

**From Warning to Wallpaper:
Why the Brain Habituates to Security Warnings**

Anderson, B., Vance, A., Kirwan, B., Jenkins, J., Eargle, D. 2016. "From Warning to Wallpaper: Why the Brain Habituates to Security Warnings and What Can Be Done about It," *Journal of Management Information Systems*, 33 (3), pp. 713–743, <http://dx.doi.org/10.1080/07421222.2016.1243947>.

Abstract

Warning messages are fundamental to users' security interactions. Unfortunately, research has shown that they are largely ineffective. A key contributor to this failure is habituation: decreased response to a repeated warning. Previous research has inferred the occurrence of habituation to warnings or measured it indirectly, such as through the proxy of a related behavior. Therefore, there is a gap in our understanding of how habituation to security warnings develops in the brain. Without direct measures of habituation, we are limited in designing warnings that can mitigate its effects.

In this study, we use neurophysiological measures to directly observe habituation as it occurs in the brain and behaviorally. We also design a polymorphic warning artifact that repeatedly changes its appearance in order to resist the effects of habituation. In an experiment using functional magnetic resonance imaging (fMRI; $n = 25$), we found that our polymorphic warning was significantly more resistant to habituation than were conventional warnings in regions of the brain related to attention.

In a second experiment ($n = 80$), we implemented the top four most resistant polymorphic warnings in a realistic setting. Using mouse cursor tracking as a surrogate for attention to unobtrusively measure habituation on participants' personal computers, we found that polymorphic warnings reduced habituation compared to conventional warnings. Together, our findings reveal the substantial influence of neurobiology on users' habituation to security warnings and security behavior in general, and we offer our polymorphic warning design as an effective solution to practice.

Keywords: Security warnings, habituation, behavioral information systems security, functional magnetic resonance imaging (fMRI), mouse cursor tracking, NeuroIS

INTRODUCTION

Warning messages are one of the last lines of defense in computer security and are fundamental to users' security interactions with technology. Unfortunately, experimental research has consistently shown that they are largely ineffective [25, 74]. This is problematic, as users are increasingly targeted by attackers seeking to gain access to the information resources of organizations. Consequently, researchers have actively sought to understand how users interact with security warnings and why warnings are so pervasively ignored [9].

A key contributor to the disregard of security warnings is habituation, which is “decreased response to repeated stimulation” [79, p. 419]. Through this phenomenon, warnings that were once salient become virtually unnoticeable, like familiar wallpaper. Although habituation has been inferred as a factor in many security-warning studies [e.g., 29, 68], little research has specifically investigated habituation in the context of warnings. Further, because habituation is difficult to directly observe using conventional methods, research that does investigate habituation has done so indirectly by observing its influence on security-related behaviors rather than by measuring habituation itself [e.g., 10, 12]. Therefore, there is a gap in our understanding of how habituation to security warnings occurs in the brain, in addition to behavioral habituation. Without direct measures of habituation, researchers are limited in their efforts to design warnings that can mitigate its effects.

This study addresses this gap by using NeuroIS—the use of neurophysiological tools to explore how the functioning of the brain affects IS-related behaviors—to open the “black box” of the brain [26] in order to observe habituation as it develops in response to security warnings. In doing so, we respond to the call of Crossler *et al.* [18] for the application of NeuroIS methods to yield fresh insights into information-security behaviors. Specifically, we point to the *repetition suppression* effect, the reduction of neural responses to stimuli that are viewed repeatedly [36], as the neural manifestation of habituation to stimuli [45]. We also draw on mouse cursor movement indicators of attention as a behavioral–attentional manifestation of habituation to stimuli. By investigating

behavioral–attentional indicators of habituation and how repetition suppression occurs in the brain, we can take a more precise approach to the design of security warnings that are resistant to the effects of habituation.

Accordingly, the research objective of this paper is twofold: (1) to directly observe habituation as it occurs in the brain and is manifested in behavioral-attentional indicators using neurophysiological tools, and, using these measures, (2) to design a security-warning artifact that is substantially more resistant to habituation than were the conventional warnings. Our research questions to pursue these objectives are:

RQ1. How does habituation occur in the brain in response to security warnings?

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

We investigated these questions in two laboratory experiments using complementary neurophysiological measures: functional magnetic resonance imaging (fMRI; Experiment 1) and mouse cursor tracking (Experiment 2). This follows the guidance of Dimoka et al. that “no single neurophysiological measure is usually sufficient on its own, and it is advisable to use many data sources to triangulate across measures” [28, p. 694]. In Experiment 1, we designed a polymorphic warning artifact that repeatedly changes its appearance as a means of sustaining attention. In doing so, we made 12 graphical variations to the warning derived from findings in the warning science literature [86]. Next, we performed an fMRI experiment to evaluate our polymorphic warning design and found that it was substantially more resistant to habituation than were conventional warnings. Further, we used the fMRI results to identify the graphical variations most resistant to habituation.

In Experiment 2, we implemented our polymorphic warning artifact in the Google Chrome browser as part of a realistic task to enhance the ecological validity of our results. This study explored how participants responded to Google Chrome permission warnings on their personal laptops. In addition, we used mouse cursor tracking to measure changes in attention over time due to

habituation. We found that our polymorphic warning resulted in reduced habituation when compared with the standard permission warning, corroborating the fMRI results of Experiment 1.

This paper is organized as follows: In the next section, we review the literature relating to security warnings, habituation, and NeuroIS methods for measuring habituation. We then discuss the hypotheses, experimental procedures, and analyses for Experiments 1 and 2, respectively. Finally, we present a general discussion of our findings for both experiments and summarize the implications.

LITERATURE REVIEW

Given its strong security implications, habituation is frequently cited as a key contributor to users' failure to heed warnings [38, 61]. However, most studies infer the presence of habituation, rather than empirically examine it. For example, Egelman et al. [29] found a correlation between user disregard for warnings and user recognition of warnings as previously viewed, and attributed this correlation to habituation. 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], in other words, habituation.

Those few studies that do empirically examine habituation to security warnings have done so indirectly. For example, Brustoloni and Villamarín-Salomón [12] found that compared to conventional warnings, a security warning that randomized the position of its option buttons resulted in users ignoring the message less frequently in risky situations. Bravo-Lillo et al. [10] 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 percent of the users in their study immediately recognized the change in the dialog message.

From the studies reviewed above, it is clear that despite habituation to security warnings being widely acknowledged as a problem, little empirical research has specifically examined this area of information security. Moreover, the lack of direct measures of habituation limits our ability to design security warnings that are resistant to habituation. Therefore, measuring the occurrence of

habituation in the brain presents an opportunity to quantify the antecedents and onset of habituation and its impacts on security-warning behavior. This is consistent with the recommendation of vom Brocke and Liang [82], who emphasize the importance of selecting NeuroIS research questions that, first and foremost, answer problems of importance to the IS community, and, secondly, benefit from studies using neurophysiological measures. With this objective in mind, we next review the application of NeuroIS to the study of behavioral information security.

In the brain, a diminished neural response to a repeated stimulus is called repetition suppression, a robust manifestation of habituation, and has been observed in humans across a variety of brain regions and experimental situations [for review, see 36, 69]. The exact cause of repetition suppression may vary according to brain region. In the visual processing system—a brain network highly relevant to the processing of visual warnings¹—repetition suppression appears to be related to repetition priming and may reflect the facilitation of stimulus processing with repeated exposures [24]. This type of response is highly influenced by the similarity of the repeated stimuli, where a more robust suppression is associated with stimuli that share a greater number of similarities [53].

Interestingly, research has found that neural measures of repetition suppression can be more sensitive to habituation than behavioral measures are. For example, Motley and Kirwan [59] demonstrated that the hippocampus differentiates subtle changes to stimuli (in this case, stimuli rotated by 15°) even when participants' behavioral responses indicate that they are unaware of the changes. Therefore, research shows that the repetition suppression effect is an effective and sensitive way to assess the underlying neural computational processes of behavior [e.g., 48, 87]. In this study, we map the sensitivity of different brain regions to stimulus repetition of security warnings.

¹ See <http://neurosynth.org/analyses/terms/visual>; https://wikipedia.org/wiki/Visual_cortex.

Mouse Cursor Tracking: A Behavioral–Attentional View of Habituation

Mouse cursor tracking has been proposed as a cost-effective and unobtrusive instrument to measure users' attention in an HCI context [e.g., 31, 60]. Mouse cursor movements can be used to “infer user attention in complex web pages containing images, text and varied content” [60, p. 2693]. For example, research has shown that eye-gaze and cursor movement patterns are highly correlated [20, 46]. Freeman et al. suggested that the movements of the hand “offer continuous streams of output that can reveal ongoing dynamics of processing, potentially capturing the mind in motion with fine-grained temporal sensitivity” [31, p. 1]. Accordingly, hundreds of recent studies have chosen mouse cursor tracking as a methodology for studying various cognitive and emotional processes. See Appendix C for examples of how mouse cursor tracking has been used to measure attention.

Importantly, mouse cursor tracking can measure changes in attention [60, 65] and therefore can be used to measure diminished attention in response to repeated warnings, the form of habituation relevant to our context [42]. Because mouse cursor tracking can be measured unbeknownst to users on their own laptops (through JavaScript embedded in a website), it offers a superior method for unobtrusively measuring habituation in users' natural environments compared to other methodologies (e.g., fMRI, eye tracking, EEG, etc.). In this paper, we use both fMRI and mouse cursor tracking to provide a more holistic view of habituation to security warnings.

THEORY AND HYPOTHESIS DEVELOPMENT

In this section, we first provide an overview of habituation theory that we will use to help support our hypotheses. We then introduce our polymorphic warning design. Finally, we discuss two different sets of hypotheses: fMRI hypotheses and mouse cursor tracking hypotheses. These hypotheses are tested in Experiment 1 and Experiment 2, respectively.

Overview of Habituation Theory

Two prominent complementary theories explain the process of habituation: the stimulus-model comparator theory (SMCT) [73], and the dual-process theory DPT [37]. SMCT [73] claims that when

people see a repeated stimulus, they form a mental model of the stimulus. As a new stimulus is observed, people compare the incoming stimulus with the mental model. If the model and stimulus are similar, the brain's amplifying system inhibits the behavioral response, resulting in a response decrement—an attenuation of the response. However, if the incoming stimulus does not match the model, an orienting reflex occurs [5]. In this situation, the brain's amplifying system will release the inhibition and the response strength will increase.

DPT proposes that when people see a stimulus, one of two processes will occur: habituation or sensitization [63]. When people encounter a stimulus, the brain compares the stimulus to any existing mental representations of the stimulus. If the stimulus and a model are similar, the habituation process occurs, resulting in a response decrement. However, if the stimulus is novel, sensitization occurs, which counterbalances habituation and increases one's response to the stimuli. We draw on the common theoretical elements of the mental model comparisons to support our hypotheses [80]. proposes that when people see a stimulus, one of two processes will occur: habituation or sensitization [63]. When people encounter a stimulus, the brain compares the stimulus to any existing mental representations of the stimulus. If the stimulus and a model are similar, the habituation process occurs, resulting in a response decrement. However, if the stimulus is novel, sensitization occurs, which counterbalances habituation and increases one's response to the stimuli. We draw on the common elements of both of these theories to support our hypotheses [80].

Polymorphic Warning Design

One of our research objectives was to design a security-warning artifact substantially more resistant to habituation than conventional warnings. Several approaches have been suggested to increase attention to security warnings and thus reduce habituation. For example, following design guidelines found in the warning literature, Bravo-Lillo et al. [10, 11] designed and tested a variety of warning *attractors*, which are graphical elements that draw attention to salient information in warning dialogs. These studies demonstrated that habituation could be reduced through UI design.

In a different approach, Wogalter stated that “habituation can occur even with well-designed warnings, but better designed warnings with salient features can slow the habituation process. Where feasible, changing the warning’s appearance may be useful in reinvigorating attention switch previously lost because of habituation” [85, p. 55]. In line with this reasoning, Brustoloni and Villamarín-Salomón [12] proposed *polymorphic warnings*, which change appearance to reduce habituation. In contrast, warnings are typically *static*, meaning they do not change appearance on repeated exposures. However, they only tested changing the order of options on the warning dialog, acknowledging that “the design space for polymorphic dialogs is vast” [12, p. 4]. In addition, it is not clear whether polymorphic warnings actually reduce habituation, as they did not measure habituation directly. Thus, our research objectives are to (1) use fMRI to determine whether a polymorphic warning design is able to reduce habituation in the brain in response to repeated warning exposures, and (2) identify the polymorphic variations that are most successful at reducing habituation.

Development

We relied on the warning science literature to develop 12 graphical variations expected to capture attention. Our polymorphic warning artifact rotated through the graphical variations on each subsequent exposure. These variations are described in Table 1 and depicted in Figure 1.

Table 1. Polymorphic variations and their support from the literature	
Text Appearance	Support
Color of text	Braun et al. [7], Laughery et al. [49]
Highlighting of text	Strawbridge [77], Young and Wogalter [88]
Message Content	Support
Pictorial symbols (e.g., an exclamation point)	Kalsher et al. [44], Sojourner and Wogalter [72]
Signal word (e.g., “warning” or “danger”)	Kalsher et al. [43], Silver and Wogalter [71]
Warning Appearance	Support
Color	Braun and Silver [8], Rudin-Brown et al. [66]
Contrast (e.g., white on black)	Sanders and McCormick [67], Young [89]
Ordering of options	Brustoloni and Villamarín-Salomón [12], De Keukelaere et al. [22]
Size	Vigilante and Wogalter [81], Wogalter and Vigilante [84]
Animation	Support
Jiggle, Scale/Zoom, Twirl/Spin	Bravo-Lillo et al. [10], Furnell [32], Leung [50]



Figure 1. Polymorphic warning design variants.

Experiment 1. Hypotheses

According to theory on habituation (see the section titled “Overview of Habituation Theory”), when users repeatedly see a warning, they will generate a robust mental model of the warning across exposures. On each successive exposure to the warning, users will rely more on the mental model of the warning as opposed to the actual warning presented, resulting in a response decrement to the actual warning [37, 73]. Accordingly, our first objective is to examine the nature of the repetition suppression effect in neural responses when users view repeated presentations of security warning stimuli. Once we establish the neural signature of habituation to security warnings in general, we will be able to determine if polymorphic warnings deviate from the normal repetition suppression effect.

In a previous study [4], researchers observed robust repetition suppression across a number of brain regions, most prominently in the visual processing system in response to repeated viewings of static computer warning stimuli. Accordingly, in the present study we hypothesize:

H1: For static and polymorphic warnings, attention will decrease in terms of neural activation across repeated exposures.

In the current study, we expect that the polymorphic warnings will maintain high levels of attention across repetitions while static images will be associated with decreases in attentional processing. This will be evident in a number of ways. First, we predict that polymorphic warnings will be more resistant to repetition suppression than static images, especially in regions associated with visual spatial attention (e.g., regions in the frontoparietal attention network) [16]. Second, we predict that regions associated with inattention will be *more* active with repeated exposures to static warnings when compared to polymorphic warnings. The default-mode network² is a network of brain regions that is more active when participants are allowed to rest and let their thoughts wander [13]. The default-mode network has been shown to be more active when participants are not attending to

² <http://neurosynth.org/analyses/terms/default>; https://wikipedia.org/wiki/Default_mode_network.

the presented stimulus [55]. If polymorphic stimuli are effective at maintaining participants' attention, then default-mode network activity should be relatively low compared to activity for static stimuli with repeated exposures. Accordingly, we expect that polymorphic warnings will be associated with increased activity in attention networks and decrease activity in the default-mode network, but static warning images will be associated with the opposite effect. The presence of an interaction between stimulus repetition and stimulus type would indicate that polymorphic warnings are more resistant to the repetition suppression effect than static warning images. Our hypothesis is therefore:

H2: Attention will decrease in terms of neural activation less for polymorphic than for static warnings across repeated exposures.

This will be evident in two ways: first, repeated presentations of static warnings will result in greater *decreases* in activation in regions associated with *attention* compared to presentations of polymorphic warnings. Second, repeated presentations of static warnings will result in greater *increases* in activation in regions associated with *inattention* compared to presentations of polymorphic warnings.

Finally, we expect that animated polymorphic warnings will exhibit a differential neural response relative to still (non-animated) polymorphic warnings and will be more resistant to habituation. This is because animation should activate regions of the brain associated with visual motion processing, such as the human motion complex (hMT+) area³ [35], in addition to brain regions that process the content of still images. Additionally, early visual processing areas, such as the primary visual cortex, are responsive to motion [78]. Thus, animated polymorphic warnings should exhibit an additive effect over that of still polymorphic warnings.

³ hMT+ or human middle temporal complex; <http://neurosynth.org/analyses/terms/motion>; https://en.wikipedia.org/wiki/Visual_modularity.

H3: Attention will be greater in terms of neural activation for animated polymorphic warnings than for non-animated polymorphic warnings.

In other words, there will be a higher neural response upon successive viewings for animated polymorphic warnings as compared with non-animated polymorphic warnings in regions that demonstrate repetition suppression, in addition to greater neural activation in motion-sensitive regions, such as the hMT+.

Experiment 2. Hypotheses

The purpose of Experiment 2 is to test the findings of Experiment 1 in a more realistic context, enhancing the ecological validity of the study as a whole using mouse cursor movements to measure habituation. Mouse cursor movements can be used as an unobtrusive surrogate of users' attention. For example, mouse cursor movements are highly sensitive to stimuli and information that, even briefly, capture users' attention [60]. Attention can be measured through various movement statistics including greater area under the curve, slower average speed, and slower initial acceleration (see methodology section for a description of these statistics). This relationship between attention and mouse cursor movements has been validated in various settings [e.g., 20, 83].

With repeated exposures to security warnings, users will rely more on the mental model of the warning as opposed to the actual warning presented, resulting in a decreased response [37, 73]. This response decrement includes paying less attention to the warning, and responding to the warnings more reactively and automatically, which will influence mouse cursor movements in several ways. First, as information on the warning captures users' attention less, mouse cursor movements will deviate less toward that information [60], moving more directly to execute a learned movement response—a predetermined movement response to the warning (e.g., dismissing the warning) [70]. Further, when users rely on the mental model more as opposed to giving attention to the warning, movements will reactively begin to execute the learned response, which results in faster

acceleration and speed of movements [54, 70]. These mouse-movement indicators of lower attention will amplify over repeated exposures to the warning—an indicator of habituation. In summary, we hypothesize:

H4: For both static and polymorphic warnings, attention will decrease in terms of mouse cursor movements (greater area under the curve, slower average speed, and slower initial acceleration) across repeated exposures.

However, polymorphic warnings should garner more attention than static warnings in general, resulting in higher mouse cursor movement indicators of attention for polymorphic warnings. When a warning's appearance is updated, this causes the warning to differ from the existing mental model of warnings in general. As a result, users will experience an orientation reflex that will cause them to pay more attention to the warning [73]. In other words, the novelty of the warning causes sensitization that draws users' attention to the warning [63]. This will result in greater deviation of mouse cursor movements from directly responding to the warning as users examine the different parts of the warning to determine what is new and view the warning's message [60]. Further, as users examine the warning, their movements will not reactively respond to the warning as much, neither accelerating nor moving as quickly [54, 70]. In summary, we predict:

H5: Attention will be greater in terms of mouse cursor movements (greater area under the curve, slower average speed, and slower initial acceleration) for polymorphic warnings than for static warnings.

As users continue to see polymorphic warnings, they will also begin to habituate to them (see H4) [84]. However, we hypothesize that this habituation will occur more slowly for polymorphic warnings than for static warnings. By continually changing the appearance of the warning, the mental model of the polymorphic warning will be less stable compare to the mental models of seeing repeated static warnings. Thus, when users see a subsequent polymorphic warning, it is less likely to

match its mental model, resulting in novelty and less habituation over time [73]. Further, as it continues to change its appearance, the polymorphic warning will further conflict with its mental model, causing sensitization that counters or slows habituation over time [63]. In summary, in addition to polymorphic warnings garnering more attention in general (H5), we predict that the decrease in attention due to habituation over time will be less for polymorphic warnings than for static warnings:

H6: Attention will decrease less in terms of mouse cursor movements (greater area under the curve, slower average speed, and slower initial acceleration) for polymorphic warnings than for static warnings across repeated exposures.

METHODOLOGY

To test our hypotheses, we conducted two experiments. Experiment 1 tested the fMRI-related hypotheses (H1–H3). Experiment 2 tested the mouse-tracking-related hypotheses (H4–H6). We will now describe the methodology for each experiment.

Experiment 1: fMRI

Measures

Functional MRI is a method of choice in decision neuroscience because of its superior ability to identify areas of the brain that are activated during decision making and other behavioral tasks [27, 64]. Although certain neurophysiological tools, such as EEG [21, 47, 51] or eye-tracking [57], offer greater temporal resolution for examining habituation, these methods cannot examine activity in specific brain regions as fMRI can. Functional MRI measures neural activity indirectly by tracking changes in the level of blood oxygenation, which are driven by changes in the metabolic demands of active neural populations. This phenomenon is known as the *blood oxygen level dependent* (BOLD) effect, and its magnitude is proportional to the degree of underlying neural activation [52]. By measuring the BOLD effect, researchers can both identify distinct regions of the brain where activity

is correlated with specific emotions (e.g., fear or uncertainty) and cognitive processes (e.g., perception or memory retrieval), and evaluate the degree of activation in these regions. Thus, fMRI is well suited to the investigation of the neural underpinnings of how habituation to security warnings occurs in the brain.

Experimental Design

We used an event-related, within-subject experimental design in which we compared the response to polymorphic warning images with the response to static warning images. Our experimental design is graphically depicted in online Appendix A and consists of four steps for each participant. In Step 1, we randomly split a pool of 40 warnings between a polymorphic and static warning treatment. In Step 2, the warnings in the static treatment were repeated 13 times. For the polymorphic treatment, warnings were also repeated 13 times, albeit that, with each repeated exposure, a different polymorphic variation was displayed in random order. For example, the first exposure of a polymorphic warning might have featured a red background, the second exposure might have included a yellow-and-black striped border, and so on. The number 13 was chosen by calculating the maximum number of polymorphic stimuli that we could show to participants while not extending total scan time beyond estimated subject toleration limits of approximately 30 minutes. To be parallel, we included 13 repetitions for static warnings as well. In Step 3, 40 images of general software applications were randomly intermixed with the warning images and displayed one time each. Trials in which these images were presented served as a baseline in the fMRI model. Accordingly, deviations from baseline, or "0", in our fMRI parameter estimates represent deviations from an active task condition (i.e., viewing images) rather than a passive rest condition. We chose this condition as a baseline in order to keep participants engaged in the task and to avoid unconstrained mental activity in the task [see 75]. There were a total of 560 images (20 polymorphic warnings \times 13 variations [for each of the polymorphic warnings] + 20 static warnings \times 13 repetitions [identical, with no variation] + 40 software images \times 1 exposure each) to be displayed in

the experiment. Finally, in Step 4, the 560 images were randomized for each participant across five blocks of 6.5 minutes each (with a 2-minute break in between) in order to relieve participants' fatigue during the tasks. This randomization ensured that any effects observed would be due to the polymorphic treatment, rather than the appearance of the warnings themselves. Further, it ensured an effective jitter between stimuli of similar types and avoided autocorrelation in the experimental design [see 19]. For each stimulus, participants were shown images for 3 seconds each with a 0.5-second interstimulus interval (ISI). Technical details of the fMRI scanner, experimental procedures, and analyses are documented in online Appendix A.

Experimental Task

For each visually presented stimulus, participants used a keypad to indicate if the image shown was (1) *identical* to one seen before the task, (2) *similar* yet different from one seen previously in the task, or (3) *new*, never seen before in the task. This question was used to provide behavioral performance data that ensured that participants were appropriately engaged in the task. Our subsequent analysis of the responses (“identical,” “similar,” and “new”) to each stimulus type (novel, static, and polymorphic manipulations) revealed that participants performed the task as expected. At the conclusion of the fMRI scan, each participant was led to an adjoining room to complete a brief post-test feedback survey. Additionally, to ensure manipulation validity [76], the post-test survey included a manipulation check question that displayed a polymorphic warning as it rotated through its variations. Each participant was asked if he or she noticed the treatment during the task. All but one of the participants reported that they had noticed the experimental treatment, which indicated successful overall manipulation. Following Straub et al., we elected to retain the participant who reported that he was not manipulated to provide “a more robust testing of the hypotheses” [76, p. 408].

Participants

Twenty-five participants were recruited from the university community and were screened for MRI compatibility. Additionally, we screened participants to require native English speakers, corrected-normal visual acuity, and right-handedness and excluded those taking psychotropic medications or with color blindness. Each participant signed an informed consent form in accordance with the IRB protocol. Of the 25 participants, 21 were male and 4 were female. Participants were age 20–27, with a mean of 23.68. We conducted a pilot study that revealed a large estimated effect size for the repetition effect (partial $\eta^2 = 0.7$). Using this estimated effect size, an a priori power analysis indicated that we would need four subjects to achieve power greater than 0.8, indicating that a sample size of 25 is more than adequate [4].

Data Analysis

Our first research question centered on the effect of repetition on measures of neural activation. Accordingly, we performed a repeated-measures ANOVA on the whole-brain fMRI data with repetition number and stimulus type (polymorphic, static) as factors. This allowed us to identify those brain regions that (1) had differential responses to subsequent repetitions (i.e., that displayed a main effect of repetition number in their neural responses as we hypothesized in H1) and (2) had differential responses to subsequent repetitions that were modulated by stimulus type (i.e., that displayed an interaction between repetition number and stimulus type as we hypothesized in H2). Accordingly, for each brain region that displayed a main effect or an interaction (i.e., the regions of interest, or ROIs), we extracted mean parameter estimates for all the voxels in the ROI to perform follow-up repeated-measures analyses, such as examining the direction and strength of any linear trends with repeated presentations.

In our first analysis, we identified 10 regions where there was a significant main effect of repetition number (thresholded with a voxel-wise $p < 0.001$ and spatial extent > 40 contiguous voxels; overall $p < 0.001$). These regions are listed in Table 2. Since these regions were defined as

merely demonstrating a main effect of repetition (i.e., activity during at least two different repetitions differed significantly from each other), we extracted mean parameter estimates from each region to further characterize the nature of the neural responses over repeated stimulus presentations. As noted in Table 2, each region demonstrated a significant linear trend, and in all but one region this linear trend was negative. As an example, Figure 2 depicts the activations in the left and right middle occipital gyrus and the change in activation on the right as a function of repeated stimulus exposures. In regions associated with visual-spatial attention (i.e., left and right occipital lobe and superior parietal lobule), there was a main effect of stimulus type, but no stimulus type by repetition interaction. These findings are consistent with H1, which was that for both polymorphic and static warning images, measures of neural activity would indicate decreased attention with repeated exposures.

Table 2. Regions demonstrating a main effect of repetition and associated inferential statistics

Region	#Voxels	Coordinates			Main Effect Stimulus		Linear Trend Repetition		
		x	y	z	F(1,22)	p	Direction	F(1,22)	p
L. Dorsolateral Prefrontal Cortex	512	50	-8	33	9.56	.005	Negative	42.63	<.001
R. Middle Occipital Gyrus	352	-29	89	12	53.60	<.001	Negative	23.26	<.001
B. Medial Frontal Gyrus	250	2	-17	48	10.62	.004	Negative	37.75	<.001
L. Middle Occipital Gyrus	241	29	89	18	26.57	<.001	Negative	32.68	<.001
R. Fusiform Gyrus	190	-26	56	-4	27.62	<.001	Negative	25.90	<.001
L. Superior Parietal Lobule	153	35	59	54	24.99	<.001	Negative	24.23	<.001
R. Superior Parietal Lobule	98	-32	68	48	31.07	<.001	Negative	22.34	<.001
L. Fusiform Gyrus	71	32	50	-10	25.99	<.001	Negative	18.69	<.001
R. Postcentral Gyrus (ventral)	54	-56	26	15	3.59	.071	Positive	27.44	<.001
L. Postcentral Gyrus (dorsal)	54	47	29	54	10.28	.004	Negative	16.63	<.001

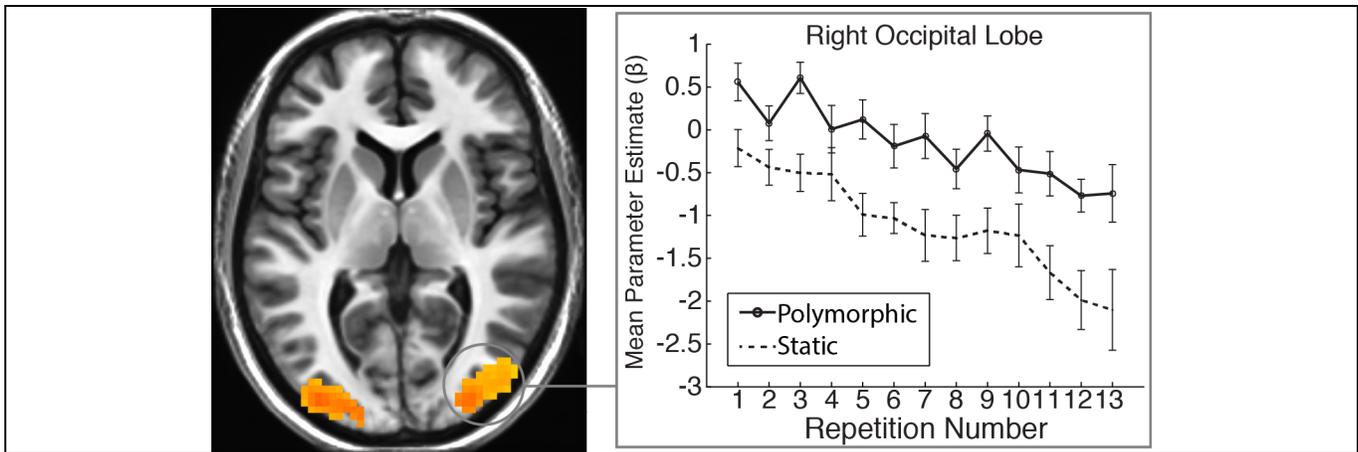


Figure 2. Example regions with fMRI activation changes with subsequent presentations of polymorphic and static warning images. Error bars represent standard error of the mean.

Our next hypothesis concerned the interaction between stimulus repetition and warning type (H2). Accordingly, we first performed a repeated-measures ANOVA on the whole-brain fMRI data with stimulus type (polymorphic, static) and repetition number as factors. We then performed repeated-measures analyses of activations from each region of interest to test our hypotheses. In the whole-brain analysis, we found four regions where there was a stimulus type \times repetition number interaction (Table 3), indicating that these regions reacted differently to repeated exposures to static and polymorphic warnings. Two regions, the left and right superior parietal cortex (Figure 3, upper panels), have been implicated in attentional processing [15]. We extracted mean parameter estimates (betas) for these functionally defined regions of interest to further characterize the nature of the interaction between warning type and repetition number using two-way repeated-measures ANOVAs with warning type (static, polymorphic) and repetition number as factors. In these regions, activation was significantly higher for polymorphic warnings than for static warnings, consistent with sustained

Table 3. Analysis results for warning type x repetition interaction				
	#Voxels	X	Y	Z
Ventral Medial Prefrontal	248	2	-59	12
Left Retrosplenial	51	8	53	18
Left Superior Parietal	45	29	62	54
Right Superior Parietal	42	-38	50	51

attentional processing and reduced repetition suppression for the polymorphic warnings (main effect of stimulus type, left: $F[1,22] = 19.10, p < 0.001$; right: $F[1,22] = 22.69, p < 0.001$).

Consistent with H2, there was sustained attentional processing for the polymorphic warnings with later repetitions as evidenced by a significant repetition by stimulus type interaction on the left ($F(12,264) = 3.07, p < 0.001$) and on the right ($F(12,264) = 3.37, p < 0.001$). The two other significant regions, the bilateral medial prefrontal cortex (MPFC) and the left retrosplenial cortex (Figure 3, lower panels), have both been shown to be major nodes in the default-mode network [13]. In both regions, activation was higher for static images than for polymorphic images (main effect of stimulus type, MPFC: $F(1,22) = 21.54, p < 0.001$; retrosplenial cortex: $F(1,22) = 5.10, p < 0.001$). Critically, this difference in activation emerged in later repetitions, consistent with H2, predicting that there would be a significant interaction between stimulus type and repetition number in the

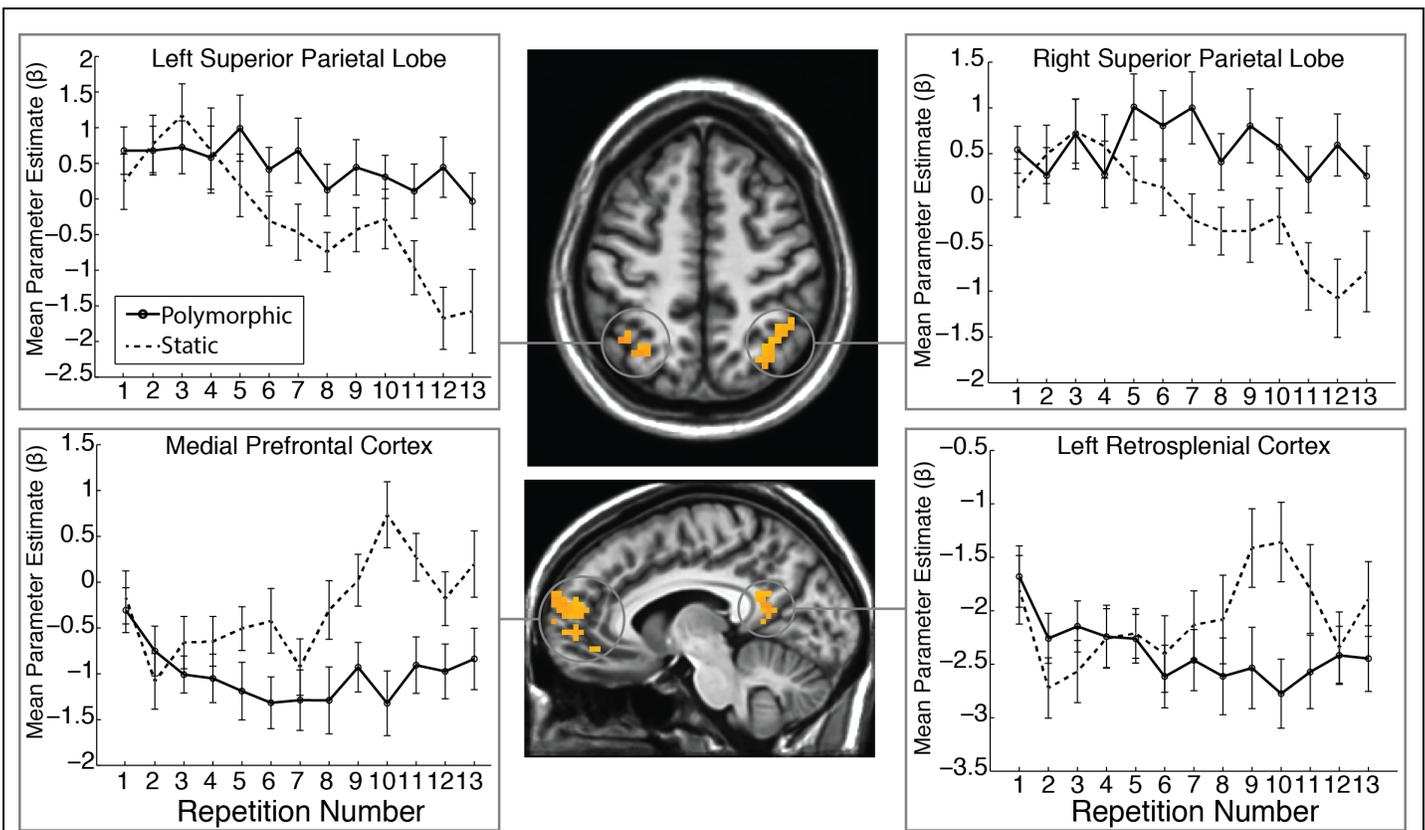


Figure 3. fMRI activation changes with subsequent presentations of polymorphic and static warning images. Error bars represent standard error of the mean.

default-mode network (interaction, MPFC: $F(12,264) = 3.92, p < 0.001$; retrosplenial cortex: $F(12,264) = 3.39, p < 0.001$). Thus, H2 was supported.

An alternate interpretation of these activations is that they reflect changes in risk perception (e.g., [58]) and value processing (e.g., [6]) rather than changes in attentional processing. This interpretation is especially compelling in the decreasing activation patterns in the superior parietal cortex as habituation might make warnings appear less severe and thus less risky. If this interpretation is correct, then the reduction in risk perception may be overcome by polymorphisms. We offer this interpretation cautiously, however, as we did not explicitly manipulate risk nor did we measure subjective risk or value perceptions.

We next performed an analysis to identify regions that exhibited BOLD activation differences between animated and non-animated warnings across all voxels in the brain by performing a *t*-test. Ten clusters were identified where activity significantly differed between these conditions (see Table B1 in online Appendix B). Importantly, and consistent with our third hypothesis (H3), we observed greater activation for animated warnings in the right area hMT+ (Figure B1 in online Appendix B), an area involved in visual motion processing ($t(21) = 5.54, p < 0.001$).

Finally, we examined the specific polymorphic warnings collapsing across presentation order to determine which polymorphic variations were most resistant to the repetition suppression effect. To quantify the most resistant variations for use in Experiment 2, we performed a *t*-test contrasting polymorphic and static warning images. Resulting statistical parameter maps were thresholded using a false-discovery rate of 0.05 and a spatial extent threshold of 1,080 mm³. The largest cluster of activation comprised the bilateral visual processing system (including the occipital, dorsal parietal, and inferior temporal lobes). The mean fMRI activations in this region of interest for each polymorphic variation were then ranked (see Figure 4). The highest scoring variations (“jiggle,” “scale,” “window color,” and “symbol”) were included in Experiment 2.

Summary of Results

As expected, we found a decrease in neural activation in regions associated with visual-spatial attention with repeated exposures to the warning stimuli (H1). Further, we found an interaction in portions of the frontoparietal attention network where fMRI activation remained relatively consistent for polymorphic warning images across repetitions while the activation for static warning images decreased (H2). We also found regions in the medial prefrontal cortex and retrosplenial cortex where activation was higher for later repetitions of static images than for polymorphic images. These additional regions have been commonly observed in the default-mode network, a network of regions that are activated when participants are allowed to engage in non-directed mental activity [34]. Accordingly, this increased activation for later repetitions of the static images but not the polymorphic images may indicate less inattention for the polymorphic images than for the static images (H2).

The animated polymorphic warnings showed the greatest effect. By using animations such as a slight jiggle or zoom upon the appearance of the warning, more brain regions were activated. This

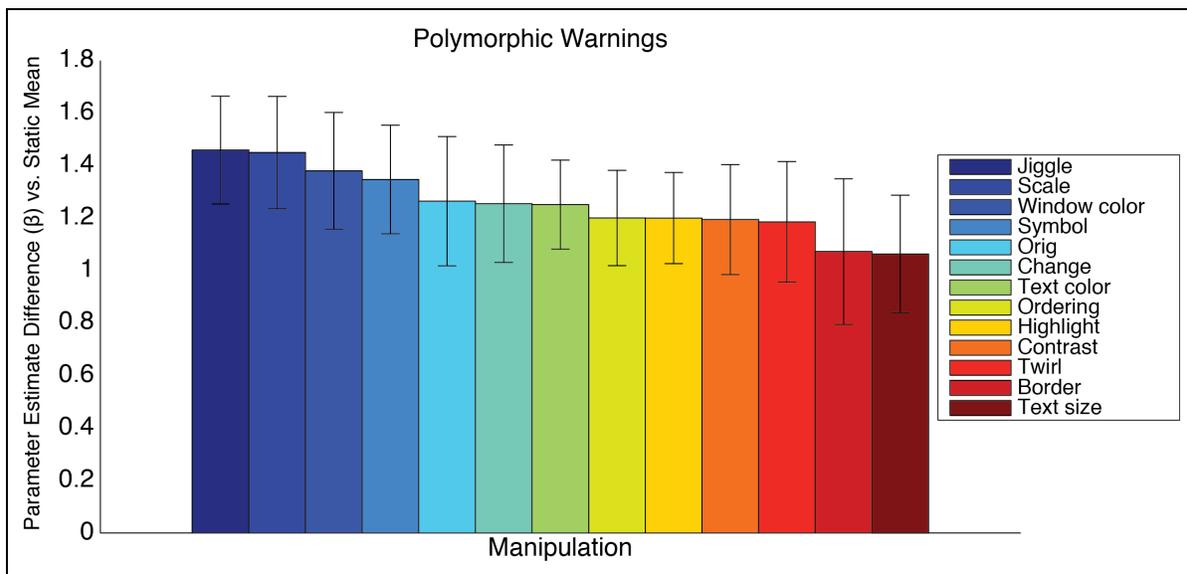


Figure 4. Mean fMRI activation difference for each polymorphic warning variation vs. the mean activation for the static warnings in the region that displayed the greatest difference between polymorphic and static warning images. Error bars represent standard error of the mean.

was true not only in the traditional bilateral visual system associated with still images, but also in the hMT+ region associated with moving images (H3). In summary, brain activity generally remained higher across repetitions for the polymorphic warnings, supporting the efficacy of our polymorphic warning artifact.

Experiment 2: Mouse Cursor Tracking

To improve the ecological validity beyond the limitations of fMRI, we designed a behavioral laboratory experiment in which participants responded to security warnings in the natural course of completing a task on their personal laptops. While users responded to security warnings, mouse cursor movement data (i.e., how a person responded using the laptop touchpad) were collected through embedded JavaScript and stored in an online database (see online Appendix C for a review of mouse cursor tracking). Mouse tracking was used because of its completely unobtrusive nature to enhance ecological validity—that is, mouse activity can be captured on users’ personal computers without additional hardware, software, or configuration [40]. These data were analyzed to explore whether users paid greater attention to the polymorphic warnings over time compared to the static warnings. By using a second neurophysiological methodology, we triangulated the results of the fMRI experiment in Experiment 1 [28].

Measures

We collected mouse cursor movements when the warning was displayed via embedded JavaScript to calculate three indicators of attention: area-under-the-curve (AUC), initial acceleration, and average speed. These measures are detailed in online Appendix E, and summarized in Table 4.

Table 4. Summary of mouse cursor indicators of attention/habituation	
	Indication of Attention / Habituation
AUC	Greater AUC = Greater Attention
Initial Acceleration	Lower Initial Acceleration = Greater Attention
Average Speed	Lower Average Speed = Greater Attention

Experimental Design

To test the hypotheses, we conducted an experiment wherein participants responded to security warnings in the natural course of completing a task on their personal laptops. We selected Google Chrome extension permission warnings as the warning type. Extensions, also referred to as “add-ons” or “plug-ins,” add functionality to the web browser but often require access to data on the computer to perform their functions. For this reason, whenever an extension is installed by a user, Chrome displays a warning that alerts the user to the specific access that the extension requires [33]. For example, an extension that accesses an application-programming interface (API) to detect the browser’s physical location must first raise the permission warning, “Add X (extension)? It can: Detect your physical location.” During the course of the experiment, participants were required to search for and install 20 simulated weather extensions and then evaluate them for usability and aesthetics. Because a permission warning is displayed whenever a Google extension is installed, the participants also received 20 permission warnings. In this way, participants were naturally habituated while completing the task.

We implemented a repeated-measure (participants evaluated 20 simulated weather extensions), between-subject experimental design. Participants were randomly assigned to the polymorphic warning treatment or the static warning treatment (a between-subject manipulation) group. The polymorphic treatment group received polymorphic warnings that changed appearance from one exposure to the next. The control group received conventional Google Chrome permission warnings that were static, meaning they did not change appearance. The warnings are shown in online Appendix D.

Experimental Task

The experiment took place in a behavioral laboratory. To heighten perceived risk, participants were asked to use their personal laptop and were required to read the following disclaimer: “The researchers have not tested all possible weather extensions that you may encounter. The university is

not responsible for any potential malicious software that may be installed as a result of this study. Each participant has the responsibility to use good judgment in deciding which extensions to install and evaluate.” Furthermore, to motivate users to read the warnings, we randomly varied which permissions were being requested in each warning.

Next, participants were asked to use Chrome and search Google.com to find 20 different weather extensions and evaluate them for (1) ease of use, (2) visual appeal, (3) relevancy of content, and (4) their intention to use the extension in the future. However, we performed a man-in-the-middle attack to manipulate the Google search results. If a query related to Chrome weather extensions was shown, the search returned our spoofed results. Each of the links in the spoofed results pointed to sites under our control with legitimate URLs. Our Chrome extension transparently redirected HTTPS traffic to HTTP to avoid SSL⁴ certificate errors. Finally, we caused traffic to the Chrome Web Store to redirect to a “down for maintenance” page. Consequently, participants were required to use the in-line Chrome extension installation process on individual web sites (for technical details of Experiment 2, see online Appendix F).

The manipulated Google search results included more than 40 different weather extensions. Participants were required to evaluate only 20 of these. If a participant received a permission warning requesting unreasonable permissions, the participant was free not to install the extension. Participants had ample time to go back to the Google search results to find an alternative extension.

While responding to each warning, an embedded JavaScript library in the warning captured the users’ cursor movements (the timestamp and x -, y -coordinate of each movement at a millisecond precision rate). These raw cursor movement data were sent to an online database through the Asynchronous JavaScript and eXtensible markup language (AJAX) call and were later used to

⁴ Secure Sockets Layer

calculate the dependent variables for each observation: AUC, initial acceleration, and average speed (calculations are discussed in the Analysis section).

After evaluating 20 extensions and finishing the experiment, each participant completed a post-test survey. The survey captured information about the participant's gender, age, operating system, the browser normally used, and the usual frequency of extension installations. It also asked whether the participant noticed anything unusual about the Domain Name Server (DNS) (a result of our manipulated search results) and queried the participant's perceived risk of ignoring security warnings, perceived severity of being infected by malware, and perceived susceptibility to malware infection for inclusion as control variables in the analysis. Per the approved IRB protocol, all participants were then debriefed, notified of the true purpose of the experiment (exploring the influence of security warnings and habituation on behavior), and told that all of the extensions were sanctioned and did not gather any personal information.

Participants

Eighty subjects participated in the final version of the study. The participants were allowed to visit and revisit a website during the course of the experiment to view the warning for that website's browser extension again. The system recorded if a person tried to reinstall an extension for inclusion in the analysis. As a result, each participant on average encountered slightly more than 20 warnings (exactly 21.2). For the analysis, data were limited to the first 20 warnings for each participant because the number of observations per repetition beyond 20 was too small for reliable analysis (per the instructions, most participants stopped at 20). Because participants were required to use their own laptops, technical limitations on a few of the laptops inhibited data collection through JavaScript. As a result, we had analyzable data from 76 participants across 1,466 warnings (40 in the treatment group and 36 in the control group). An a priori power analysis calculated with G*Power 3.1.7 suggested that a sample size of 1,073 observations was adequate for a small effect size and that 111

Table 5. Descriptive statistics for AUC, average speed, and initial acceleration (raw data was rescaled and transformed)

	Mean	SD	Min	Max
AUC				
Polymorphic	1.556	4.714	0.488	10.000
Static	0.511	0.514	0.134	6.410
Initial Acceleration				
Polymorphic	4.017e-05	3.698e-05	1.740e-07	2.330e-04
Static	8.048e-05	8.802e-05	3.800e-07	6.370e-04
Average Speed				
Polymorphic	8.620e-04	1.049e-03	7.430e-05	1.670e-02
Static	1.524e-03	1.589e-03	6.610e-05	1.729e-02

was adequate for a medium effect size. Thus, our sample size was more than adequate.

Approximately 63 percent of the participants were male, and the average age was 22.35.

Data Analysis

The Lowess curve [14] for the resulting statistics (AUC, initial acceleration, and average speed) was plotted for each treatment across time (see Figure 5). The descriptive statistics for each mouse-movement are displayed in Table 5. We first performed a check for manipulation validity. To explore whether the polymorphic warnings were perceived, we asked participants, “During the experiment, did you notice that some of the warnings changed their appearance like the above image?” (The polymorphic warning rotating through its variations was shown). None of the participants in the control group reported “yes,” and 95 percent of the participants in the treatment group reported “yes.” This difference was significant in an independent sample *t*-test: $t(74) = 25.807, p < 0.001$.

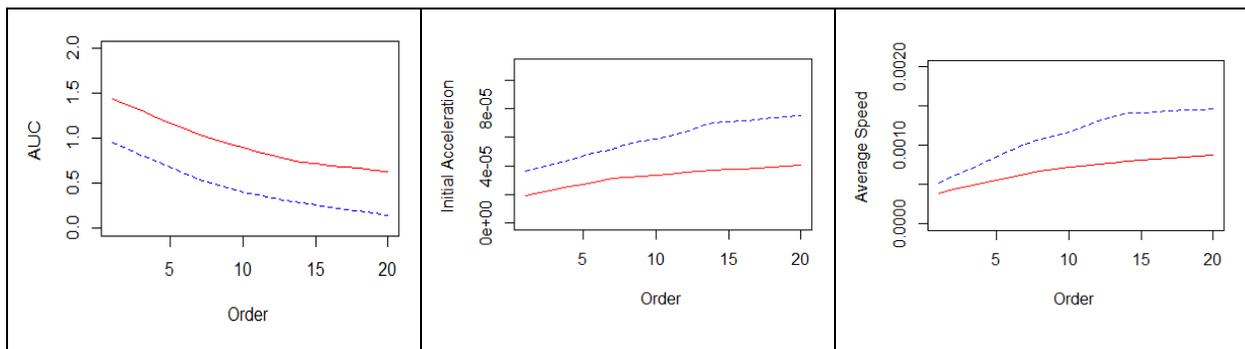


Figure 5. Lowess curve of AUC, average speed, and initial acceleration by order (solid line = polymorphic warning, dotted line = static warning)

After preparing the data and conducting the manipulation checks, we specified a multivariate general linear model to test the hypotheses. The model featured three dependent variables: AUC, acceleration, and speed for each observation ($n = 1,466$). Based on the Lowess curves, we log-transformed these dependent variables to form linear relationships. We included several independent variables. First, we included the warning order number to model habituation across multiple warning exposures, and we included the treatment group to represent polymorphic (coded as 1) and static (coded as 0) conditions. To explore the interaction between habituation and treatments, we included an interaction term for treatment \times order. Next, we controlled for the participant identifier to allow variability across participants. This accounts for other individual differences that may influence mouse cursor movements. Finally, we controlled for whether the warning was being reinstalled (coded as 1) or being installed for the first time (coded as 0)⁵. The results are shown in Table 6. The r^2 for AUC was 0.095, for initial acceleration was 0.224, and for average speed was 0.132.

As a supplemental analysis, we performed a graded motor response analysis [20] in the time course of participants responding to security warnings (see online Appendix G). In summary, this analysis allows us to explore how much users deviated from the idealized response trajectory throughout their response, rather than just an aggregated measure of total deviation at the end, which is the case with AUC. Thus, this supplemental analysis lends understanding in terms of whether warning information sustains greater attention throughout the movement for polymorphic warnings than for static warnings. For brevity, the analysis and results are described in online Appendix E. In summary, corroborating the findings described above regarding AUC, the results of the supplementary analysis provide strong support that deviation was greater for participants in the polymorphic treatment than for participants in the static warning treatment almost throughout the entire movement interaction.

⁵ As participants were allowed to complete the task naturally for ecological validity, they could potentially revisit websites, reinstall extensions, and see the warning for those extensions again.

Table 6. Experiment 2 results

Dependent Variables	Independent Variables	Parameter Estimate	F	p
AUC	(Intercept)	0.001	300.434 (1, 1362)	< 0.001
AUC	Treatment	0.967	64.728 (1, 1362)	< 0.001
AUC	Order	-0.053	14.474 (1,1362)	< 0.001
AUC	isReinstalled	0.689	7.991 (1, 1362)	< 0.01
AUC	Treatment*Order	0.028	32.133 (1, 1362)	< 0.001
AUC	Participant	<i>a</i>	1.041 (79, 1362)	< 0.01
Initial Acceleration	(Intercept)	3.512e-05	1354.441 (1, 1362)	< 0.001
Initial Acceleration	Treatment	-3.782e-6	156.188 (1, 1362)	< 0.001
Initial Acceleration	Order	4.882e-06	136.552 (1, 1362)	< 0.001
Initial Acceleration	isReinstalled	5.364e-05	59.235 (1, 1362)	< 0.001
Initial Acceleration	Treatment*Order	-3.782e-06	9.289 (1, 1362)	< 0.01
Initial Acceleration	Participant	<i>a</i>	1.746 (79, 1362)	< 0.001
Average Speed	(Intercept)	0.001	1275.828 (1, 1362)	< 0.001
Average Speed	Treatment	-0.001	93.289 (1, 1362)	< 0.001
Average Speed	Order	7.528e-05	83.262 (1, 1362)	< 0.001
Average Speed	isReinstalled	0.001	25.129 (1, 1362)	< 0.001
Average Speed	Treatment*Order	-4.064e-05	18.072 (1, 1362)	< 0.001
Average Speed	Participant	<i>a</i>	1.450 (79, 1362)	> 0.05

a Each participant id has its own parameter estimate, adjusting for natural individual differences in mouse movement between participants

Summary of Results

The results of the analyses support the hypotheses. The warning order number had a strong negative correlation with AUC ($p < 0.001$), positive correlation with initial acceleration ($p < 0.001$), and positive correlation with average speed ($p < 0.001$). As habituation increased with subsequent exposures (i.e., with warning order number), H4 was supported for all three indicators. The results also indicated that the use of polymorphic warnings increases AUC ($p < 0.001$), decreases initial acceleration ($p < 0.001$), and decreases average speed ($p < 0.001$), suggesting that users habituate less when polymorphic warnings are shown (H5 supported for all three indicators). In addition, the results suggest that users habituate at a slower rate when polymorphic warnings are shown. The interactions between treatment and order were significant for AUC ($p < 0.001$), initial acceleration ($p < 0.01$), and average speed ($p < 0.001$) (H5 supported for all three indicators).

DISCUSSION

This research makes several contributions, which we summarize below.

Contributions of Experiment 1: fMRI experiment—Polymorphic VS. Static Warnings

First, in Experiment 1, we extended previous research on habituation by using neuroscience methods to observe the neural correlates of habituation as it occurs. While previous research has understood habituation as an automatic response, our research is the first to show empirically how habituation to security warnings occurs in the brain. In doing so, we identified the phenomenon of repetition suppression as a neurobiological explanation for why habituation to security warnings occurs.

Additionally, this study demonstrated an application of measuring the repetition suppression in the brain using fMRI. Whereas previous research measured habituation indirectly by observing its effects, such as inattentive behaviors [10], this study measured neural correlates of habituation directly as it occurs in the brain. Specifically, we showed how using a simple, repeated-exposure experimental design can permit researchers to detect the existence and size of the repetition suppression effect using the BOLD response. Using this method, we illustrated the continued drop in activation in the frontoparietal attention network for static images, but sustained activation for polymorphic images over 13 exposures. These results can provide researchers with a useful baseline of the repetition suppression effect in response to security warnings for future research. Further, these measures may be used to guide the development and testing of security warnings that are resistant to habituation and, thus, lead to safer behavior.

Second, we made an artifactual contribution by designing (in Experiment 1) and implementing (in Experiment 2) a polymorphic warning as a UI artifact. We utilized the warning science literature to derive 12 polymorphic variations that can be generically applied to a wide variety of security warnings. Whereas scholars have criticized IS research for the frequent absence of the IT artifact [3, 62], our polymorphic warning UI design artifact is central to our contribution.

Importantly, we analyzed the fMRI data to improve the UI artifact by identifying the polymorphic variations that were most effective in reducing repetition suppression. In this way, we followed the guidance of Dimoka et al., who suggested that “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” [26, p. 700, p. 700]. By using repetition suppression in the brain as our objective dependent variable, we were able to simplify the polymorphic warning to include only the top four most effective variations, including animated and non-animated forms, making our polymorphic warning more practical to implement.

Third, we demonstrated in Experiment 1 that polymorphic warnings are more resistant to repetition suppression than are static warnings. We theorized that this is because animated warnings activate areas of the brain that perceive motion (hMT+), providing an additive effect over and above the visual processing of non-animated images. Unsurprisingly, of our 12 polymorphic warnings, the two most resistant to repetition suppression were animated (“jiggle” and “scale”).

Contributions of Experiment 2: Mouse Cursor-Tracking Behavioral Experiment—

Polymorphic vs. Static Warnings

In Experiment 2, we made a methodological contribution by providing three mouse cursor-tracking indicators of decreasing attention and therefore habituation: AUC, initial acceleration, and average speed. Whereas fMRI is considered costly and labor intensive and requires a high level of specialized expertise [28], our mouse cursor-tracking measures are inexpensive to implement and can be mass deployed. Moreover, fMRI necessarily introduces artificiality into an experimental task because of the requirements of the technology (e.g., the fMRI machine is loud, and participants must perform the task lying down). In contrast, the mouse cursor-tracking measures are unobtrusive to the user and may be deployed in web-based tasks simply by including a JavaScript file within standard web pages. Further, it can be deployed without any special hardware or processes (e.g., an eye tracker and

a configuration process) that may induce artificiality. Thus, our mouse cursor-tracking measures provide behavioral security researchers with a powerful technique to assess user habituation.

Second, the results of the mouse cursor-tracking behavioral experiment corroborated the fMRI findings of Experiment 1. Specifically, we found that polymorphic warnings resulted in both *lower* habituation (as evidenced by the main effect) and *reduced* habituation (as indicated by the interaction effect of the polymorphic treatment and the warning display order) than did static warnings. Using two complimentary neurophysiological measures allowed us to validate the results of both techniques, as well as to compensate for weaknesses inherent in each method.

Overall Contributions

Finally, our findings highlight the usefulness of applying NeuroIS to the domain of behavioral information security. Because automatic or unconscious mental processes underlie much of human cognition and decision making [26], they likely play an important role in a number of other security behaviors. Therefore, our research points to promising new research directions for behavioral information security in general.

Implications for Theory and Practice

Although users are frequently cited by security researchers as careless and inattentive [39], our results show that at least part of this behavior is obligatory and unconscious as a natural consequence of how the brain works. Our findings thus add to the chorus that users are not the enemy [1]. These results also illustrate that users, in addition to having to defend against malicious actors and software, must battle their own biology. Future research should investigate other obligatory, unconscious, or automatic behaviors that undermine the information security of individuals and organizations. NeuroIS methods are uniquely qualified to examine such behaviors because of their ability to observe phenomena beneath the cognition of the user [28].

Our findings have important implications for practice in the development of interventions to reduce habituation to security warnings. Rather than relying only on interventions such as SETA

programs, which encourage greater attention and vigilance [46], our results suggest that an effective complementary measure is to develop UI design artifacts that target repetition suppression in the brain, such as the polymorphic warning tested in this study. Rather than requiring explanations and training that can require hours, our polymorphic artifact elicits positive effects in milliseconds.

Importantly, in providing this benefit, the polymorphic warning artifact presented in this study is unobtrusive and imposes no additional cost to the user. In contrast, other techniques for curbing habituation, such as imposing a time delay on security warnings before they can be dismissed [10, 12], impose a cost on the user that can be considerable over time and when aggregated over a large workforce or populace [39]. Further, our polymorphic warning artifact is simple and cost-effective to implement and can be implemented in virtually any kind of system.

Limitations and Future Research

This research is subject to a number of limitations. First, both experiments used laboratory experiments that necessarily introduced artificiality into their tasks. In particular, current technology limits the flexibility and realism of experimental tasks. Fortunately, these problems in the fMRI Experiment 1 were at least partially compensated for by Experiment 2, wherein participants conducted a more ecologically valid task. Additionally, the object of laboratory experimentation is to maximize precision and control, not external validity [23, 56]. We therefore leave to future research to apply field methodologies that can achieve greater levels of external validity.

Second, both experiments in this study were cross-sectional. Although we observed that our polymorphic warning was more resistant to habituation than static warnings in our experiment, it's possible that the effectiveness of the polymorphic may decline overtime. Further research is needed to test the effectiveness of the polymorphic warning longitudinally.

Third, although frequently cited as a problem [17], habituation is not the only factor influencing the failure of users to heed security warnings. Other factors worthy of examination include the urgency of the task at hand [41], lack of comprehension of the warning message [85], or

conscientious decisions to ignore the warnings for a variety of reasons [30, 39]. We welcome further research on the multifaceted problem of security-warning disregard.

CONCLUSION

User habituation to security warnings has long been a point of concern in the area of information security [9]. In past research, habituation has been attributed to user carelessness, inattention, and ineptitude [39]. In contrast, we demonstrate in this study that habituation is largely obligatory, as a result of the way the brain processes familiar visual stimuli. A chief implication of our results is that, because habituation occurs unconsciously at the neurobiological level, interventions designed to encourage greater attention and vigilance on the part of users, such as SETA programs, are incomplete on their own. Our findings suggest that a complementary solution is to develop UI designs that are less susceptible to habituation. The polymorphic warning artifact developed in this study is one such effective design. Accordingly, our study supports the development of more complete theories of security behavior that take into account the biology of the user and provide useful neurophysiological measures to guide the design of habituation-resistant warnings.

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*, Washington, D.C.: USENIX Association (2013), pp. 257-272.
3. Akhlaghpour, S., Wu, J., Lapointe, L. and Pinsonneault, A. The ongoing quest for the IT artifact: Looking back, moving forward. *Journal of Information Technology*, 28, 2 (2013), 150-166.
4. Anderson, B., Vance, A., Kirwan, B., Eargle, D. and Howard, S. Users aren't (necessarily) lazy: Using NeuroIS to explain habituation to security warnings. In *International Conference on Information Systems*, Auckland, New Zealand: AIS, 2014.
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. Bartra, O., McGuire, J.T. and Kable, J.W. The valuation system: A coordinate-based meta-analysis of bold fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76 (2013), 412-427.
7. Braun, C.C., Greeno, B. and Silver, N.C. Differences in behavioral compliance as a function of warning color. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*: Sage Publications, 1994, pp. 379-383.

8. 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.
9. Bravo-Lillo, C., Cranor, L.F., Downs, J., Komanduri, S. and Sleeper, M. Improving computer security dialogs. In *Proceedings of the 13th IFIP TC 13 International Conference on Human-Computer Interaction - Volume 6949 Part IV*, Lisbon, Portugal: Springer-Verlag (2011), pp. 18-35.
10. 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* Newcastle, United Kingdom: ACM (2013), pp. 1-12.
11. 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 *10th Symposium On Usable Privacy and Security (SOUPS 2014)*: USENIX Association, 2014, pp. 105-111.
12. Brustoloni, J.C. and Villamarín-Salomón, R. Improving security decisions with polymorphic and audited dialogs. In *SOUPS '07: Proceedings of the Third symposium on Usable Privacy and Security*, New York, NY, USA: ACM, 2007, pp. 76-85.
13. Buckner, R.L., Andrews-Hanna, J.R. and Schacter, D.L. The brain's default network. *Annals of the New York Academy of Sciences*, 1124, 1 (2008), 1-38.
14. Cleveland, W.S. Lowess: A program for smoothing scatterplots by robust locally weighted regression. *The American Statistician*, 35, 1 (1981), 54.
15. Corbetta, M., Kincade, J.M., Ollinger, J.M., McAvoy, M.P. and Shulman, G.L. Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nature Neuroscience*, 3, 3 (2000), 292-297.
16. Corbetta, M. and Shulman, G.L. Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3, 3 (2002), 201-215.
17. Cranor, L.F. A framework for reasoning about the human in the loop. In *UPSEC'08 Proceedings of the 1st Conference on Usability, Psychology, and Security*, San Francisco, California: USENIX Association (2008), pp. 1-15.
18. Crossler, R.E., Johnston, A.C., Lowry, P.B., Hu, Q., Warkentin, M. and Baskerville, R. Future directions for behavioral information security research. *Computers & Security*, 32, 1 (2013), 90-101.
19. Dale, A.M. Optimal experimental design for event-related fMRI. *Human Brain Mapping*, 8, 2-3 (1999), 109-114.
20. Dale, R., Kehoe, C. and Spivey, M.J. Graded motor responses in the time course of categorizing atypical exemplars. *Memory & Cognition*, 35, 1 (2007), 15-28.
21. de Guinea, A.O., Titah, R. and Léger, P.-M. Explicit and implicit antecedents of users' behavioral beliefs in information systems: A neuropsychological investigation. *Journal of Management Information Systems*, 30, 4 (2014), 179-210.
22. De Keukelaere, F., Yoshihama, S., Trent, S., Zhang, Y., Luo, L. and Zurko, M.E. Adaptive security dialogs for improved security behavior of users. In *Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part I*, Uppsala, Sweden: Springer-Verlag (2009), pp. 510-523.
23. Dennis, A.R. and Valacich, J.S. Conducting experimental research in information systems. *Communications of the Association for Information Systems*, 7, 5 (2001).
24. Desimone, R. Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences*, 93, 24 (1996), 13494-13499.
25. Dhamija, R., Tygar, J.D. and Hearst, M. Why phishing works. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Montréal, Québec, Canada: ACM (2006), pp. 581-590.
26. 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.

27. Dimoka, A. How to conduct a functional magnetic resonance (fMRI) study in social science research. *MIS Quarterly*, 36, 3 (2012), 811-840.
28. 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.
29. 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*, Florence, Italy: ACM, 2008, pp. 1065-1074.
30. Egelman, S. and Schechter, S. The importance of being earnest [in security warnings]. In: A.-R. Sadeghi (ed.), *Financial cryptography and data security*: Springer Berlin Heidelberg, 2013, pp. 52-59.
31. Freeman, J.B., Dale, R. and Farmer, T.A. Hand in motion reveals mind in motion. *Frontiers in Psychology*, 2, 59 (2011), 1-6.
32. Furnell, S. Usability versus complexity—striking the balance in end-user security. *Network Security*, 2010, 12 (2010), 13-17.
33. Google. Permission warnings—Google Chrome. https://developer.chrome.com/extensions/permission_warnings (accessed on 1/26/2015)
34. Greicius, M.D., Krasnow, B., Reiss, A.L. and Menon, V. Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*, 100, 1 (2003), 253-258.
35. Grill-Spector, K. and Malach, R. The human visual cortex. *Annual Review of Neuroscience*, 27 (2004), 649-677.
36. 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.
37. Groves, P.M. and Thompson, R.F. Habituation: A dual-process theory. *Psychological Review*, 77 (1970), 419-450.
38. 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*, Toronto, Ontario, Canada: ACM (2014), pp. 2647-2656.
39. Herley, C. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *Proceedings of the 2009 Workshop on New Security Paradigms Workshop*: ACM, 2009, pp. 133-144.
40. Hibbeln, M., Jenkins, J.L., Schneider, C., Valacich, J.S. and Weinmann, M. How is your user feeling? Inferring emotion through human-computer interaction devices. *MIS Quarterly* (forthcoming).
41. Jenkins, J.L. and Durcikova, A. What, I shouldn't have done that? The influence of training and just-in-time reminders on secure behavior. In *International Conference for Information Systems (ICIS)*, Milan, Italy: AIS (2013).
42. Kalsher, M. and Williams, K. Behavioral compliance: Theory, methodology, and result. In: M. Wogalter (ed.), *Handbook of Warnings*. Mahwah NJ: Lawrence Erlbaum Associates, 2006, pp. 313-331.
43. Kalsher, M.J., Brewster, B.M., Wogalter, M.S. and Spunar, M.E. Hazard level perceptions of current and proposed warning sign and label panels. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 39, 5 (1995), 351-355.
44. 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.
45. Kandel, E.R. The molecular biology of memory storage: A dialogue between genes and synapses. *Science*, 294, 5544 (2001), 1030-1038.

46. Karjalainen, M. and Siponen, M. Toward a new meta-theory for designing information systems (IS) security training approaches. *Journal of the Association for Information Systems*, 12, 8 (2011), 518-555.
47. Kuan, K.K.Y., Zhong, Y. and Chau, P.Y.K. Informational and normative social influence in group-buying: Evidence from self-reported and EEG data. *Journal of Management Information Systems*, 30, 4 (2014), 151-178.
48. Lacy, J.W., Yassa, M.A., Stark, S.M., Muftuler, L.T. and Stark, C.E. Distinct pattern separation related transfer functions in human CA3/dentate and ca1 revealed using high-resolution fMRI and variable mnemonic similarity. *Learning & Memory*, 18, 1 (2011), 15-18.
49. Laughery, K.R., Young, S.L., Vaubel, K.P. and Brelsford Jr, J.W. The noticeability of warnings on alcoholic beverage containers. *Journal of Public Policy & Marketing*, 12, 1 (1993), 38-56.
50. Leung, C.-M. Depress phishing by CAPTCHA with OTP. In *3rd International Conference on Anti-counterfeiting, Security, and Identification in Communication, 2009. ASID 2009.*, Hong Kong: IEEE, 2009, pp. 187-192.
51. Li, M., Jiang, Q., Tan, C.-H. and Wei, K.-K. Enhancing user-game engagement through software gaming elements. *Journal of Management Information Systems*, 30, 4 (2014), 115-150.
52. Logothetis, N.K., Pauls, J., Augath, M., Trinath, T. and Oeltermann, A. Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412, 6843 (2001), 150-157.
53. Lueschow, A., Miller, E.K. and Desimone, R. Inferior temporal mechanisms for invariant object recognition. *Cerebral Cortex*, 4, 5 (1994), 523-531.
54. MacKay, D.M. Cerebral organization and the conscious control of action. In: J. Eccles (ed.), *Brain and conscious experience*. Berlin Heidelberg: Springer, 1965, pp. 422-445.
55. Mason, M.F., Norton, M.I., Van Horn, J.D., Wegner, D.M., Grafton, S.T. and Macrae, C.N. Wandering minds: The default network and stimulus-independent thought. *Science*, 315, 5810 (2007), 393-395.
56. McGrath, J.E. Dilemmatics: The study of research choices and dilemmas. *American Behavioral Scientist*, 25, 2 (1981), 179-210.
57. Meservy, T.O., Jensen, M.L. and Fadel, K.J. Evaluation of competing candidate solutions in electronic networks of practice. *Information Systems Research*, 25, 1 (2014), 15-34.
58. Mohr, P.N., Biele, G. and Heekeren, H.R. Neural processing of risk. *The Journal of Neuroscience*, 30, 19 (2010), 6613-6619.
59. Motley, S.E. and Kirwan, C.B. A parametric investigation of pattern separation processes in the medial temporal lobe. *The Journal of Neuroscience*, 32, 38 (2012), 13076-13084.
60. Navalpakkam, V. and Churchill, E. Mouse tracking: Measuring and predicting users' experience of web-based content. In *2012 ACM Annual Conference on Human Factors in Computing Systems*, Austin, Texas, USA: ACM (2012), pp. 2963-2972.
61. 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)*, San Diego, CA: (2014), pp. 1-16.
62. Orlikowski, W.J. and Iacono, C.S. Research commentary: Desperately seeking the "IT" in IT research—a call to theorizing the IT artifact. *Information Systems Research*, 12, 2 (2001), 121-134.
63. 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.
64. Riedl, R., Mohr, P.N.C., Kenning, P.H., Davis, F.D. and Heekeren, H.R. Trusting humans and avatars: A brain imaging study based on evolution theory. *Journal of Management Information Systems*, 30, 4 (2014), 83-114.
65. Rodden, K., Fu, X., Aula, A. and Spiro, I. Eye-mouse coordination patterns on web search results pages. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, Florence, Italy: ACM (2008), pp. 2997-3002.

66. 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.
67. Sanders, M.S. and McCormick, E.J. Human factors in engineering and design. New York: McGraw-Hill, 1987.
68. Schechter, S.E., Dhamija, R., Ozment, A. and Fischer, I. The emperor's new security indicators. In *Security and Privacy, 2007. SP'07. IEEE Symposium on*, Berkeley, CA: IEEE, 2007, pp. 51-65.
69. Segaert, K., Weber, K., de Lange, F.P., Petersson, K.M. and Hagoort, P. The suppression of repetition enhancement: A review of fMRI studies. *Neuropsychologia*, 51, 1 (2013), 59-66.
70. Seidler, R., Noll, D. and Thiers, G. Feedforward and feedback processes in motor control. *NeuroImage*, 22, 4 (2004), 1775-1783.
71. Silver, N.C. and Wogalter, M.S. Broadening the range of signal words. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Denver, Colorado: Sage, 1989, pp. 555-559.
72. 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.
73. Sokolov, E. Higher nervous functions: The orienting reflex. *Annual Review of Physiology*, 25 (1963), 545-580.
74. Sotirakopoulos, A., Hawkey, K. and Beznosov, K. On the challenges in usable security lab studies: Lessons learned from replicating a study on SSL warnings. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, Menlo Park, CA: ACM, 2011, pp. 3:1-3:18.
75. Stark, C.E.L. and Squire, L.R. When zero is not zero: The problem of ambiguous baseline conditions in fMRI. *Proceedings of the National Academy of Sciences*, 98, 22 (2001), 12760-12766.
76. Straub, D., Boudreau, M.-C. and Gefen, D. Validation guidelines for IS positivist research. *Communications of the Association for Information Systems*, 13, 24 (2004), 380-427.
77. Strawbridge, J.A. The influence of position, highlighting, and imbedding on warning effectiveness. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Dayton, OH: Sage, 1986, pp. 716-720.
78. Sunaert, S., Van Hecke, P., Marchal, G. and Orban, G.A. Motion-responsive regions of the human brain. *Experimental Brain Research*, 127, 4 (1999), 355-370.
79. Thompson, R.F. and Spencer, W.A. Habituation: A model phenomenon for the study of neuronal substrates of behavior. *Psychological Review*, 73, 1 (1966), 16-43.
80. Thompson, R.F. Habituation: A history. *Neurobiology of Learning and Memory*, 92, 2 (2009), 127-134.
81. Vigilante, W.J. and Wogalter, M. Effects of label format on knowledge acquisition and perceived readability by younger and older adults. *Ergonomics*, 46, 4 (2003), 327-344.
82. vom Brocke, J. and Liang, T.-P. Guidelines for neuroscience studies in information systems research. *Journal of Management Information Systems*, 30, 4 (2014), 211-234.
83. Welsh, T.N. and Elliott, D. Movement trajectories in the presence of a distracting stimulus: Evidence for a response activation model of selective reaching. *The Quarterly Journal of Experimental Psychology Section A*, 57, 6 (2004), 1031-1057.
84. Wogalter, M. and Vigilante Jr, W.J. Attention switch and maintenance. In: M.S. Wogalter (ed.), *Handbook of Warnings*. Mahwah NJ: Lawrence Erlbaum Associates, 2006, pp. 245-266.
85. Wogalter, M.S. Communication-human information processing (C-HIP) model. In: M.S. Wogalter (ed.), *Handbook of Warnings*. Mahwah, NJ: Lawrence Erlbaum Associates, 2006, pp. 51-61.
86. Wogalter, M.S. Purposes and scope of warnings. In: M.S. Wogalter (ed.), *Handbook of Warnings*. Mahwah, NJ: Lawrence Erlbaum Associates, 2006, pp. 3-9.
87. Yassa, M.A., Stark, S.M., Bakker, A., Albert, M.S., Gallagher, M. and Stark, C.E. High-resolution structural and functional MRI of hippocampal CA3 and dentate gyrus in patients with amnesic mild cognitive impairment. *NeuroImage*, 51, 3 (2010), 1242-1252.

88. Young, S.L. and Wogalter, M. Effects of conspicuous print and pictorial icons on comprehension and memory of instruction manual warnings. *Human Factors*, 32 (1990), 637-649.
89. Young, S.L. Increasing the noticeability of warnings: Effects of pictorial, color, signal icon and border. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, San Francisco, VA: Sage, 1991, pp. 580-584.

AUTHOR BIOGRAPHIES

Bonnie Brinton Anderson is an Associate Professor of Information Systems and Director of the Master of Information Systems Management (MISM) program in the Marriott School of Management at Brigham Young University. She received her PhD from Carnegie Mellon University. Her work has been published or accepted in *Information Systems Research*, *Journal of Management Information Systems*, *Journal of the Association for Information Systems*, *European Journal of Information Systems*, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, *Decision Support Systems*, *Electronic Commerce Research*, *Expert Systems with Applications*; *Electronic Commerce Research*, *Communications of the ACM*, *Information Sciences*, *IEEE Transactions: Systems, Men, and Cybernetics*, *The Journal of Systems and Software*, and other outlets. She currently researches the intersection of decision neuroscience and behavioral information security.

Anthony Vance is as an Associate Professor of Information Systems in the Marriott School of Management of Brigham Young University. He has earned Ph.D. degrees in Information Systems from Georgia State University, USA; the University of Paris—Dauphine, France; and the University of Oulu, Finland. His previous experience includes working as a visiting research professor in the Information Systems Security Research Center at the University of Oulu. His work is published in outlets such as *MIS Quarterly*, *Information Systems Research*, *Journal of Management Information Systems*, *Journal of the Association for Information Systems*, *European Journal of Information Systems*, *Journal of the American Society for Information Science and Technology*, and *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. His research focuses on behavioral and neuroscience applications to information security. He currently is an associate editor at *MIS*

Quarterly and serves on the editorial board of *Journal of the Association for Information Systems*.

C. Brock Kirwan is an Associate Professor of Psychology and Neuroscience at Brigham Young University. He received his PhD in Psychological and Brain Sciences from Johns Hopkins University in 2006. Dr. Kirwan has a decade of experience conducting fMRI scans with patient populations at Johns Hopkins University, the University of California, San Diego, the University of Utah, and now BYU. He has published numerous papers reporting fMRI and neuropsychological results in journals such as *Science*, *Proceedings of the National Academy of Sciences*, *Neuron*, and the *Journal of Neuroscience*, as well as journals in information systems such as *Information Systems Research*, *Journal of Management Information Systems*, *Journal of the Association for Information Systems*, and *European Journal of Information Systems*.

Jeff Jenkins is an Assistant Professor of Information Systems at the Marriott School of Management, Brigham Young University. He graduated with a Ph.D. in Management Information Systems from the University of Arizona. His active research includes human-computer interaction and behavioral information security. In a human-computer interaction context, Jeff's research explores how to infer human states using computer input devices such as the computer mouse, keyboard, or touchscreen. His research has been published in various journals and conference proceedings, including *MIS Quarterly*, *Information Systems Research*, *Journal of Management Information Systems*, *European Journal of Information Systems*, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, *Computers in Human Behavior*, and others. Prior to earning his Ph.D., Jeff was a Software Engineer in both the public and private sectors.

David Eargle is a doctoral candidate in the Information Systems and Technology Management Area at the University of Pittsburgh in the Katz Graduate School of Business. He is currently a NSF Graduate Research Fellow. He completed a joint baccalaureate-master's program in information systems management at Brigham Young University, completing the IS PhD preparation program and graduating magna cum laude with University Honors. His research interests include human-computer interaction and information security. He has coauthored several articles in these areas using neurophysiological and other methodologies in outlets such as *Information Systems Research*, *Journal of Management Information Systems*, *Journal of the Association for Information Systems*, *European Journal of Information Systems*, the *International Conference on Information Systems*, and the *Hawaii International Conference on System Sciences*, along with the *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.

ONLINE APPENDICES

APPENDIX A: FMRI TECHNICAL DETAILS FOR EXPERIMENT 1

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 were acquired with a T1-weighted magnetization-prepared rapid acquisition including a gradient-echo (MP-RAGE) sequence with the following parameters: TE = 2.26 ms, flip angle = 9°, slices = 176, slice thickness = 1.0 mm, matrix size = 256 × 215, and voxel size = 1 mm × 0.98 mm × 0.98 mm. Functional scans were acquired with a gradient-echo, echo-planar, T2*-weighted pulse sequence with the following parameters: TR = 2000 ms, TE = 28 ms, flip angle = 90°, slices = 40, slice thickness = 4.0 mm (no skip), matrix size = 64 × 64, and voxel size = 3.44 mm × 3.44 mm × 3 mm.

Protocol

Participants were given a verbal briefing about the MRI procedures and the task, and were then situated supine in the scanner. Visual stimuli were viewed using a mirror attached to the head coil reflecting a large monitor outside the scanner that was configured to display images in reverse so that they appeared normal when viewed through the mirror. We first performed a 10-second localizer scan, followed by a 7-minute structural scan. Following these, we started the experimental task (see Figure A1). We used E-Prime software to display the stimuli and synchronize the display events and scanner software. Total time in the scanner was 55 minutes.

Note: For each participant, perform steps 1–4.

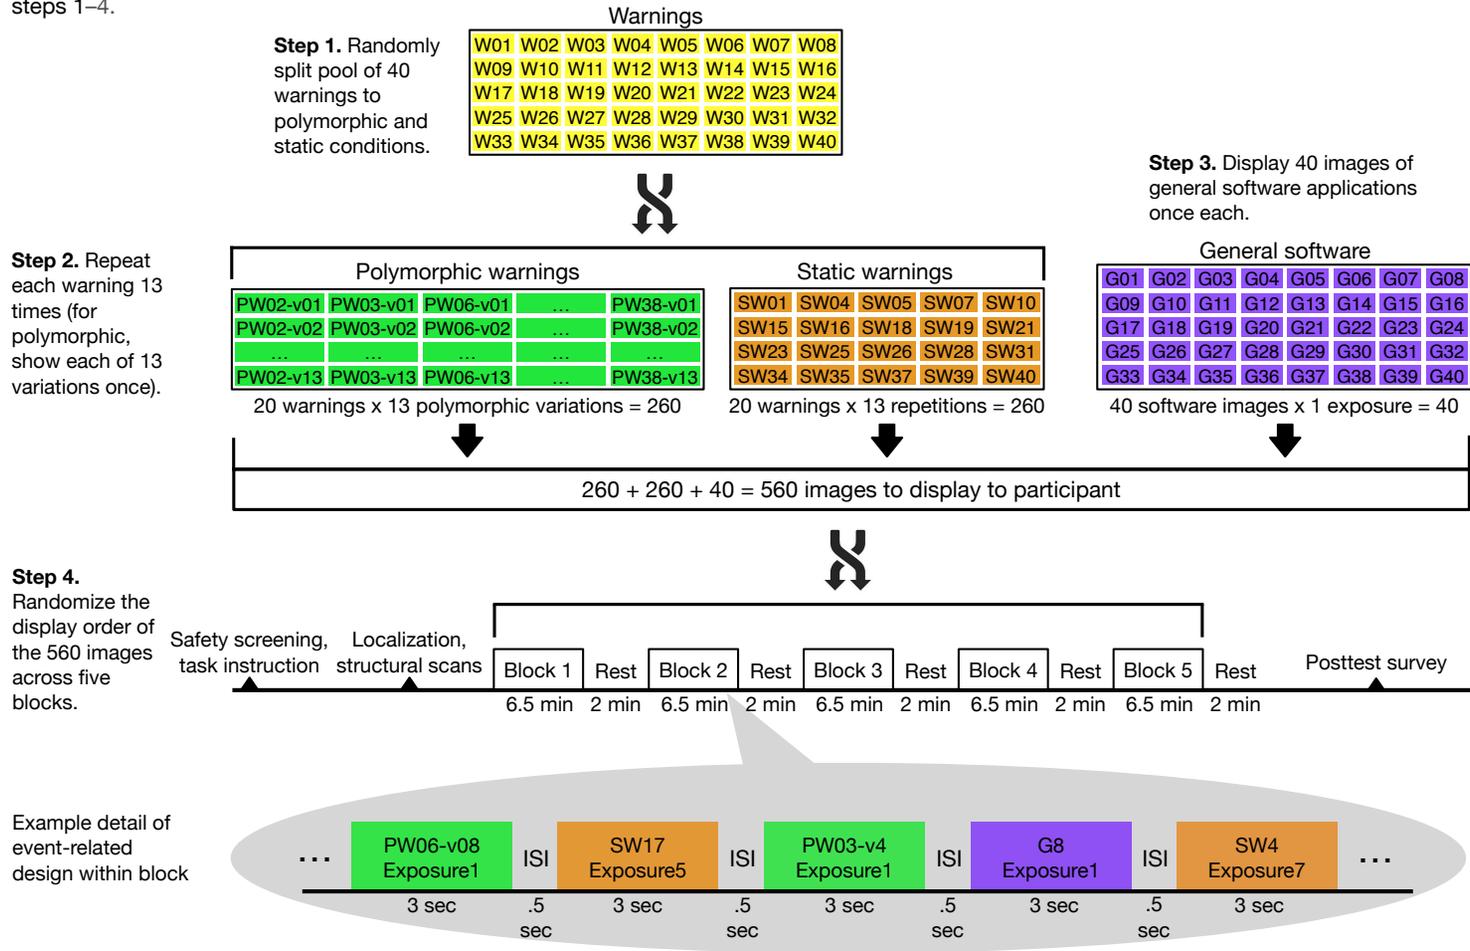


Figure A1. Graphical representation of fMRI experimental procedure for Experiment 1.
ISI = Interstimulus interval

Analysis

MRI data were analyzed with the Analysis of Functional Images (AFNI) suite of programs [11]. Briefly, 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. Data from each run were 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. As in previous studies [e.g., 44], spatial normalization was accomplished by first warping the structural scan to the Talairach atlas [63] followed by warping to a template brain with Advanced Neuroimaging Tools (ANTs). For the single-subject ("first level") analyses, we performed multiple regression of the form $y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 \dots \beta_n x_n + \epsilon$, where "y" is the observed fMRI timecourse for each voxel and each "x" term is a vector regressor representing either conditions of interest (e.g., stimulus type or repetition number) or a nuisance regressor (e.g., motion or scanner drift). We created separate behavioral regressors coding for each repetition number for both static warnings and polymorphic warnings (total of 26 task regressors) in addition to nuisance variables coding for motion (3 rotations and 3 directions of translation) and scanner drift (4 polynomial regressors for each of the 5 scan runs). The single-presentation general computing screenshots served as the implicit baseline for the regression analysis (see "Experimental Design"). Stimulus events were modeled using a stick function convolved with the canonical hemodynamic response. Resulting parameter estimates (beta values) were blurred with a 5-mm FWHM Gaussian kernel. Parameter estimates for the conditions of interest were then entered into group-level analyses, such as ANOVAs or *t*-tests (please see "Experiment 1: Data Analysis" for detailed descriptions of group-level analyses), which were used to determine functional regions of interest (ROIs). Once functional ROIs were identified, we extracted mean parameter estimates within each ROI for further investigation in order to characterize the direction and strength of interactions and other effects. All whole-brain

voxel-wise tests were corrected for multiple comparisons using a false-discovery rate of 0.05 and a spatial extent threshold of 20 contiguous voxels (540 mm³).

Two regression analyses were then conducted. In the first, polymorphic warning images were grouped according to repetition number regardless of the specific polymorphic variation. This was done to allow us to perform an analysis of repetition effects. The second analysis grouped polymorphic warning images according to specific variation, regardless of presentation order.

APPENDIX B: SUPPLEMENTARY ANALYSES FOR EXPERIMENT 1

Table B1. Analysis results for animated vs. non-animated comparison						
	#Voxels	X	Y	Z	t(21)	p
Contrast: Animated > Non-Animated						
R. Cuneus	280	-2	86	24	4.75	<.001
R. Temporoparietal Junction (hMT+)	262	-56	44	30	5.54	<.001
L. Retrosplenial Cortex	208	5	56	21	4.74	<.001
R. Anterior Cingulate	20	-2	-32	-7	4.49	<.001
R. Posterior Middle Temporal Gyrus	20	-47	65	9	3.52	.002
Contrast: Non-Animated > Animated						
L. Posterior Occipital Lobe	123	14	98	-1	4.50	<.001
L. Lingual Gyrus	113	38	71	-16	4.28	<.001
R. Posterior Occipital Lobe	111	-14	95	-4	4.99	<.001
R. Inferior Parietal Lobule	42	-32	47	45	3.31	.003
L. Inferior Parietal Lobule	27	38	53	48	3.76	<.001

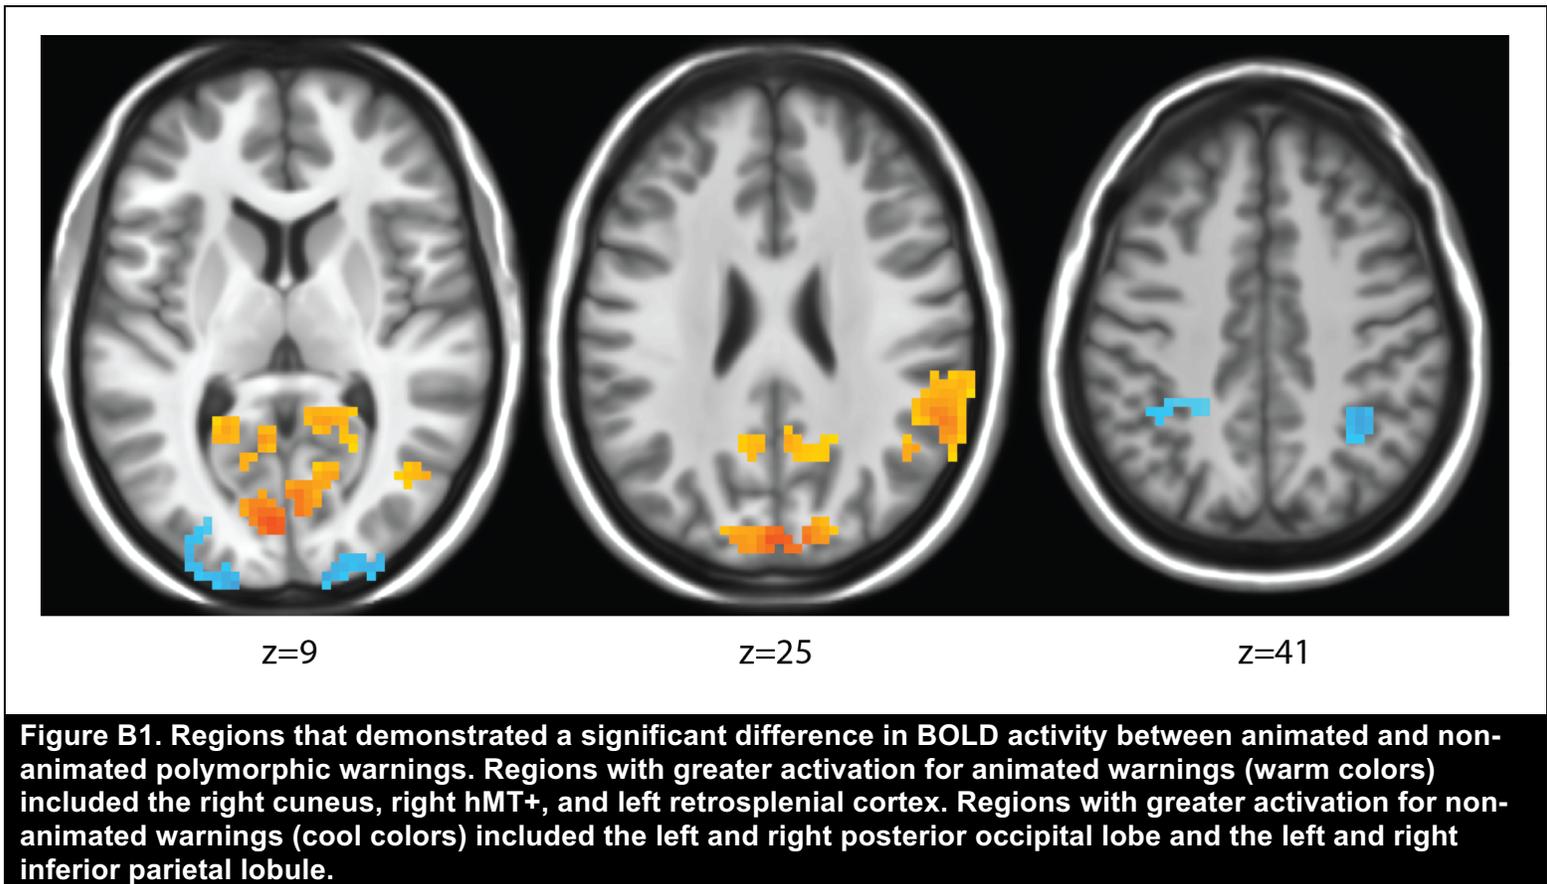


Figure B1. Regions that demonstrated a significant difference in BOLD activity between animated and non-animated polymorphic warnings. Regions with greater activation for animated warnings (warm colors) included the right cuneus, right hMT+, and left retrosplenial cortex. Regions with greater activation for non-animated warnings (cool colors) included the left and right posterior occipital lobe and the left and right inferior parietal lobule.

APPENDIX C: A BRIEF HISTORY AND DESCRIPTION OF MOUSE CURSOR TRACKING

Origins in Hand and Finger Tracking Research

As a measure of neuromotor and psychological outcomes, mouse cursor tracking fits within the larger domain of hand and finger tracking research. The tracking of the hand and fingers (i.e., fine motor movements) has long been used to provide insight into human cognitive processes. A 1961 review of fine motor movement tracking as an instrument for psychological experiments described the field as consisting of “several hundred task-oriented tracking” publications [1, p. 56]. In this early research, custom input devices, such as joy sticks [64] and steering-wheel-like devices [28], were used to measure motor movement characteristics, including precision, speed, manual dexterity, and reaction time, to name a few [see 51 for a summary]. Researchers used these motor movement characteristics to gain insight into various cognitive and neuromotor processes, such as response orientation [19, 20], cognitive integration [29], and change anticipation [53].

Mouse Tracking and the Advent of Personal Computing

As the adoption of personal computers (equipped with hand-held input devices, such as the computer mouse) drastically increased in the 1990s and 2000s, so did the opportunities for studying cognitive processes via hand movements. The term “mouse cursor tracking” (sometimes just referred to as “mouse tracking”) was coined, referring to the measurement of cursor positions and timestamps of movements on the computer screen (which could be manipulated by the computer mouse or another input device, such as a track pad, pointing stick, or touch screen). Researchers initially explored using mouse cursor tracking as a cost-effective alternative to eye tracking to denote where people devote their attention in a HCI context [7, 8, 30]. For example, research has shown that eye-gaze and cursor movement patterns are highly correlated [8, 30, 49]. When scanning search results, the cursor often follows the eye and marks promising search hits (i.e., the cursor pointer stops or lingers near such information), suggesting where users devote their attention [54]. Likewise, users often move the

Table C1. Examples of recent cognitive and emotional processes examined through monitoring mouse cursor/hand movements.

cursor while viewing web pages, suggesting that the cursor may indicate where users focus their attention [45]. In selecting menu items, the cursor often tags potential targets (i.e., hovers over the link) before selecting an item [10]. Monitoring user clicks can also assess the relevance of search results [34]. Finally, by continuously recording cursor position, researchers can assess the user's awareness, attraction, and avoidance of content (e.g., avoiding ads, not looking at text because of frustration, or struggling reading the text) [46]. Consequently, mouse cursor tracking is often applied as a usability assessment tool for visualizing cursor movements on webpages [2, 36] and developing heat maps that indicate where users devote their attention [3, 37].

Advancement of Mouse Tracking as a Neurophysiological Method

As the ability to assess more fine-grained measurements and mouse cursor movements improved, research expanded the use of mouse cursor tracking to explore a more diverse set of neuromotor and psychological responses. In a concise review of mouse cursor-tracking literature, Freeman *et al.* [25, p. 1] suggested that the movements of the hand “offer continuous streams of output that can reveal ongoing dynamics of processing, potentially capturing the mind in motion with fine-grained temporal sensitivity.” Accordingly, hundreds of recent studies have chosen mouse cursor tracking as a methodology for studying various cognitive and emotional processes (see Table E1 for recent examples). Many of these studies have focused on how people respond immediately after seeing a stimulus and the cognitive process of decision-making.

Cognitive process examined through mouse cursor/hand movements	Citation
Attitude formation, concealment of racial prejudices	[71]
Attraction toward distracting stimuli	[59, 60]
Decision conflict	[40, 48]
Decision making	[14, 39]
Deception	[15, 69]
Detection of dual cognitive processing	[26]
Dynamic competition in classification tasks	[12, 21, 22, 24]
Emotional reactions	[33, 38, 55, 73, 74]
Increased cognitive processing	[24]
Language learning, processing, or interpretation	[4, 5, 17, 61]
Learning	[13, 75]
Mathematical processing	[18]
Memory recall	[50]
Metacognition	[41]
Perception formation of people	[9, 27]
Semantic priming	[57]
Search/Recognition	[58]
Spatial knowledge development	[67]
Subconscious/Implicit/Anticipatory processing	[6, 66, 72]
Task switching	[68]

How Mouse Cursor Tracking is Commonly Implemented

Mouse cursor tracking is typically performed by embedding JavaScript into a webpage (e.g., JQuery) or by using a desktop application, such as MouseTracker [23]. For example, JQuery (a common and freely available JavaScript library) can capture the x, y coordinate and timestamp for mouse cursor movements on the computer screen. Various statistics can be calculated on the characteristics of trajectories and movements to learn about cognitive and neuromotor processes from this voluminous raw data. For example, characteristics of the trajectory include the x- and y-locations of the cursor during different points of the interaction, the number of direction changes along the trajectory, or the deviation from an idealized response trajectory (a straight line connecting the starting and ending points of a movement). Two measures of deviation from the idealized response trajectory include area-under-the-curve (the geometric area between the idealized response trajectory and the actual trajectory; AUC) and maximum deviation (the longest perpendicular line between the idealized response trajectory and the actual trajectory). Examples of movement characteristics include the

speed, the acceleration at different points, and the angle of movement, to name a few. A more exhaustive discussion of mouse cursor-tracking measures and their calculations was presented by [23, 32].

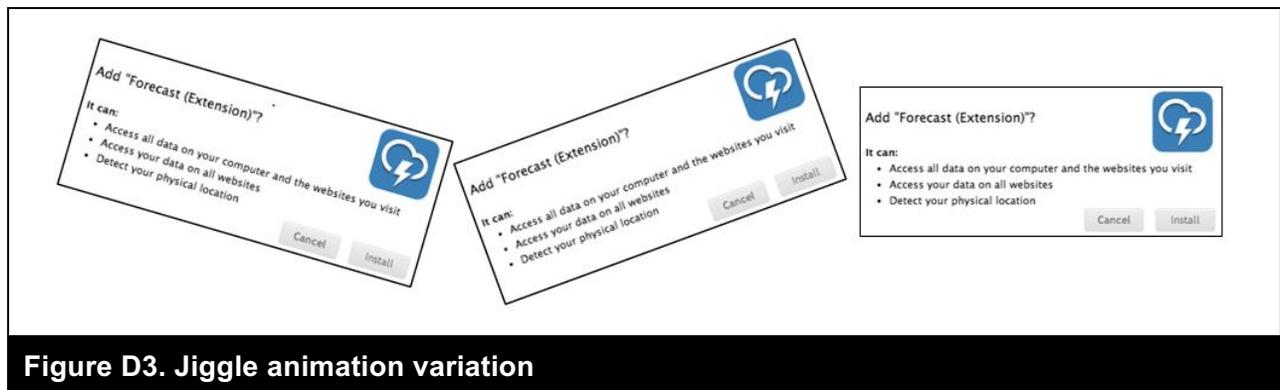
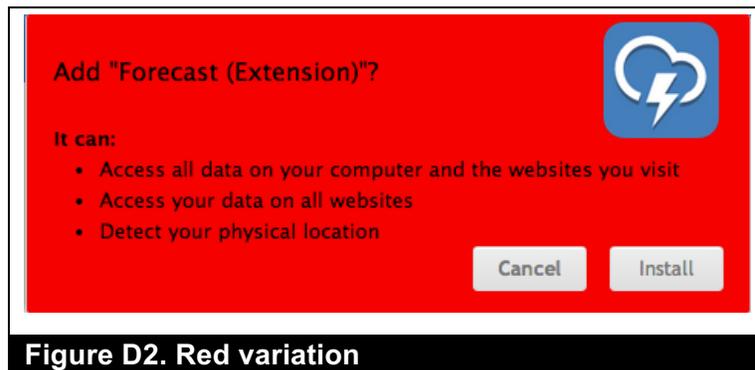
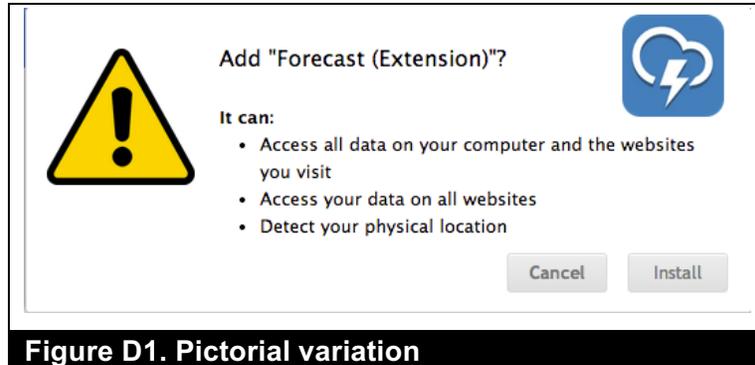
Pros and Cons of Mouse Cursor Tracking

Mouse cursor tracking has pros and cons as a research instrument. The method allows researchers to model many aspects of attention, but it cannot completely replace gaze captured through an eye tracker. For example, a user's eye-gaze fixation may change (move to another stimulus that catches their attention) without moving the mouse. In such circumstances, although a prolonged eye-gaze fixation may indicate attention or interest, a prolonged cursor fixation may not [35].

On the other hand, mouse cursor tracking can be performed at almost no cost using free JavaScript libraries that can be embedded in normal web pages. Furthermore, mouse cursor tracking can be performed in a natural environment, such as the user's personal computer as he or she interacts with websites, thereby improving the ecological validity of the research. Further, as previously discussed, analyzing mouse cursor movements may provide insights into cognitive process aside from attention (e.g., decision conflict, emotion, memory recall). Hence, mouse cursor tracking has been described as measuring "high-fidelity, real-time motor traces of the mind [that] can reveal 'hidden' cognitive states that are otherwise not availed by traditional measures" [25, p. 2].

APPENDIX D: SUPPLEMENTARY FIGURES FOR EXPERIMENT 2

Chrome Permission Warning Manipulations Used



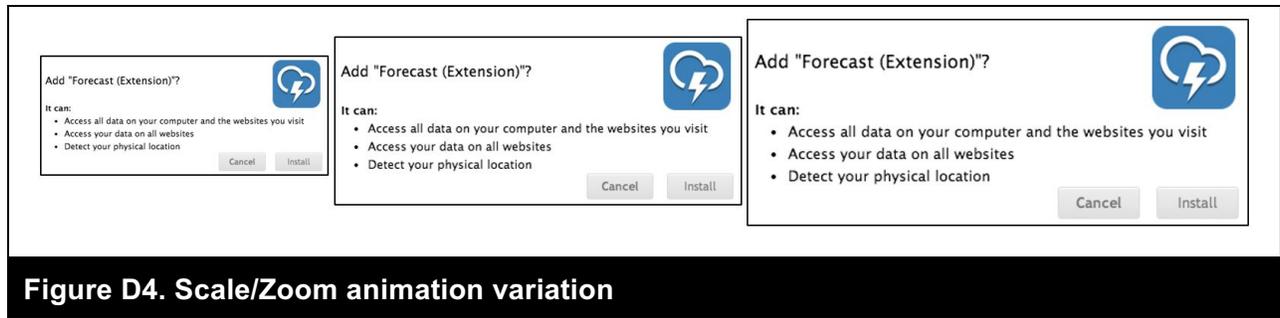


Figure D4. Scale/Zoom animation variation

Chrome Screenshots

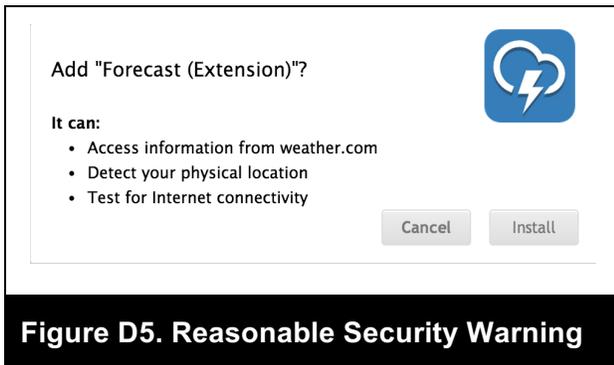


Figure D5. Reasonable Security Warning

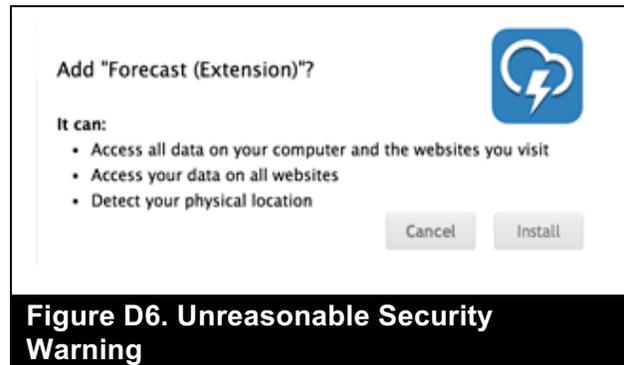


Figure D6. Unreasonable Security Warning

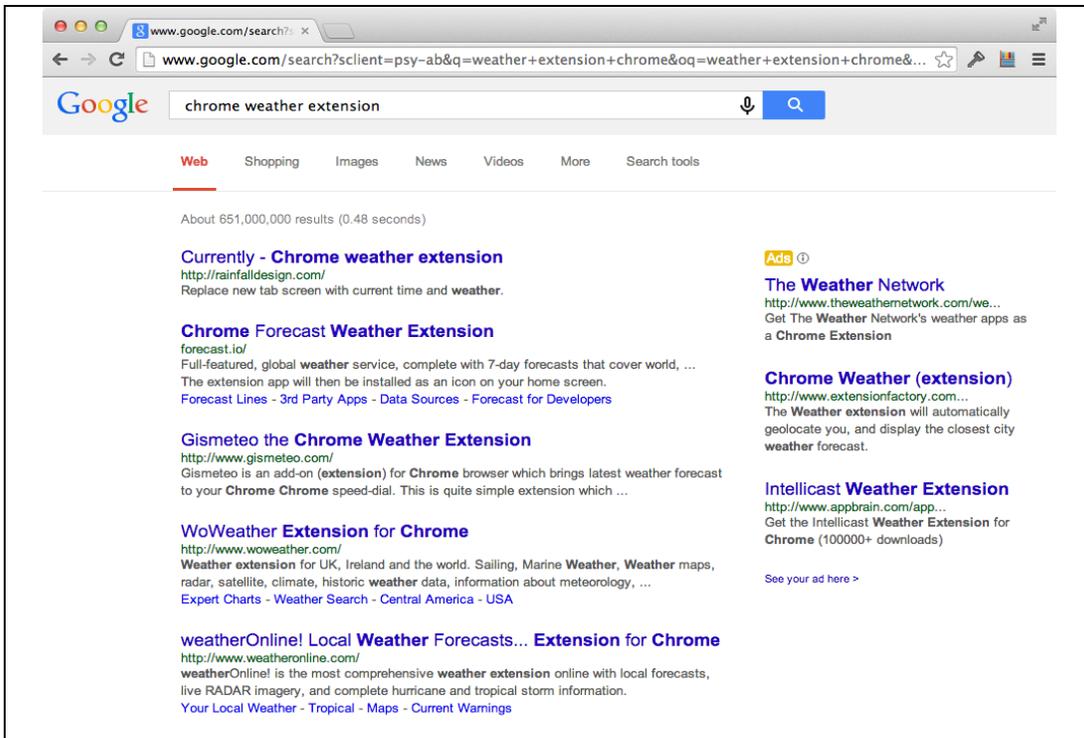


Figure D7. Sample manipulated Google search results

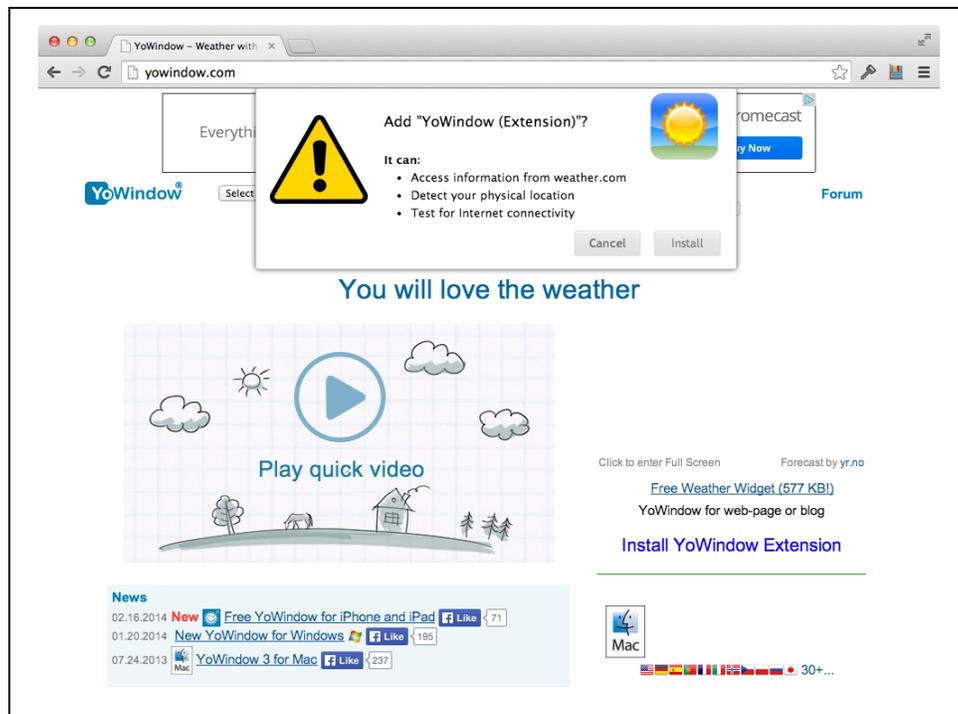


Figure D8. Sample warning on from search results

AUC Curve

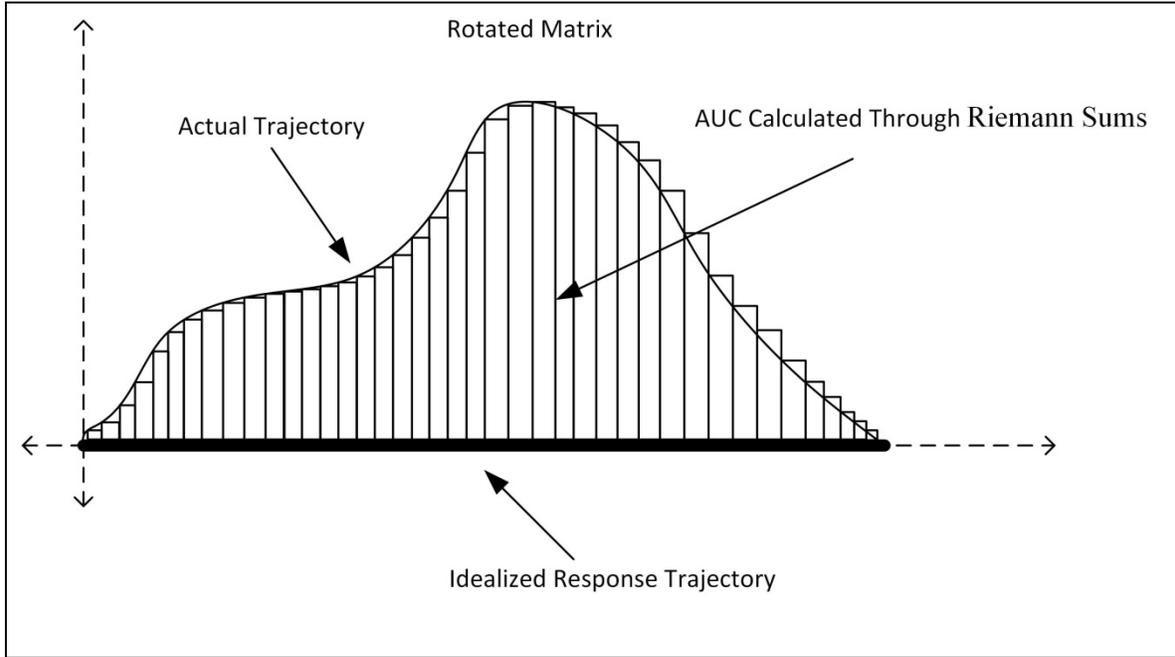


Figure D9. Calculation of AUC

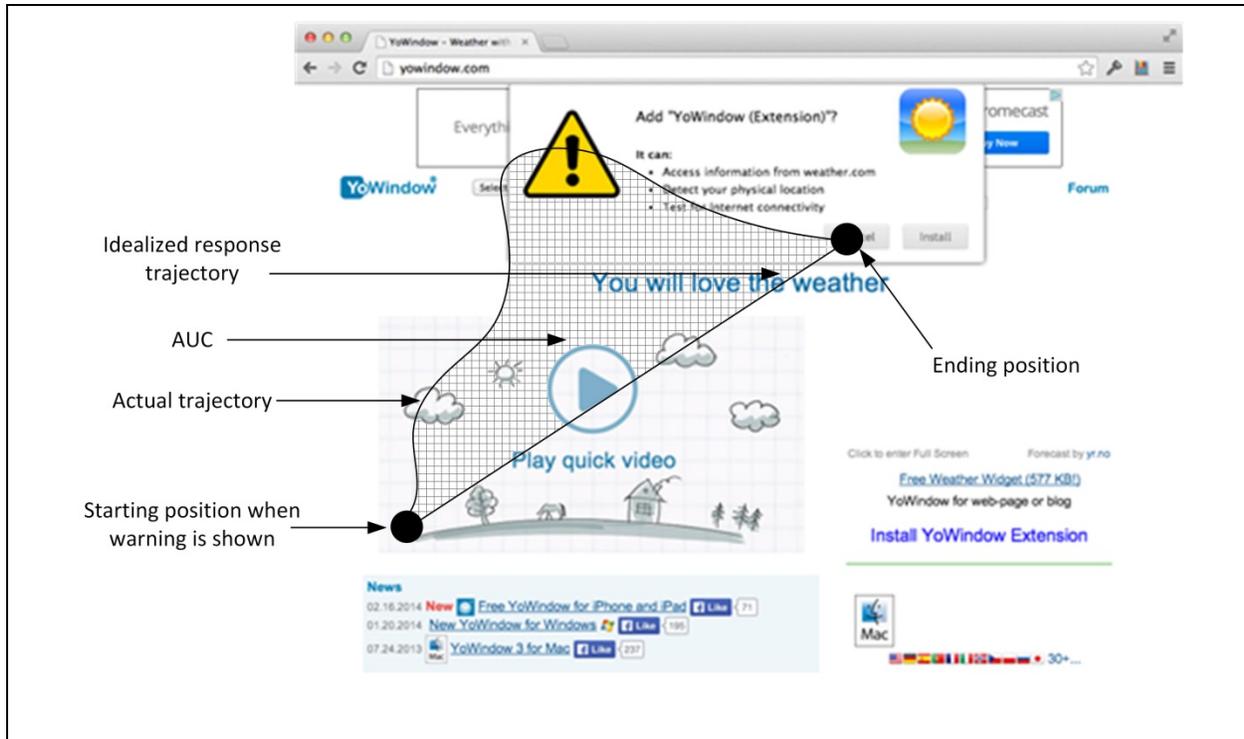


Figure D10. Example area-under-the-curve (AUC) for a security warning

APPENDIX E: MOUSE CURSOR MEASUREMENTS FOR EXPERIMENT 2

This appendix describes our calculations of area-under-the-curve (AUC), initial acceleration, and average speed. To facilitate these calculations, we rescaled all x -, y -coordinate pairs to a 4×3 rectangle (-2 to 2 and -1.5 to 1.5), which retains the aspect ratio of most computer screens. The starting position of the cursor when the warning was first shown was mapped at point 0,0, and the other coordinate pairs were remapped accordingly. The graph was then rotated so that the endpoint appeared on the x -axis in addition to the starting point at 0,0 (facilitating the calculation of the idealized response trajectory for AUC described below) using the following transformation for all coordinate pairs¹:

$$(x', y', z') = ((x \cos \theta + y \sin \theta), (-x \sin \theta + y \cos \theta), z' = z),$$

whereas θ represents the angle about the z -axis and is calculated through computing the tangent of the right triangle: $\tan(\theta) = \text{opposite}/\text{adjacent}$; the opposite is the y -coordinate of the endpoint, and the adjacent is the x -coordinate of the endpoint in the rescaled data.

Based on this rescaled and transformed data, we first calculated AUC as an indicator of attention. AUC is an aggregate indicator of how much information (e.g., warning information) attracted users' attention while moving (e.g., while dismissing a warning). AUC is the geometric area between a user's movement trajectory and the idealized response trajectory [23]. An *idealized response trajectory* is a straight line between a user's starting position (when a warning is shown) and the final destination point (e.g., the button is clicked to dismiss the warning). When information from a warning catches users' attention, it causes them to deviate from this straight line (the idealized response trajectory) as the cursor moves toward the warning information and explores the page. AUC is a measure of the total deviation from the idealized response trajectory; that is, the total attention

¹ See <http://www.mathematics-online.org/inhalt/erlaeuterung/erlaeuterung209/>

given to other information [23]. AUC was estimated through calculating Riemann sums,² which approximates the area between an unknown curve and the x -axis through creating rectangles for each x .

Second, we calculated initial acceleration toward dismissing the warning as an indicator of attention and habituation. Waiting for the body to perceive, comprehend, and program a response to a stimulus (a process often referred to as the *feedback loop*) is relatively time consuming. Therefore, if a person encounters a familiar stimulus (i.e., matches a mental model), the body does not wait for the feedback loop to complete before responding. Rather, based on past experience, the body automatically programs and executes a motor response to the stimulus [56]. Compared to movements guided by the feedback loop, feedforward movements occur very “rapidly, as there is no need to account for the delay of feedback loops” [56, p. 1775, p. 1775]. For example, if a person encounters a security warning that she or he is habituated to, the nervous system automatically starts programming a motor response even before the person reads the security information, resulting in faster initial acceleration. The more a person relies on a mental model (i.e., the more habituated a person is), the more the brain will rely on feedforward movements rather than the feedback loop [62], and the faster the initial acceleration will be.

Initial acceleration was calculated for a 150-ms window starting 75 ms after the warning was first displayed. On average, the brain requires approximately 70–80 ms to recognize a visual stimulus change. This is known as the *visual awareness period* [16]. To account for this reaction time, we began calculating acceleration immediately after the visual awareness period (around 75 ms). We stopped calculating acceleration 150 ms after this visual awareness period to gauge participants’ initial reactions to the visual stimulus before they were able to cognitively process the warning’s

² See http://mathinsight.org/calculating_area_under_curve_riemann_sums

content [65]. Acceleration was calculated using a standard formula: $acceleration = \frac{\Delta Velocity}{\Delta Time}$

whereas velocity was calculated as: $velocity = \frac{\Delta Distance}{\Delta Time}$.

Finally, we calculated movement speed as an indicator of attention and habituation. When a warning garners a user's attention under conditions of low habituation, the user engages in a comprehension process, that is, a process of understanding the intended meaning of the warning [70]. The comprehension process requires mental effort and consumes the user's limited cognitive resources [31]. Likewise, the brain requires cognitive resources to visually guide the hand (i.e., the mouse cursor) to its destination [42, 43]. When this cognitive capacity is decreased, hand movement precision also decreases, resulting in greater deviations from one's intended trajectory. One way the brain automatically compensates for this decrease in precision is to decrease the speed of movements [42, 43]. Hand (i.e., cursor) movement speed and precision are inversely related [52]; when speed is reduced, movement precision is increased as the brain has more time to perceive and program needed corrections, allowing it to operate optimally under cognitive demanding situations [42, 43]. However, when attention is low and habituation occurs, the comprehension stage is short or even non-existent [31]. The cognitive resources that would have been required to comprehend the security information are therefore available for other tasks, such as guiding the mouse cursor to a desired destination. As a result, the brain utilizes these resources to enable faster movements, which minimizes average total movement time [42, 43]. This average speed was calculated for the entire interaction (from the time the warning was shown until the user clicked Install or Cancel) using the following formula:

$$average\ speed = \frac{Total\ Distance\ Traveled}{Total\ Movement\ Time}$$

APPENDIX F: TECHNICAL DETAILS FOR THE MAN-IN-THE-MIDDLE ATTACK IN EXPERIMENT 2

For the web-browsing task in Experiment 2, if participants searched Google with a query related to Chrome weather extensions, the search returned our spoofed results. Furthermore, each of the links in the spoofed results pointed to sites under our control, despite their legitimate URLs. To accomplish this technological sleight of hand, we carried out a man-in-the-middle attack [47]. This was accomplished by having all participants access the Internet through our router in the laboratory, rather than having them connect through the regular campus network. Our router, upon assigning an Internet protocol (IP) address via Dynamic Host Configuration Protocol (DHCP), poisoned each participant's DNS configuration to use only the server under our control.³ Our DNS server, in turn, directed all traffic for www.google.com and the 40 weather extension uniform resource locators (URLs) to our own web server. We configured our web server to examine each hijacked web traffic request, and to determine whether we would serve our spoofed results, such as our fake Google search results and versions of the weather extension pages, or pass the traffic through to the legitimate host. This way, only Google search results matching terms similar to "Chrome weather extension" would result in our spoofed traffic being served. Any search query not of interest to us would return results of a legitimate, live Google search.

We caused traffic headed for the Chrome Extension Web Store (<https://chrome.google.com/webstore>) to redirect to an error page hosted on our server. This was done via en-route URL rewrites and HTML DOM manipulation. The error page announced to visitors that the Chrome Web Store was down for maintenance. It was important that we blocked participants from visiting the Chrome Web Store, since we were spoofing the in-line Chrome extension installation process as opposed to the in-store process.

³ A DNS server is responsible for resolving domain names, such as www.google.com, to server IP addresses.

We could not spoof Google's HTTPS certificates because we did not have a root-signing certificate. To circumvent SSL certificate errors, our Chrome extension also transparently redirected HTTPS traffic to HTTP. For instance, traffic to <https://www.google.com> was rewritten as <http://www.google.com>.

APPENDIX G: GRADED MOTOR RESPONSE ANALYSIS

As a supplemental analysis, we performed a graded motor response analysis in the time course of participants responding to security warnings. This analysis allows us to explore how much people deviate from the idealized response trajectory throughout their response, rather than just an aggregated measure of total deviation at the end, which is the case with AUC. Thus, this supplemental analysis lends understanding of whether warning information sustains greater attention throughout the movement for polymorphic warnings than static warnings. The analysis was adapted from Dale *et al.* [12]. Similarly to Dale *et al.* [12], all trajectories were first normalized to 101 time steps and translated to begin at the x, y-coordinate of (0,0). Each time step consisted of an averaged x, y-coordinate for that segment of the response trajectory. This normalization is required to average the full trajectories from multiple responses, allowing us to compare the movement trajectories of polymorphic warnings and static warnings.

Whereas Dale *et al.* [12] and similar studies have compared trajectories in a very controlled setting (e.g., moving the mouse always from the bottom middle of the screen to the upper corners of the screen to answer a question), our study was in a less controlled environment (people moved the mouse from wherever it was when the warning appeared to a location to dismiss the warning). As such, a few modifications were made to the traditional graded motor response analysis to accommodate this difference. Primarily, rather than using the x-coordinate as the primary measure of deviation between trajectories in different treatments, we used the distance from the actual trajectory to the idealized response trajectory for each time step as a measure of deviation. This allowed us to calculate deviation even if people were moving to dismiss warnings from different locations on the page.

To do this, we derived an equation for a straight line (the idealized response trajectory) between the point where the cursor was when the warning first appeared to the endpoint where the

user ultimately dismissed the warning. For each time step, we calculated the distance between the averaged x, y-position of the actual trajectory to the idealized response trajectory equation.

Consistent with past literature, we then explored whether the amount of deviation at each time step was statistically different using *t*-tests. Eight consecutive time slots that are statistically different between trajectories are considered a conservative indicator of significance to avoid alpha slippage [12]. Our analysis demonstrated statistically significant differences ($p < 0.05$) from time step 5 to 88 (94 consecutive time stamps), providing strong support that deviation was greater for participants in the polymorphic treatment than for participants in the static warning treatment almost throughout the entire interaction. The results are visualized in Figure H1.

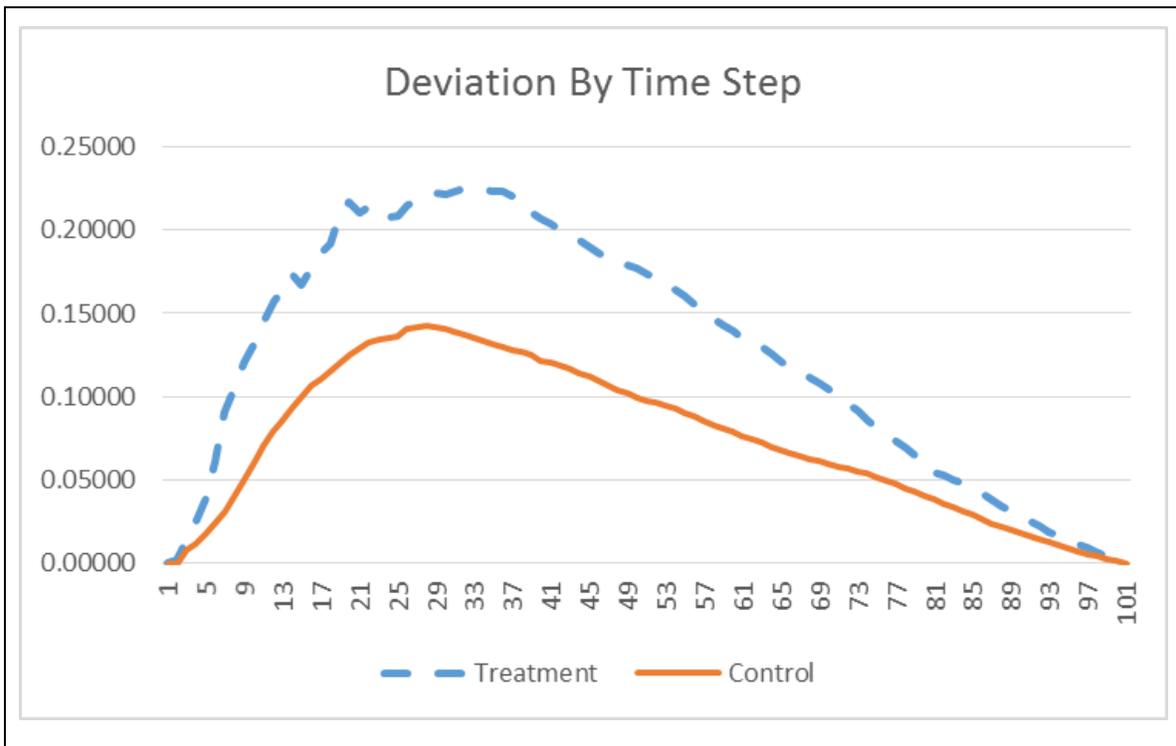


Figure F1. Graded motor response analysis. The x-axis shows the number of time steps in the analysis; the y-axis depicts the amount of deviation from the idealized response trajectory by time over the total movement.

ONLINE APPENDICES REFERENCES

1. Adams, J.A. Human tracking behavior. *Psychological Bulletin*, 58, 1 (1961), 55-79.
2. Arroyo, E., Selker, T. and Wei, W. Usability tool for analysis of web designs using mouse tracks. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems*, Montreal, Canada: ACM, 2006, pp. 484-489.
3. Atterer, R. and Lorenzi, P. A heatmap-based visualization for navigation within large web pages. In *Proceedings of the 5th Nordic conference on Human-Computer Interaction*, Lund, Sweden: ACM, 2008, pp. 407-410.
4. Barca, L. and Pezzulo, G. Unfolding visual lexical decision in time. *PLoS One*, 7, 4 (2012), e35932.
5. Bartolotti, J. and Marian, V. Language learning and control in monolinguals and bilinguals. *Cognitive Science*, 36, 6 (2012), 1129-1147.
6. Bruhn, P. Emergence of spontaneous anticipatory hand movements in a probabilistic environment. *Advances in Cognitive Psychology*, 9, 2 (2013), 62-73.
7. Byrne, M.D., Anderson, J.R., Douglass, S. and Matessa, M. Eye tracking the visual search of click-down menus. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, Pittsburgh, Pennsylvania, USA: ACM (1999), pp. 402-409.
8. Chen, M.C., Anderson, J.R. and Sohn, M.H. What can a mouse cursor tell us more?: Correlation of eye/mouse movements on web browsing. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, Seattle, Washington: ACM (2001), pp. 281-282.
9. Cloutier, J., Freeman, J.B. and Ambady, N. Investigating the early stages of person perception: The asymmetry of social categorization by sex vs. Age. *PLoS One*, 9, 1 (2014), e84677.
10. Cox, A.L. and Silva, M.M. The role of mouse movements in interactive search. In *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*, Vancouver, British Columbia: Cognitive Science Society, 2006, pp. 1156-1161.
11. Cox, R.W. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29, 3 (1996), 162-173.
12. Dale, R., Kehoe, C. and Spivey, M.J. Graded motor responses in the time course of categorizing atypical exemplars. *Memory & Cognition*, 35, 1 (2007), 15-28.
13. Dale, R., Roche, J., Snyder, K. and McCall, R. Exploring action dynamics as an index of paired-associate learning. *PLoS One*, 3, 3 (2008), e1728.
14. Dshemuchadse, M., Scherbaum, S. and Goschke, T. How decisions emerge: Action dynamics in intertemporal decision making. *Journal of Experimental Psychology*, 142, 1 (2013), 93-100.
15. Duran, N.D., Dale, R. and McNamara, D.S. The action dynamics of overcoming the truth. *Psychonomic Bulletin & Review*, 17, 4 (2010), 486-491.
16. Eagleman, D.M. and Sejnowski, T.J. Motion integration and postdiction in visual awareness. *Science*, 287, 5460 (2000), 2036-2038.
17. Farmer, T.A., Cargill, S.A., Hindy, N.C., Dale, R. and Spivey, M.J. Tracking the continuity of language comprehension: Computer mouse trajectories suggest parallel syntactic processing. *Cognitive Science*, 31, 5 (2007), 889-909.
18. Faulkenberry, T.J. How the hand mirrors the mind: The embodiment of numerical cognition. In *Proceedings of the 40th Annual Meeting of the Research Council on Mathematics Learning*, Tulsa, Oklahoma, 2013, p. 205.
19. Fleishman, E.A. and Hempel Jr, W.E. Factorial analysis of complex psychomotor performance and related skills. *Journal of Applied Psychology*, 40, 2 (1956), 96-104.
20. Fleishman, E.A. Factor structure in relation to task difficulty in psychomotor performance. *Educational and Psychological Measurement*, 17, 1 (1957), 522-532.
21. Freeman, J.B., Ambady, N., Rule, N.O. and Johnson, K.L. Will a category cue attract you? Motor output reveals dynamic competition across person construal. *Journal of Experimental Psychology*, 137, 4 (2008), 673-690.

22. Freeman, J.B. and Ambady, N. Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological Science*, 20, 10 (2009), 1183-1188.
23. Freeman, J.B. and Ambady, N. Mousetracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behavior Research Methods*, 42, 1 (2010), 226-241.
24. Freeman, J.B. and Ambady, N. When two become one: Temporally dynamic integration of the face and voice. *Journal of Experimental Social Psychology*, 47, 1 (2011), 259-263.
25. Freeman, J.B., Dale, R. and Farmer, T.A. Hand in motion reveals mind in motion. *Frontiers in Psychology*, 2, 59 (2011), 1-6.
26. Freeman, J.B. and Dale, R. Assessing bimodality to detect the presence of a dual cognitive process. *Behavior Research Methods*, 45, 1 (2012), 1-15.
27. Freeman, J.B. Abrupt category shifts during real-time person perception. *Psychonomic Bulletin & Review*, 21, 1 (2014), 85-92.
28. Grossman, E. The information-capacity of the human motor-system in pursuit tracking. *Quarterly Journal of Experimental Psychology*, 12, 1 (1960), 01-16.
29. Guilford, J.P. and Lacey, J.I. Army air forces aviation psychology program research reports: Printed classification tests. DTIC Document, 1947.
30. Guo, Q. and Agichtein, E. Towards predicting web searcher gaze position from mouse movements. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, Austin, TX: ACM, 2010, pp. 3601-3606.
31. Hancock, H., Bowles, C., Rogers, W. and Fisk, A. Comprehension and retention of warning information. In: M. Wogalter (ed.), *Handbook of Warnings*. Mahwah NJ: Lawrence Erlbaum Associates, 2006, pp. 267-277.
32. Hehman, E., Stoller, R.M. and Freeman, J.B. Advanced mouse-tracking analytic techniques for enhancing psychological science. *Psychological Science*, 20, 10 (2014), 1183-1188.
33. Hibbeln, M., Jenkins, J.L., Schneider, C., Valacich, J.S. and Weinmann, M. How is your user feeling? Inferring emotion through human-computer interaction devices. *MIS Quarterly* (forthcoming).
34. Huang, J., White, R.W. and Dumais, S. No clicks, no problem: Using cursor movements to understand and improve search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Vancouver, BC, Canada: ACM, 2011, pp. 1225-1234.
35. Huang, J., White, R. and Buscher, G. User see, user point: Gaze and cursor alignment in web search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Austin, Texas, USA: ACM (2012), pp. 1341-1350.
36. Lagun, D. and Agichtein, E. Viewser: Enabling large-scale remote user studies of web search examination and interaction. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Beijing, China: ACM, 2011, pp. 365-374.
37. Lettner, F. and Holzmann, C. Heat maps as a usability tool for multi-touch interaction in mobile applications. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia*, Ulm, Germany: ACM, 2012, pp. 1-2.
38. Maehr, W. eMotion: Estimation of user's emotional state by mouse motions. Saarbrücken, Germany: VDM Verlag, 2008.
39. Martens, M.A., Hasinski, A.E., Andridge, R.R. and Cunningham, W.A. Continuous cognitive dynamics of the evaluation of trustworthiness in williams syndrome. *Frontiers in Psychology*, 3, 160 (2012), 1-9.
40. McKinstry, C., Dale, R. and Spivey, M.J. Action dynamics reveal parallel competition in decision making. *Psychological Science*, 19, 1 (2008), 22-24.
41. Metcalfe, J., Eich, T.S. and Miele, D.B. Metacognition of agency: Proximal action and distal outcome. *Experimental Brain Research*, 224, 1 (2013), 1-12.
42. Meyer, D.E., Abrams, R.A., Kornblum, S., Wright, C.E. and Keith Smith, J. Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review*, 95, 3 (1988), 340-370.

43. Meyer, D.E., Smith, J.E.K., Kornblum, S., Abrams, R.A. and Wright, C.E. Speed-accuracy tradeoffs in rapid aimed movements: Toward a theory of rapid voluntary action. In: M. Jeannerod (ed.), *Attention and Performance XIV* Hillsdale, NJ: Lawrence Erlbaum Associates, 1990, pp. 173-226.
44. Motley, S.E. and Kirwan, C.B. A parametric investigation of pattern separation processes in the medial temporal lobe. *The Journal of Neuroscience*, 32, 38 (2012), 13076-13084.
45. Mueller, F. and Lockerd, A. Cheese: Tracking mouse movement activity on websites, a tool for user modeling. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, Seattle, Washington: ACM (2001), pp. 279-280.
46. Navalpakkam, V. and Churchill, E. Mouse tracking: Measuring and predicting users' experience of web-based content. In *2012 ACM Annual Conference on Human Factors in Computing Systems*, Austin, Texas, USA: ACM (2012), pp. 2963-2972.
47. OWASP. Man in the middle attack. https://www.owasp.org/index.php/Man-in-the-middle_attack (accessed on 8/20/14)
48. Palmer, C.J., Paton, B., Barclay, L. and Hohwy, J. Equality, efficiency, and sufficiency: Responding to multiple parameters of distributive justice during charitable distribution. *Review of Philosophy and Psychology*, 4, 3 (2013), 1-16.
49. Pan, B., Hembrooke, H.A., Gay, G.K., Granka, L.A., Feusner, M.K. and Newman, J.K. The determinants of web page viewing behavior: An eye-tracking study. In *2004 Symposium on Eye Tracking Research & Applications*, San Antonio, Texas: ACM (2004), pp. 147-154.
50. Papesh, M.H. and Goldinger, S.D. Memory in motion: Movement dynamics reveal memory strength. *Psychonomic Bulletin & Review*, 19, 5 (2012), 906-913.
51. Parker Jr, J.F. and Fleishman, E.A. Ability factors and component performance measures as predictors of complex tracking behavior. *Psychological Monographs: General and Applied*, 74, 16 (1960), 1-36.
52. Plamondon, R. and Alimi, A.M. Speed/accuracy trade-offs in target-directed movements. *Behavioral and Brain Sciences*, 20, 2 (1997), 279-303.
53. Poulton, E. The basis of perceptual anticipation in tracking. *British Journal of Psychology. General Section*, 43, 4 (1952), 295-302.
54. Rodden, K., Fu, X., Aula, A. and Spiro, I. Eye-mouse coordination patterns on web search results pages. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, Florence, Italy: ACM (2008), pp. 2997-3002.
55. Rodrigues, M., Gonçalves, S., Carneiro, D., Novais, P. and Fdez-Riverola, F. Keystrokes and clicks: Measuring stress on e-learning students. In: J. Casillas, F.J. Martinez-Lopez, R. Vicari and F.D.I. Prieta (eds.), *Management intelligent systems*. Switzerland: Springer, 2013, pp. 119-126.
56. Seidler, R., Noll, D. and Thiers, G. Feedforward and feedback processes in motor control. *NeuroImage*, 22, 4 (2004), 1775-1783.
57. Shah, K.N., Rathod, K.R. and Agravat, S.J. A survey on human computer interaction mechanism using finger tracking. *ArXiv e-prints*, 1402.0693 (2014).
58. Solman, G.J.F., Cheyne, J.A. and Smilek, D. Found and missed: Failing to recognize a search target despite moving it. *Cognition*, 123, 1 (2012), 100-118.
59. Song, J.H. and Nakayama, K. Role of focal attention on latencies and trajectories of visually guided manual pointing. *Journal of Vision*, 6, 9 (2006), 982-995.
60. Song, J.H. and Nakayama, K. Target selection in visual search as revealed by movement trajectories. *Vision Research*, 48, 7 (2008), 853-861.
61. Spivey, M.J., Grosjean, M. and Knoblich, G. Continuous attraction toward phonological competitors. *National Academy of Sciences of the United States of America*, 102, 29 (2005), 10393-10398.
62. Staddon, J.E.R. and Higa, J.J. Multiple time scales in simple habituation. *Psychological Review*, 103, 4 (1996), 720-733.
63. Talairach, J. and Tournoux, P. Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system: An approach to cerebral imaging. Stuttgart: Thieme, 1988.

64. Taylor, F.V. and Birmingham, H.P. Studies of tracking behavior. II. The acceleration pattern of quick manual corrective responses. *Journal of Experimental Psychology*, 38, 6 (1948), 783-795.
65. Thorpe, S., Fize, D. and Marlot, C. Speed of processing in the human visual system. *Nature*, 381, 6582 (1996), 520-522.
66. Tower-Richardi, S.M., Brunyé, T.T., Gagnon, S.A., Mahoney, C.R. and Taylor, H.A. Abstract spatial concept priming dynamically influences real-world actions. *Frontiers in Psychology*, 27, 361 (2012), 1-12.
67. Wang, Q., Taylor, H.A. and Brunyé, T.T. When going the right way is hard to do: Distinct phases of action compatibility in spatial knowledge development. *Acta psychologica*, 139, 3 (2012), 449-457.
68. Weaver, S.M. and Arrington, C.M. Tracking the multitasking mind. *Zeitschrift für Psychologie*, 221, 1 (2013), 51-60.
69. Weis, L. The honest hand - how computer mouse trajectories might capture deceptive processes. University of Aberdeen, 2012.
70. Wogalter, M. and Vigilante Jr, W.J. Attention switch and maintenance. In: M.S. Wogalter (ed.), *Handbook of Warnings*. Mahwah NJ: Lawrence Erlbaum Associates, 2006, pp. 245-266.
71. Wojnowicz, M.T., Ferguson, M.J., Dale, R. and Spivey, M.J. The self-organization of explicit attitudes. *Psychological Science*, 20, 11 (2009), 1428-1435.
72. Yu, Z., Wang, F., Wang, D. and Bastin, M. Beyond reaction times: Incorporating mouse-tracking measures into the implicit association test to examine its underlying process. *Social Cognition*, 30, 3 (2012), 289-306.
73. Zimmermann, P., Guttormsen, S., Danuser, B. and Gomez, P. Affective computing-a rationale for measuring mood with mouse and keyboard. *International Journal of Occupational Safety and Ergonomics*, 9, 4 (2003), 539-551.
74. Zimmermann, P., Gomez, P., Danuser, B. and Schär, S. Extending usability: Putting affect into the user-experience. In *Proceedings of the 4th Nordic Conference on Human-Computer Interaction*: ACM, 2006, pp. 27-32.
75. Zushi, M., Miyazaki, Y. and Norizuki, K. Web application for recording learners' mouse trajectories and retrieving their study logs for data analysis. *Knowledge Management & E-Learning: An International Journal*, 4, 1 (2012), 37-50.