation. Thus, our findings echo the message that users are not the enemy [1]. Instead, security interventions should be designed to operate within the bounds of human limitations. In this study we have demonstrated that one such limitation is habituation.

Although past studies have examined how habituation influences users response to warnings, they have not examined how habituation influences response to warnings across time [e.g., 2; 4; 9; 10; 29; 45; 51]. Rather, past studies have examined habituation to security warnings in cross-sectional experimental designs. In this study, we have extended this research by examining how people habituate to warnings over a 5-day workweek, providing a more complete understanding of how habituation evolves over time. We found that habituation begins to occur after only the second exposure to a warning, and increases throughout the workweek.

Our longitudinal design also allowed us to examine response recovery, and whether or not there is an increase in response to a security warning after the warning is withheld for a time. Response recovery is a common characteristic of habituation [37], but 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. This supports the idea that habituation can be effectively mitigated by allowing time between displaying warnings.

Furthermore, this study provides a more complete measure of habituation than past studies, and can thus aid in the future evaluation of warning designs to deter habituation. Security warnings are primarily graphical UI elements. However, no research yet exists on how the visual processing of security warnings changes over time and how this corresponds to established fMRI measures of habituation. This paper marks the first attempt to study habituation by measuring both fMRI and eye-tracking data simultaneously. We found that the eye tracking results closely mimic the fMRI results, suggesting that eye tracking is a cost-effective alternative to fMRI for studying habituation to warnings. This finding will enable future researchers to conduct more ecologically valid habituation studies that use eye tracking in a normal computing environment.

Importantly, we showed that the process of habituation over the course of a workweek can be mitigated by changing a warning's appearance. Past studies have shown that polymorphic warnings are effective in reducing habituation in a cross-sectional setting [4]. Here we have extended this work by showing polymorphic warnings do not lose their novelty over time. Rather, our results demonstrate that polymorphic warnings continue to be more resistant to habituation over a 5-day period as compared to static warnings.

We also identified specific regions of the brain that indicate habituation and recovery in relation to security warnings. This contribution deepens our understanding of security interventions by "mapping them into brain areas with existing functional or neurological connotations from the cognitive neuroscience literature," and could "become the means for assessing the performance of IT designs and help guide the design of future IT systems" [22, pp. 691, 694]. In our study, we found that habituation to security warnings was implicated in areas associated with visual processing (bilateral visual stream) and attentional control (the fronto-parietal network). In this manner, we showed that habituation to security warnings affects multiple areas and functions of the brain simultaneously.

Finally, our results highlighted the benefit of applying neuroscience theories and methods to the domain of information security. Although much valuable research has investigated psychological processes that lead to insecure behavior [e.g., 19; 34; 47], our research suggests a more fundamental explanation of why users sometimes behave insecurely: human biology. That is, users may behave insecurely because secure behaviors conflict with natural neurobiological processes. This suggests that future work should apply neurophysiological tools to design security interventions that better accord with human neurobiology.

8 Limitations and Future Research

Our research is subject to a number of limitations. First, although fMRI provides neural insights into the process of habituation, it also introduces artificiality into the experiment. For example, participants viewed 260 warnings during 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 [3; 9]. 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. [3; 9]

Secondly, our tests measured habituation of attention to warnings, rather than how habituation affects actual security behavior [53]. 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 Anderson et al. [4], which demonstrated that their polymorphic warning design was effective in reducing habituation behavior, as measured by mouse cursor tracking.

Third, 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 signficant 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.

9 Conclusion

Past research has found that security warnings are frequently ignored, rendering systems vulnerable to security threats. In this study we have explained how habituation contributes to the failure of security warnings. 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 user response to security messages. This study also demonstrates that neurophysiological tools are useful in understanding the biology of the user as it relates to the design of security interventions.

Appendix: fMRI and Eye-Tracking Experimental 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 × .977 × .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 × 3 × 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 was 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.

We first performed a 10-second localizer scan, followed by a seven-minute structural scan. Following these scans, we started the experimental task. We used SR Research Experiment Builder software to display the stimuli and to synchronize the display events and scanner software. The total scan time was 26.6 minutes for each day. Upon completion of the scan on Days 1–4, participants were thanked and reminded of the next day's scan. At the end of the 5-day period, participants were again thanked, debriefed, and given $60 compensation. All ex-post tests revealed that no subjects needed to be excluded (e.g., due to abnormalities or excessive movement).

Task

The task used an event-related design. Stimuli consisted of four different categories: static warnings, polymorphic warnings, unique general software images, and repeated general software images. The polymorphic and static warnings were created by randomly splitting a pool of 40 warnings between the two categories, which resulted in 20 images in each. Static warnings were repeated during the task four times per day. Polymorphic warnings also appeared four times for each warning, but each appearance displayed a different variation of the warning (i.e., the window jiggled, the window scaled, the window's changed color, or the symbol in the window changed). Both static and polymorphic warnings were presented, for a total of 80 trials within each scan session (20 images × 4 repetitions = 80 trials).

General software images (e.g., control panel, installation windows) were also displayed to compare RS of warnings to RS of general software, as well as to provide a baseline measure of unique presentations. A total of 120 software images were randomly divided into two sets: unique and repeated general software images. The repeated general software images consisted of 20 images that were repeated four times each day. The warning images were displayed in a similar manner, resulting in a total of 80 trials. The remaining 100 images were further divided into five groups of 20 images, each in the unique image group, in order to serve as a baseline condition; all activations were portioned out based on these trials. One group of 20 images was displayed each day, and each image was displayed a single time. In total, 260 images were displayed to each participant during a scan session.

Image size was normalized to subtend approximately 8.5° of visual angle on the images' longest axis. Each image was displayed for three seconds, with an intertrial interval of 0.5 seconds. The 260 images were randomly divided into two blocks of 7.7 minutes each, with a two-minute break in between each block to reduce participant fatigue.

Engagement Check

To ensure that participants were attentive to the task in the scanner, they were instructed to rate the severity of the content of each item as it was presented to them; the answer choices available to them were "extremely severe," "somewhat severe," "somewhat not severe," or "extremely not severe." Answers were given at any time during each trial by pressing a button on an MRI-compatible button box. Following the guidance of Dimoka [21], we performed two checks to ensure that participants were engaged in the task. First, we explored whether participants ranked each stimulus on the severity scale or ignored the ranking. We found that participants ranked stimuli 99.8% of the time, which is a strong indicator of engagement. Second, we explored whether participants ranked the security warnings as more severe than the software prompts, which would suggest that participants were giving thoughtful responses. A t-test indicated that participants did indeed report that the security warnings (m = 2.998, sd = 1.478) were more severe than the software prompts [(m = 2.366, sd = 1.826), t (13463) = 25.236, p < .001, d = 0.435].

fMRI Data Analysis Details

Functional data was 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 lowfrequency motion. A three-dimensional automated image registration routine, 3dVolreg [17], 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 [57].

REFERENCES

[1] ADAMS, A. and SASSE, M.A. Users are not the enemy. *Communications of the ACM 42*, 12 (1999), 40-46.

[2] AKHAWE, D. and FELT, A.P. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Proceedings of the 22nd USENIX Conference on Security* (2013).

[3] AMER, T.S. and MARIS, J.-M.B. Signal words and signal icons in application control and information technology exception messages—hazard matching and habituation effects. *Journal of Information Systems 21*, 2 (2007), 1-25.

[4] ANDERSON, B.B., KIRWAN, C.B., JENKINS, J.L., EARGLE, D., HOWARD, S. and VANCE, A. How polymorphic warnings reduce habituation in the brain: Insights from an fMRI study. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015).

[5] BARRY, R.J. Habituation of the orienting reflex and the development of preliminary process theory. *Neurobiology of Learning and Memory 92*, 2 (2009), 235-242.

[6] BERMAN, M.G., JONIDES, J. and LEWIS, R.L. In search of decay in verbal short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition 35*, 2 (2009), 317-333.

[7] BRAUN, C.C. and SILVER, N.C. Interaction of signal word and colour on warning labels: Differences in perceived hazard and behavioural compliance. *Ergonomics 38*, 11 (1995), 2207-2220.

[8] BRAVO-LILLO, C., CRANOR, L., KOMANDURI, S., SCHECHTER, S. and SLEEPER, M. Harder to ignore? Revisiting pop-up fatigue and approaches to prevent it In (2014).

[9] BRAVO-LILLO, C., KOMANDURI, S., CRANOR, L.F., REEDER, R.W., SLEEPER, M., DOWNS, J. and SCHECHTER, S. Your attention please: Designing security-decision UIs to make genuine risks harder to ignore. In *Proceedings of the Ninth Symposium on Usable Privacy and Security* (2013).

[10] BRUSTOLONI, J.C. and VILLAMARÍN-SALOMÓN, R. Improving security decisions with polymorphic and audited dialogs. In (2007).

[11] ÇEVIK, M.Ö. Habituation, sensitization, and pavlovian conditioning. *Frontiers in Integrative Neuroscience 8* (2014), 13.

[12] CHEN, G., ADLEMAN, N.E., SAAD, Z.S., LEIBENLUFT, E. and COX, R.W. Applications of multivariate modeling to neuroimaging group analysis: A comprehensive alternative to univariate general linear model. *NeuroImage 99* (2014), 571-588.

[13] CHRISTOFFERSEN, G.R.J. 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 (1997), 45-66.

[14] COLOMBO, J. and MITCHELL, D.W. Infant visual habituation. *Neurobiology of Learning and Memory 92*, 2 (2009), 225-234.

[15] COOKE, S.F., KOMOROWSKI, R.W., KAPLAN, E.S., GAVORNIK, J.P. and BEAR, M.F. Visual recognition memory, manifested as long-term habituation, requires synaptic plasticity in v1. *Nature Neuroscience 18*, 2 (2015), 262-271.

[16] COX, R.W. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research 29*, 3 (1996), 162-173.

[17] COX, R.W. and JESMANOWICZ, A. Real-time 3d image registration for functional MRI. *Magnetic Resonance in Medicine, 42* (1999), 1014-1018.

[18] CRANOR, L.F. A framework for reasoning about the human in the loop. In (2008).

[19] D'ARCY, J. and HERATH, T. A review and analysis of deterrence theory in the IS security literature: Making sense of the disparate findings. *European Journal of Information Systems 20*, 6 (2011), 643-658.

[20] DIMOKA, A. How to conduct a functional magnetic resonance (fMRI) study in social science research. *MIS Quarterly 36*, 3 (2012), 811-840.

[21] DIMOKA, A., BANKER, R.D., BENBASAT, I., DAVIS, F.D., DENNIS, A.R., GEFEN, D., GUPTA, A., LSCHEBECK, A., KENNING, P.H., PAVLOU, P.A., MULLER-PUTZ, G., RIEDL, R., VOM BROCKE, J. and WEBER, B. On the use of neurophysiological tools in IS research: Developing a research agenda for NeuroIS. *MIS Quarterly 36*, 3 (2012), 679-702.

[22] DIMOKA, A., PAVLOU, P.A. and DAVIS, F.D. Research commentary-NeuroIS: The potential of cognitive neuroscience for information systems research. *Information Systems Research 22*, 4 (2011), 687-702.

[23] EGELMAN, S., CRANOR, L.F. and HONG, J. 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* (2008).

[24] FURNELL, S. Usability versus complexity—striking the balance in end-user security. *Network Security 2010*, 12 (2010), 13-17.

[25] GEFEN, D. and STRAUB, D.W. Gender differences in the perception and use of e-mail: An extension to the technology acceptance model. *MIS Quarterly 21*, 4 (1997), 389-400.

[26] GRILL-SPECTOR, K., HENSON, R. and MARTIN, A. Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences 10*, 1 (2006), 14-23.

[27] GRIZZARD, M., TAMBORINI, R., SHERRY, J.L., WEBER, R., PRABHU, S., HAHN, L. and IDZIK, P. 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 (2015), 64-87.

[28] GROVES, P.M. and THOMPSON, R.F. Habituation: A dual-process theory. *Psychological review 77* (1970), 419-450.

[29] HARBACH, M., HETTIG, M., WEBER, S. and SMITH, M. Using personal examples to improve risk communication for security & privacy decisions. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems* (2014).

[30] HEISZ, J.J. and SHORE, D.I. More efficient scanning for familiar faces. *Journal of Vision 8*, 1 (2008).

[31] KALSHER, M. and WILLIAMS, K. Behavioral compliance: Theory, methodology, and result. In *Handbook of warnings*. Lawrence Erlbaum Associates, Mahwah NJ, 313-331, 2006.

[32] KALSHER, M.J., WOGALTER, M.S. and RACICOT, B.M. Pharmaceutical container labels: Enhancing preference perceptions with alternative designs and pictorials. *International Journal of Industrial Ergonomics 18*, 1 (1996), 83-90.

[33] LEUNG, C.-M. Depress phishing by CAPTCHA with OTP. In *Proceedings of the Third International Conference on Anti-counterfeiting, Security, and Identification in Communication (ASID 2009)* (2009).

[34] LIANG, H. and XUE, Y. Avoidance of information technology threats: A theoretical perspective. *MIS Quarterly 33*, 1 (2009), 71-90.

[35] NEUPANE, A., SAXENA, N., KURUVILLA, K., GEORGESCU, M. and KANA, R. Neural signatures of user-centered security: An fMRI study of phishing, and malware warnings. In *The Network and Distributed System Security Symposium (NDSS)* (2014).

[36] RANKIN, C.H. Introduction to special issue of neurobiology of learning and memory on habituation. *Neurobiology of Learning and Memory 92*, 2 (2009), 125-126.

[37] RANKIN, C.H., ABRAMS, T., BARRY, R.J., BHATNAGAR, S., CLAYTON, D.F., COLOMBO, J., COPPOLA, G., GEYER, M.A., GLANZMAN, D.L., MARSLAND, S., MCSWEENEY, F.K., WILSON, D.A., WU, C.-F. and THOMPSON, R.F. Habituation revisited: An updated and revised description of the behavioral characteristics of habituation. *Neurobiology of Learning and Memory 92*, 2 (2009), 135-138.

[38] RIEDL, R., DAVIS, F.D. and HEVNER, A.R. Towards a NeuroIS research methodology:

Intensifying the discussion on methods, tools, and measurement. *Journal of the Association for Information Systems 15*, 10 (2014), i-xxxv.

[39] RIEDL, R., HUBERT, M. and KENNING, P. Are there neural gender differences in online trust? An fMRI study on the perceived trustworthiness of ebay offers. *MIS Quarterly 34*, 2 (2010), 397-428.

[40] RIEDL, R., RANDOLPH, A.B., VOM BROCKE, J., LÉGER, P.-M. and DIMOKA, A. The potential of neuroscience for human-computer interaction research. In *Proceedings of the 9th Annual Workshop on HCI Research in MIS* (2010).

[41] ROSENHOLTZ, R., LI, Y. and NAKANO, L. Measuring visual clutter. *Journal of Vision 7*, 2 (2007), 1-22.

[42] RUDIN-BROWN, C.M., GREENLEY, M.P., BARONE, A., ARMSTRONG, J., SALWAY, A.F. and NORRIS, B.J. The design of child restraint system (CRS) labels and warnings affects overall CRS usability. *Traffic Injury Prevention 5*, 1 (2004), 8-17.

[43] RYAN, J.D., ALTHOFF, R.R., WHITLOW, S. and COHEN, N.J. Amnesia is a deficit in relational memory. *Psychological Science 11*, 6 (2000), 454-461.

[44] SCHAUB, F., BALEBAKO, R., DURITY, A.L. and CRANOR, L.F. A design space for effective privacy notices. *To appear in the* (2015).

[45] SCHECHTER, S.E., DHAMIJA, R., OZMENT, A. and FISCHER, I. The emperor's new security indicators. In *IEEE Symposium on Security and Privacy, 2007. SP '07* (2007).

[46] SCHMID, S., WILSON, D.A. and RANKIN, C.H. Habituation mechanisms and their importance for cognitive function. *Frontiers in Integrative Neuroscience 8* (2014), 97.

[47] SIPONEN, M. and VANCE, A. Neutralization: New insights into the problem of employee information systems security policy violations. *MIS Quarterly 34*, 3 (2010), 487-502.

[48] SOJOURNER, R.J. and WOGALTER, M.S. The influence of pictorials on evaluations of prescription medication instructions. *Drug Information Journal 31*, 3 (1997), 963-972.

[49] SOKOLOV, E. Higher nervous functions: The orienting reflux. *Annual Review of Physiology 25* (1963), 545-580.

[50] STEMBERT, N., PADMOS, A., BARGH, M.S., CHOENNI, S. and JANSEN, F. 2015. A study of preventing email (spear) phishing by enabling human intelligence. In *Intelligence and Security Informatics Conference (EISIC), 2015 European* IEEE, 113-120.

[51] SUNSHINE, J., EGELMAN, S., ALMUHIMEDI, H., ATRI, N. and CRANOR, L.F. Crying wolf: An empirical study of SSL warning effectiveness. In (2009).

[52] THOMPSON, R.F. Habituation: A history. *Neurobiology of Learning and Memory 92*, 2 (2009), 127-134.

[53] VANCE, A., ANDERSON, B.B., KIRWAN, C.B. and EARGLE, D. Using measures of risk perception to predict information security behavior: Insights from electroencephalography (EEG). *Journal of the Association for Information Systems 15*, 10 (2014), 679-722.

[54] VENKATESH, V. and MORRIS, M.G. Why don't men ever stop to ask for directions? Gender, social influence, and their role in technology acceptance and usage behavior. *MIS Quarterly 24*, 1 (2000), 115-139.

[55] VENKATRAMAN, V., DIMOKA, A., PAVLOU, P.A., VO, K., HAMPTON, W., BOLLINGER, B., HERSHFIELD, H.E., ISHIHARA, M. and WINER, R.S. Predicting advertising success beyond traditional measures: New insights from neurophysiological methods and market response modeling. *Journal of Marketing Research 52*, 4 (2015), 436-452.

[56] WOGALTER, M.S. Communication–human information processing (C-HIP) model. In *Handbook of warnings*, M.S. WOGALTER Ed. Lawrence Erlbaum Associates, Mahwah, NJ, 51-61, 2006.

[57] XIONG, J.H., GAO, J.H., LANCASTER, J.L. and FOX, P.T. Clustered pixels analysis for functional MRI activation studies of the human brain. *Human Brain Mapping 3*, 4 (1995), 287-301.