What Does the Brain Tell Us About the Timing of Security Messages? Insights from an fMRI and Web-based Experiment

Abstract

We explore how adherence to security messages is affected by dual-task interference (DTI), a cognitive limitation in which even simple tasks cannot be simultaneously performed without significant performance loss. Using functional magnetic resonance imaging (fMRI), we show that displaying a security message at a high-DTI time substantially reduces neural activation in the medial temporal lobe (MTL)—a brain region associated with memory—and significantly reduces adherence. In contrast, displaying the security message at a low-DTI time yields substantially higher adherence. We applied these findings to the context of a web browser, specifically the Chrome Cleanup Tool (CCT) message in Google Chrome for Windows. We identified low- and high-DTI times common to browsing and found that displaying the CCT message at low-DTI times resulted in 36% higher adherence on average compared to high-DTI times. Our findings indicate that, where possible, displaying security messages at low-DTI times can result in substantially higher adherence.

1 Introduction

The negative impact of interruptions is well recognized in the HCI literature [24-26]. Nevertheless, most security messages are displayed without any regard for the workflow of the user. Although some security messages are tied to a specific event or user action (such as SSL and malware browser warnings), others are not (e.g., software update prompts and malware scan notifications; see Figure 1). This latter class of security messages can be delayed until a time when the user is better equipped to respond.

In this study, we explore how adherence to security messages is affected by *dual-task interference* (DTI), a limitation of the human cognitive system [41]. Much like a computer's central processing unit (CPU), the human brain processes many tasks serially and, therefore, must rapidly switch attention between multiple tasks that are performed at the same time [33]. However, unlike a CPU, research shows that when people attempt even simple tasks simultaneously, the tasks can "interfere with each other quite drastically, even though they are neither intellectually challenging nor physically incompatible" [33].

Although DTI has been researched in a variety of contexts [41], it is unclear (1) whether interruptive

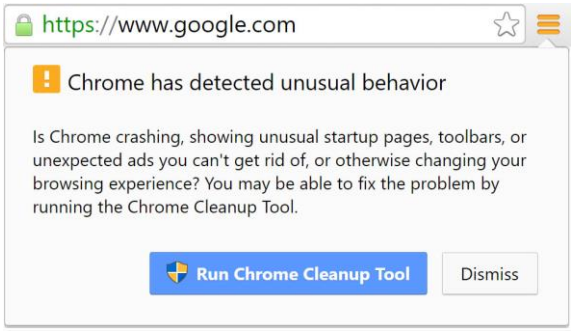

Figure 1. The Chrome Cleanup Tool message in Google Chrome for Windows version 45.

security messages induce DTI and if this in turn affects adherence. Further, it is unknown (2) whether finessing the timing of security messages can reduce DTI and improve adherence. The objectives of this study are to address these two gaps.

First, we conducted a functional magnetic resonance imaging (fMRI) experiment to observe DTI as it occurs in the brain in response to a security message that interrupts a competing primary task. Neuroscience has proved a useful lens through which to observe security behavior [29; 45], particularly fMRI [3; 30]. We found that (1) adherence to the security message is significantly worse during high-DTI times than when responding to the security message is the only task; (2) in a whole-brain analysis, activation in the *medial temporal lobe* (MTL)—a brain region associated with declarative memory—is much lower when compared to responding to the security message only; and (3) reduced activity in the MTL significantly predicts the degradation of performance relative to the security message.

Second, HCI research on interruptions suggests that the severity of an interruption can be reduced by introducing the interruption at a more opportune moment (i.e., when DTI is low) [1; 14; 20]. Accordingly, we examined whether DTI can be reduced by displaying the security message at a low-DTI time. In this case, we show that (1) the level of adherence is not different when responding to the message is the sole task, (2) activation in the MTL is much higher compared to the high-DTI condition, and (3) the change in activation in the MTL between the low- and high-DTI conditions significantly predicts the difference in adherence between the two conditions.

We applied our fMRI findings to the context of the Chrome Cleanup Tool (CCT) message in Google Chrome for Windows version 45, which is currently displayed without regard for the activity of the user. In collaboration with members of the Google Chrome security team, we identified six low- and high-DTI times common to the web browsing experience. Using 572 participants via Amazon Mechanical Turk (MTurk), we tested how adherence to the CCT is affected depending on whether it interrupts at a low- or high-DTI time. We found that on average, displaying the CCT message at low-DTI times resulted in an adherence rate of 63% compared to 23% for high-DTI times. Our findings indicate that, for those security messages that can be delayed, displaying security messages at low-DTI times can substantially improve adherence.

2 Literature Review

Understanding why users frequently fail to appropriately respond to security messages is a significant concern for security researchers [12]. Failing to adhere to a security message has been shown to be due to a variety of factors, such as habituation [3], lack of comprehension [12], and conscious decisions to ignore security messages [11].

Security messages often interrupt users while they are performing other primary tasks on a computer, such as completing a work-related project or using the computer for entertainment [46]. A research area known as interruption science [31] largely supports the notion that interruptions decrease users' performance of their primary task [e.g., 8; 37]. Interruptions during computing tasks result in reduced productivity [27], increased stress [26], and increased time being required to complete the task [21].

Scholars provide several theoretical arguments explaining why interruptions decrease performance on primary tasks. For example, interruptions decrease the time available to work on the primary task [22], cause people to overlook information cues in the primary task [37], and cause people to forget their primary goal [8]. Research has identified three main factors that influence the severity of interruptions: (1) the delay of interruption, (2) the complexity of the interrupting secondary task, and (3) the timing of the interruption [4]. In the case of timing, researchers have sought to minimize the impact of an interruption by presenting the interruption at an opportune moment [1; 27]. An ongoing stream of research has examined how to best identify such moments. Examples of approaches used include observing mouse and keyboard events [18], monitoring audio and video [32], textual analysis of incoming messages [14], statistical modeling [1; 14], and physiological approaches, such as pupil dilation [19], table-top pressure [43], and electroencephalography (EEG) and electrodermal activity (EDA) [48]. This research demonstrates that interruptibility (the interval between primary tasks) can be measured using a variety of techniques with high accuracy, and that the impact of interruptions on primary tasks can be reduced through intelligent timing.

Little research, however, examines why performance of the secondary task (i.e., the interruption) may also decrease. Unlike the secondary tasks in prior literature, in the context of security, it is critical that users attend to and perform well on the secondary task as failure to comply may result in more serious consequences than those of the primary task. Figure 2 describes the gap we address in this study.

We posit that DTI is a theoretical lens for explaining why adherence may decrease when security messages interrupt other tasks. DTI is a neural phenomenon that explains why people have trouble performing two or more relatively simple tasks concurrently [33]. It explains performance decrements in a variety of contexts, including driving while talking on the phone [40], searching congruently for multiple pieces of information [28], processing visual and verbal information together [17], and texting while walking [34]. Normally, people are not aware of tasks interfering with each other unless the two tasks are cognitively difficult, physically incompatible, or evoke negative emotional reactions. However, studies demonstrate that just the opposite is actually true. For example, when people are involved in even simple cognitive tasks, they cannot process information or perform behaviors related to other tasks as quickly or effectively [e.g., 23].

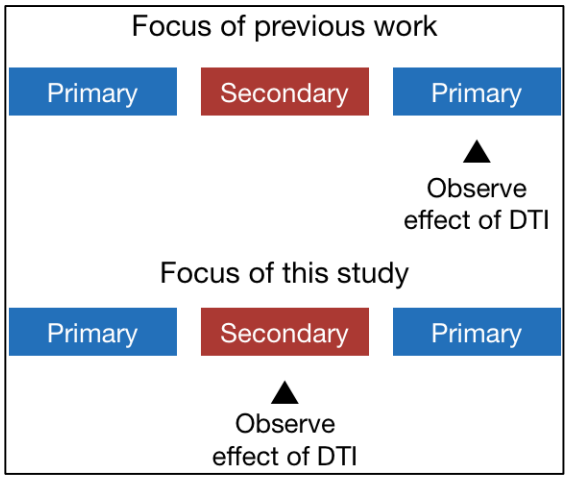

Figure 2. Observing the effect of dual-task interference (DTI) on the secondary security task rather than the primary task.

DTI typically occurs under one of two paradigms: bisensory and divided attention [41]. Under the bisensory paradigm, people engage in two tasks simultaneously, such as walking and talking at the same time. In contrast, under the divided attention paradigm, participants must switch attention between stimuli, such as in the context of this study, when a security message interrupts a primary task.

From a neurocognitive perspective, research has found that DTI may result from tasks competing for the same brain functions [35]. Effectively, humans have limited cognitive abilities; when multiple tasks compete for the same cognitive functions, fewer resources are available for each individual task, and so task performance subsequently decreases. This view is commonly known as the *capacity sharing model* [44]. In the context of security, although a security message may interrupt and demand attention from a user for a short time, people do not necessarily free adequate cognitive resources associated with the primary task to respond effectively to the message [e.g., 13]. DTI occurs as cognitive functions are still engaged in the primary task while responding to security messages.

Security literature suggests this relationship. Specifically, Yee [47] suggests that "interrupting users with prompts presents security decisions in a terrible context: it teaches users that security issues obstruct their main task and trains them to dismiss prompts quickly and carelessly." One reason users may choose to dismiss security messages quickly and carelessly within this context is that it is cognitively difficult for them to switch between their primary task and optimally addressing the security message. Bravo-Lillo et al. [5] suggest that users often ignore or suboptimally address interrupting prompts because users have a limited cognitive ability to switch between tasks. Consequently, users unconsciously ignore and fail to comprehend security messages [9; 13].

In this paper, we hypothesize that the capacity sharing model in the DTI paradigm helps account for security message adherence. We test our hypotheses both behaviorally and through the use of fMRI. By doing so, we establish a theoretical foundation for designing and presenting security messages to reduce the effects of DTI. We extend previous literature that shows interruptions may decrease performance of primary tasks by explaining why performance may also decrease for the interrupting secondary task (e.g., the security message).

3 Hypotheses

We predict that when a security message interrupts another task, security message adherence will decrease. The capacity sharing model explains that people have finite cognitive resources. The utilization of cognitive resources may inhibit a person from using the same or related brain regions for a secondary task [44]. Applying this to a security message setting, when systems show security messages during another task that utilizes working memory, this may inhibit a user's ability to process that security message and recall information from declarative memory. This neurocognitive limitation will decrease security message adherence, as users may not be able to recall or process the information necessary to respond appropriately. In summary,

H1. When a security message (secondary task) interrupts a primary task, security message adherence decreases.

Conversely, when a security message does not interrupt another task (e.g., is presented between primary tasks), responding to the security message does not compete for cognitive resources. One has access to the necessary cognitive resources to respond more appropriately to the security message, resulting in greater security message adherence. In summary,

H2. When a security message (secondary task) is displayed between (rather than during) primary tasks, security message adherence increases.

4 Experiment 1—fMRI

We used fMRI to test our hypotheses. We used fMRI to validate our behavioral data that DTI was in fact happening in the brain and because fMRI has superior spatial resolution to locate neural correlates of DTI. Further, fMRI is an accepted method in neuroscience for studying DTI [41] and has already yielded important insights into information security behaviors [3; 30].

From a neural perspective, we examined how DTI influences activation in the medial temporal lobe (MTL), the brain region associated with declarative memory, which is needed to represent information to respond properly to the security message. The MTL declarative memory system includes the hippocampus and adjacent cortical structures (entorhinal, perirhinal, and parahippocampal cortices).

We used a working memory task as the primary task to induce DTI. The system first asked participants to memorize (encode) a 7-digit code. This is similar to other common working memory tasks, such as dialing a 7-digit phone number. After a short encoding time, participants were given a brief rehearsal period in which they were required to maintain what they learned in working

memory. Finally, the system asked participants to retrieve the code. Following the reasoning supporting our hypotheses, we predicted that if an interruption occurred during the rehearsal period, performance would be likely to decrease because the working memory rehearsal process would be disrupted.

We chose a task that consumed working memory because many computer tasks have similarly high working memory demands. Maintaining information in the working memory requires brain structures including the hippocampus and amygdala [15] and utilizes MTL areas including the perirhinal cortex (PrC), entorhinal cortex (EC), and the CA1 region of the hippocampus [36]. This suggests that other MTL processes, such as memory retrieval, may be inhibited during working memory maintenance periods. Our security task includes some elements of judgment as well as some elements of memory—common in typical security training.

We used an experimental design with two conditions: requiring a response to a security warning (1) in the middle of a primary task (the high-DTI condition), or (2) between two primary tasks (the low-DTI condition). We presented conditions using HTML in a web browser running on the MRI experimental computer. JavaScript embedded in the HTML page listened for a sync pulse from the MRI machine to coordinate the display of the stimuli and the MRI events (similar to E-Prime software).

4.1 Methodology

We utilized a repeated-measure, within-subject design that required participants to respond to security messages during low-DTI (i.e., not during a rehearsal period) and high-DTI times (i.e., during a rehearsal period). We operationalized the security messages used in this experiment as application permission warnings similar

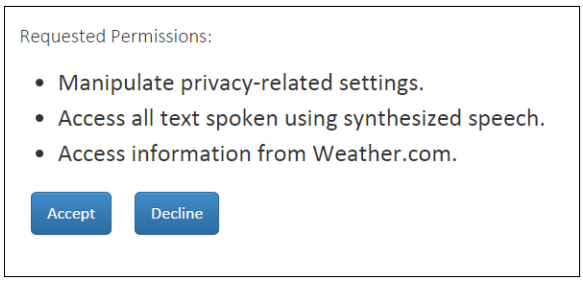

Figure 3. Example permission warning.

to those that are displayed as users install a Google Chrome browser extension (see Figure 3 for an example). The warning listed the permissions that the application was requesting.

Within the scanner, but prior to starting the experiment, participants were required to learn which permission warnings were reasonable or unreasonable in the context of a weather extension. For example, we classified the permission "Read and modify all your data on all websites you visit" as unreasonable in this context. During the experiment, we instructed participants to reject all warnings that contained any unreasonable permissions and to accept all others. To ensure that participants learned which permissions were unreasonable, they were required to complete a quiz with 100% accuracy.

After successfully finishing the training, participants completed the three conditions listed below (see Figure 4), presented in a random order. To relieve participants' fatigue during the experiment, there was a brief break in between each condition.

4.1.1 Condition A: High-DTI

In the high-DTI condition, the system presented participants with a 7-digit code. The system then asked participants to encode the code for 5 seconds. After 5

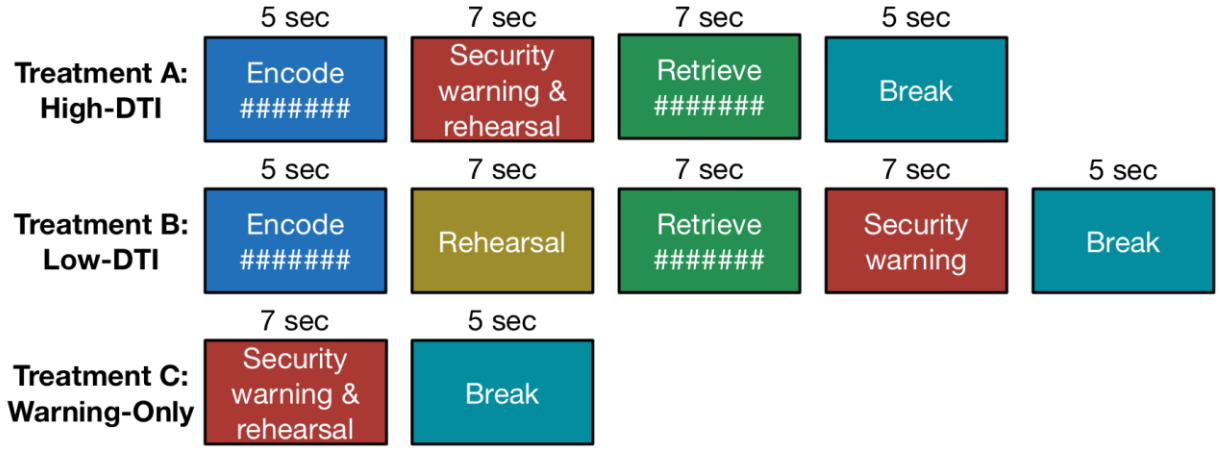

Figure 4. Experimental design.

seconds, the code disappeared and the system showed the warning during the rehearsal phase of the recall task. Participants were then given 7 seconds to click on either reject or accept based on their previous training regarding acceptable permissions. At this time, the warning disappeared and a question appeared asking participants to select the code they were most recently asked to memorize from among 5 other codes. Participants were given 7 seconds to select the code, and then they were given a break for 7 seconds to be used as a baseline in the analysis. Participants repeated this 18 times in Condition A. Since the warning was presented during a working memory maintenance period (i.e., between the encoding and retrieval screens), the performance of appropriately accepting and rejecting the warning (i.e., security message adherence) is likely influenced by DTI.

4.1.2 Condition B: Low-DTI

Condition B followed the same procedure as Condition A, except the system switched the 7-second warning page and the 7-second break page. Thus, participants first encoded the code, retained that code in their memory for the rehearsal period (to help ensure that performance differences were due to DTI rather than time differences in the conditions), retrieved the code, and then responded to the warning. The system repeated this 18 times with a 5-second break between each trial to be used as a baseline in the analysis. As the warning did not occur during the maintenance period, DTI influences performance in this condition less than in Condition A.

4.1.3 Condition C: Warning Only

In this condition, participants saw only warnings and did not receive the encode/retrieve task. Like the previous conditions, participants were given 7 seconds to respond to the warning. The system repeated this 18 times with a break between each trial, which we used as a baseline in the analysis. Since there was no memorization task, DTI likely does not influence security message adherence.

4.1.4 Pilot Test

Prior to administering our experiment in the fMRI scanner, we conducted a pilot test using Amazon's MTurk to see if the conditions resulted in a behavioral difference concerning security message adherence. In this pilot test, we followed the same procedure as that to be used in the fMRI data collection, except that participants completed only 4 repetitions in each condition.

Thirty-three individuals participated in the pilot test. Following the recommendations of Steelman et al. [39], we used a U.S.-based demographic to mirror our student-based population of native English speakers. The average age of participants was 32.2 years.

To test our manipulations, we conducted a mixed effects logistic regression model to account for the repeated measure nature of the data. In the model, we included adherence as the binary dependent variable (1 for adherence, 0 for not). We then included the participant ID and the condition order (e.g., which condition the participants saw 1st, 2nd, etc.) as random effects. We also included the conditions (as dummy variables) as fixed effects. The warning-only condition was specified as the reference or baseline group. Finally, we included the trial repetition number as a fixed effect (e.g., the *n*th warning within a given condition) to explore whether participants' behavior changed across trials as a result of habituation. Wald statistics were calculated for significance tests.

We found support for our manipulations of DTI. The low-DTI condition did not statistically differ from the warning-only condition, $z = -1.243$, $p > .05$, $\beta = -0.431$. However, the high-DTI condition had significantly lower adherence compared to the warning-only condition, $z = -3.667$, $p < .001$, $\beta = -1.213$. When treating the low-DTI condition as the reference group, the high-DTI condition also had significantly lower adherence than the low-DTI condition, $z = -2.539$, $p < .05$, $\beta = -0.782$. Furthermore, the trial repetition number did not significantly influence adherence, indicating that habituation did not influence the results, $z = 0.018$, $p > .05$, $\beta = 0.002$. The R^2 of the model was 0.176. The overall adherence percentages for each condition are shown in Table 1.

Although different conditions were presented in different blocks of the task both in the pilot test and in the subsequent fMRI task, users were not able to avoid the decrease in performance due to DTI in the high-DTI condition. This decrease in performance is consistent with our hypotheses of exceeding working memory capacity in the high-DTI condition, which we tested in the fMRI task.

4.1.5 MRI Procedure

Next, we ran the full experiment in an fMRI laboratory to investigate the neural correlates of DTI in our chosen tasks. fMRI has high spatial resolution and can localize neural activation to specific brain regions in a non-invasive manner (see Figure 5). We followed the guidelines provided by Dimoka [10] for our fMRI protocol.

Figure 5. Siemens 3T Tim Trio MRI scanner used in our experiment.

We briefed participants verbally about the MRI procedures and the experimental task. Participants viewed the experimental images on a large MR-compatible monitor at the opening of the MRI scanner by means of a mirror attached to the head coil. They used an MR-compatible trackball to interact with the system. The appendix provides technical details of the fMRI scans, experimental procedures, and the analyses. At the conclusion of the experiment, participants completed a brief post-test survey outside the MRI scanner.

4.1.6 Participants

We recruited 24 participants from the university community. We screened each participant to require MRI compatibility, native English speakers, corrected-normal visual acuity, and right-handedness. We excluded those with color blindness or those taking psychotropic medications. In accordance with the university's institutional review board (IRB) protocol, we gave all participants an informed consent form to sign. Of the 24 participants, 11 were female and 13 were male. Participant age ranged from 18–40 years of age, with a mean age of 23.7 years. We paid participants $25 for approximately 1 hour in the scanner.

4.2 Hypothesis 1 Analysis

We hypothesized that when a security message (secondary task) interrupts a primary task, security message adherence decreases. Both the fMRI and behavioral analysis support this hypothesis, as described below.

Condition	Average Security Warning Adherence
High-DTI	63.08%
Low-DTI	76.61%
Warning-Only	82.26%

Table 1. Pilot test warning performance.

4.2.1 fMRI Analysis

We examined the neural correlates of responding to security warnings under dual-task conditions by comparing activation for the high-DTI warning/rehearsal period (in which participants were required to maintain a 7-digit code in their working memory and respond to the warning stimulus) with activation for the warning in the warning-only condition using paired *t*-tests. We exclusively masked the results of this comparison with the warning vs. baseline comparison to eliminate spurious activations (such as visual responses to the stimulus and motor responses from manipulating the trackball).

In this whole-brain analysis, we observed significant clusters of activation in the MTL only. In each case, activation was greater for the warning-only and low-DTI conditions than for the high-DTI condition ($ps < .05$), suggesting that participants were utilizing the MTL less for processing the security warning in the high-DTI condition (see Figure 6).

4.2.2 Behavioral Analysis

In addition to the MRI analysis, we explored how DTI influences participants' actual security message adherence. We specified the same mixed effects logistic regression model as specified in the pilot test to account for the repeated measure nature of the data. Our results again showed that the low-DTI condition did not statistically differ from the warning-only condition, $z = -0.581$, $p > .05$, $\beta = -0.162$. However, the high-DTI condition had significantly lower adherence compared to the warning-only condition, $z = -5.807$, $p < .001$, $\beta = -1.346$. When treating the low-DTI condition as the reference group, the high-DTI condition also had significantly lower adherence than the low-DTI condition, $z = -5.581$, $p < .001$, $\beta = -1.184$. Furthermore, the trial repetition number did not significantly influence adherence, indicating that habituation did not influence the results, $z = -0.251$, $p > .05$, $\beta = -0.001$. The R^2 of the model was 0.296. The overall adherence percentages for each condition are shown in Table 2.

Condition	Average Security Warning Adherence
High-DTI	77.08%
Low-DTI	91.20%
Warning-Only	92.59%

Table 2. fMRI warning performance.

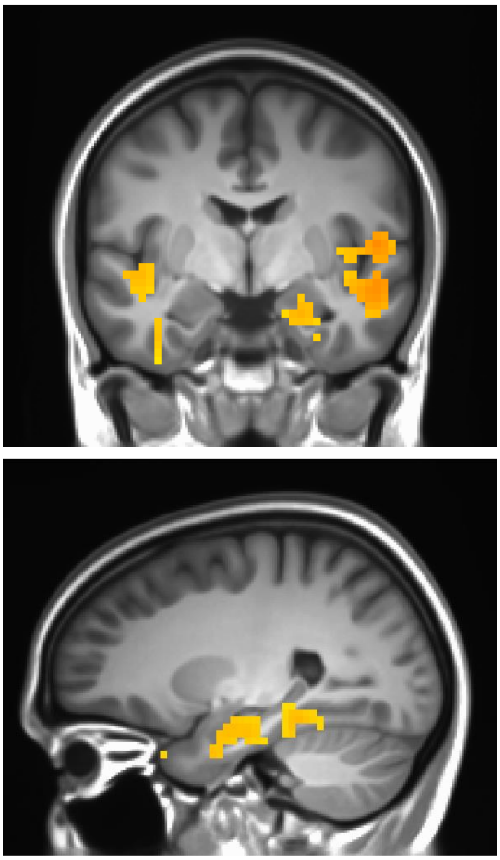

Figure 6. Decreased activity in response to the high-DTI condition compared to the warning-only condition—warm colors indicate decreased blood flow.

4.2.3 Integrated fMRI and Behavioral Analysis

Finally, we explored whether the change in MTL activation between the high-DTI and low-DTI conditions predicted participants' change regarding security message adherence. We specified a regression model with participants' change in security message adherence as the dependent variable and participants' change in MTL between the two conditions as the independent variable. As there was only one adherence score and change in MTL value for each participant, this is a between-subject comparison (non-repeated data). The results show that the change in MTL activation significantly predicts security message adherence: $\beta = -.470$, $t(23) = 2.495$, $p < .01$, $R^2 = .212$ (see Figure 7). As is evident, there is an outlier. Removing this outlier improves the model fit ($R^2 = .35$ vs. .33), indicating that the model's significance is not driven by this data point.

4.2.4 Test for Habituation

The fMRI method calls for repeated trials per treatment to ensure that brain activation levels are measured reliably. To explore whether participants' behavior changed across trials as a result of habituation, we included the trial repetition number as a fixed effect (i.e., the *n*th warning within a given condition). The results showed that the effect of the trial repetition number was insignificant ($z = -0.251$, $p > .05$), indicating that habituation likely did not influence adherence in this experiment.

4.3 Experiment 1—Discussion

Both the fMRI and behavioral analysis supported our hypotheses. First, we found that participants in the high-DTI condition exhibited less activation in the bilateral MTL than participants in the warning-only condition did (H1 supported). This suggests that DTI inhibits one's ability to utilize the MTL to retrieve information from the long-term memory necessary to accept or reject the permission warnings. We found that people had more than 15% lower security message adherence in the high-DTI condition than in the warning-only condition. We found that the change in MTL predicted participants' change in terms of security message adherence.

We also found that displaying the warning between the working memory tasks (i.e., not during the rehearsal period) improved performance (H2 supported). In the low-DTI condition, participants had more activation in the MTL than they did in the high-DTI condition. Likewise, in the high-DTI condition, participants had an

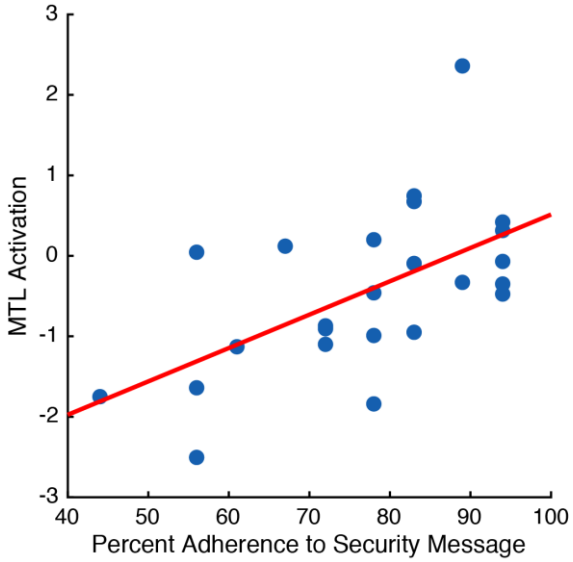

Figure 7. Correlation of MTL activation and adherence to the security message. Each dot represents a participant. The y-axis shows the change in MTL activation for each participant between the low- and high-DTI conditions (low-DTI - high-DTI). The x-axis shows adherence for each participant.

approximately 14% lower security message adherence than those in the low-DTI condition. The change in MTL in this comparison also predicted participants' change in security message adherence. Finally, we found that there was no statistical difference between the warning-only condition and the low-DTI condition.

The neural results from our fMRI study help us to understand why, neurobiologically, people fail to adhere to interruptive security messages. This methodology thus provides an advantage over behavioral data gathered solely through observation and allows us to determine optimal timing for security messages from a neurobiological perspective.

5 Experiment 2—Chrome Cleanup Tool

The main finding from Experiment 1 was that the timing of the interruption mattered, with the high-DTI condition having negative effects both neurally and behaviorally compared to the low-DTI condition. The purpose of Experiment 2 was not to replicate Experiment 1, but rather to apply the finding of Experiment 1 to a more realistic scenario, that is, that displaying security messages at a low-DTI time results in higher security message adherence than at a high-DTI time.

For our experimental context we selected the Google CCT message used in Google Chrome for Windows version 45 (see Figure 1). Google Chrome accounts for more than 56% of the global desktop browser market share [38], and so the CCT potentially impacts millions of users. The CCT detects whether malware has tampered with the host computer, such as by manipulating the browser or other Internet settings [16]. If a problem is detected, the CCT displays a message to the user asking for permission to remove the unwanted software and restore Chrome's original settings. Although the CCT message is important, it is not tied to the users' workflow and, therefore, can be delayed.

Currently, the CCT is displayed at an indeterminate time, appearing sometime after Chrome's initial launch once the CCT code has been downloaded and a scan of the host system has been completed. Because of this variability, the CCT is frequently displayed while the user is engaged in another task and is too often ignored [16].

To improve adherence to the CCT, we collaborated with a team of Google Chrome security engineers responsible for the CCT to identify three low-DTI times at which to display security messages during the browsing experience. The selection of these times was informed by (1) DTI theory and the fMRI results from Experiment 1, (2) input from the Chrome security team on moments that are frequent and generalizable across a wide variety of web-based activities, and (3) a feasibility assessment for implementation in Chrome.

Our first proposed low-DTI time is when a user finishes viewing a web-based video. The second is when users switch web domains (including when a domain is first loaded when a new browser window opens). The third is when users are waiting for a page to load (the average page load is between 4 and 20 seconds depending on location and industry [6]).

Additionally, for comparison, we chose three high-DTI times: while a person is watching an online video, while a person is typing, and while a person is moving to close the browser. These high-DTI times were selected to be in the middle of other tasks, which, based on DTI theory and the results of Experiment 1, should exhibit lower security message adherence. This resulted in a between-subject design with six conditions (three low-DTI times and three high-DTI times).

5.1 Hypotheses

We tested three low-DTI timings to compare against our three chosen high-DTI timings. For each of these times, we propose that users experience a between-task time that will exhibit low DTI and thereby high security message adherence.

H3a–c: Displaying a security message (a) when a web-based video ends, (b) when switching domains, or (c) while waiting for a page to load will result in higher security message adherence than displaying a message during high-DTI times will.

5.2 Participants

We recruited 572 participants from Amazon's MTurk, resulting in 94–97 participants per condition. Per Steelman, Hammer and Limayem [39], all participants were required to be from the United States. The average age of participants was 35.38 years old; 52% were male. All participants were required to participate in the experiment using Google Chrome on Windows, the only version of Chrome that uses the CCT. Participants were paid $1 USD for approximately a 6-minute task.

5.3 Ethics

The university IRB approved the protocols used. We complied with all requester policies for MTurk. In each experiment, participants were debriefed afterward and told the purpose of the experiment. Working with Google Chrome's security team, we received permission to exactly replicate the CCT message in our experiment. To increase realism, we did not tell participants upfront that the true purpose of the task was to examine how they respond to the CCT. Per Google's guidelines, at the end of the task, we disclosed that the CCT message they saw

was simulated and not capable of detecting malware. We then provided a link that educated participants about the real CCT message and instructed them that they should comply with the message if they saw it outside the experimental setting

5.4 CCT Message Development

The CCT prompt used in our study is identical to the real CCT prompt shown in Figure 1 except that ours was positioned slightly lower on the screen relative to the address bar. In a follow-up survey, we asked participants if they noticed anything odd about the prompt. None of the participants mentioned the small positioning difference.

5.5 Procedure

The experimental task was designed to realistically mimic a typical MTurk task. Participants were instructed that their task was to help create an archive of videos on the web. Participants were given a webpage URL to watch a 30-second commercial video. After watching the video, participants were given another URL that led them to a video archive website. On this website, they were asked to enter the URL for the video they had watched and to summarize the video in at least 25 words (the system enforced the word count). After submitting their summary, the webpage displayed the following message for 10 seconds: "Please wait while we fetch and process the video." After processing was completed, participants were given a confirmation code that they were required to enter on another webpage to receive payment.

The two websites were designed by the research team specifically for this study and included JavaScript that could trigger the CCT message. Consistent with the way the CCT is displayed in Chrome, the CCT message remained visible over the content of the webpage until the user either accepted or dismissed it, or until the webpage was closed.

The experimental system randomly assigned each participant to one condition before visiting the first webpage, resulting in a between-subject design with six different conditions. The three low-DTI conditions were displayed as follows. First, in the LDTI-1 ("After video") condition, participants saw the prompt after finishing the video. In the LDTI-2 ("Switching web domains") condition, participants saw the prompt when switching domains to the second website. For the LDTI-3 ("Waiting for page load") condition, we implemented an artificial loading delay of approximately six seconds as the second page loaded. A few seconds into the loading delay, the CCT message was shown.

The three high-DTI conditions were similar to the low-DTI conditions except that the CCT message was displayed during, rather than before or after, a task. For the HDTI-1 ("During video") condition, participants saw the CCT message 10 seconds into the 30-second video. In the HDTI-2 ("While typing") condition, participants saw the prompt while typing the description of the video (when they typed word 10 out of 25). Finally, in the HDTI-3 ("Mousing to close window") condition, participants were shown the prompt as they were in the middle of moving to close the window. We determined when people were moving to close the window based on the trajectory and velocity of movement toward the upper right-hand corner of the screen at the end of the task.

The system recorded whether participants clicked on the "Run Chrome Cleanup Tool" button, the "Dismiss" button, or ignored the message.

Participants completed a post-task survey to gather demographic information and manipulation check items. Further, the post-task survey disclosed the real purpose of the experiment—to explore how people respond to the CCT message.

5.6 Results

5.6.1 Manipulation Checks

Prior to analyzing the security message adherence for the different conditions, we conducted a manipulation check to again verify that our hypothesized times had lower DTI than the high-DTI times. In a post-task survey, we asked all participants the following question as a manipulation check for DTI: "When the above message appeared, I was busy doing other things" (with the CCT message shown above the question). Participants responded on a 7-point Likert agree–disagree scale. An ANOVA indicated that a difference existed among the different conditions: $F(5, 530) = 12.899$, $p < .001$. Using a Tukey post hoc comparison analysis, we found that people reported significantly less DTI for each of the low-DTI conditions than for the high-DTI conditions.

5.6.2 Main Analysis

To test our hypotheses, we compared each condition to every other condition, using logistic regression contrasts.

We conceptualized our dependent variable—security message disregard—as whether participants ignored the message or responded to it (a binary variable). We chose this instead of whether participants clicked on the "Run Chrome Cleanup Tool" or "Dismiss" buttons because feedback from the pilot test indicated that some people clicked "Dismiss" if they thought the prompt itself was malware (which is not an example of disregarding the message, but rather a thoughtful response). However, as

the CCT prompt does not automatically disappear when ignored, responses from the pilot test suggest that not responding to the warning was a result of not noticing or giving attention to the warning. Thus, we deem an appropriate conceptualization of security message disregard as whether participants ignored the message.

Further, because the prompt did not disappear unless the participant clicked on the "Dismiss" or "Run Chrome Cleanup Tool" button or until the end of the experiment, some messages were shown longer than others were (and because some messages were shown earlier in the experiment, they had the potential to be shown for a longer duration than others). To account for this, we also recorded how long the message was displayed to control for this duration in the analysis.

We then modeled contrasts for each condition group. Each condition was included as a dummy variable, coded as 1 if the timing was implemented and 0 if the timing was not implemented (each participant received only one timing; therefore, only one of the dummy variables was coded as 1, and all the others were coded as 0). Finally, to compare each condition to every other condition, we treated each condition separately as the reference/baseline class in this model (the condition to which every other condition is compared). Table 3 summarizes the results. In this table, difference in percentage is shown (percentage on the x minus percentage on the y), followed by the significance level of whether there was a difference between the intersecting groups based on the analysis. The Nagelkerke R^2 for the model was .247. Table 4 summarizes the conditions in descending order by security message adherence.

	LDTI-1				
LDTI-2	-.03 ns	LDTI-2			
LDTI-3	.22**	.24**	LDTI-3		
HDTI-1	-.36***	-.33***	-.57***	HDTI-1	
HDTI-2	-.34***	-.32***	-.56***	.01 ns	HDTI-2
HDTI-3	-.31***	-.28***	-.52***	.05 ns	.03 ns

Values equal the difference in the percentage of adherence between the group on the treatment on the x-axis and the treatment on the y-axis. The asterisk indicates whether this difference was significant based on the logistic regression analysis ($p < .05$, ** $p < .01$, *** $p < .001$, ns = non-significant).

Table 3. Change in percentage (x − y) and significance among low-DTI and high-DTI conditions.

5.6.3 Qualitative Support

In a post survey, we allowed participants to comment on the experiment. Several comments provided qualitative support that high DTI decreased adherence. For example, one participant stated, "I remember seeing it [the security message] and thinking that I was too busy to be concerned about it and ignored it." Another participant explained, "You were so preoccupied with completing the HIT that you didn't even notice the warning (which is my case)." A third participant stated, "I was so busy doing the task I did not notice the warning. I would normally have noticed right away, I think." These comments and others suggest that high DTI negatively impacts security message adherence.

6 Experiment 2 Discussion

We found that displaying the message during each of the low-DTI conditions resulted in significantly higher adherence than in all of the high-DTI conditions. Thus, hypotheses H3a–c were supported. This was consistent both with DTI theory and the main finding from Experiment 1 that the timing of the interruption mattered.

Among the low-DTI times, we found that conditions LDTI-1–2 ("after video," "switching web domains") were not significantly different from each other. However, they had significantly lower security message adherence than LDTI-2 ("during page load"). In summary, our experiment provides support that intelligent timing of security messages can mitigate the DTI effect.

7 General Discussion

The objectives of this study were twofold. First, we sought to explain how DTI occurs in the brain in response to interruptive security messages, and how this in turn affects adherence. We designed an fMRI experiment (Experiment 1) that examined how activation in the MTL changes when users are interrupted by security messages in the middle of a task or between a competing primary tasks. We found that activation in the MTL is lower when users are interrupted in the middle of a primary task (high DTI) compared to when they are between primary tasks (low DTI). This suggests that high-DTI conditions decrease users' ability to utilize the MTL for processing security messages. We also showed that this decrease in MTL activation was associated with decreased security message adherence (see Table 2).

Our second objective was to determine whether finessing the timing of security messages can reduce DTI and improve adherence. We found in the fMRI experiment that participants who received the security message between primary tasks experienced activation in the MTL similar to when they attended to the security message as the only task. This in turn resulted in greater security message adherence. Further, we showed that the increase in MTL activation between the low- and high-DTI conditions directly predicted participants' increased adherence (see Figure 6).

Finally, we applied our findings from the fMRI experiment to the context of the CCT message in Google Chrome (Experiment 2). We found that participants had higher adherence when the CCT was displayed at low-DTI times compared to high-DTI times (see Tables 3 and 4).

7.1 Contributions

We empirically explore how DTI affects security message adherence. In past research, DTI scholars have largely focused on how interruptions decrease performance of primary tasks (e.g., work productivity, accuracy). However, in a security context, it is vital to understand how DTI influences performance on the security message—the interruption itself.

In Experiment 1, we used fMRI to show that DTI suppresses activity in the MTL region of the brain, which decreases one's ability to retrieve the necessary information from declarative memory to respond properly to the security message. For example, when the security message was presented at a high-DTI time, participants displayed an approximately 15% decrease in adherence compared to a low-DTI time. In contrast, when presenting security messages at a low-DTI time, participants exhibited essentially the same level of adherence as they did if their exclusive task was to respond to the security message. These results provide a sound theoretical foundation for designing security messages that avoid DTI and improve adherence.

Code	Condition	Adherence
LDTI-3	Waiting for Page Load	77.89%
LDTI-1	After Video	56.25%
LDTI-2	Switching Domains	53.68%
	LDTI average	**62.61%**
HDTI-3	On the Way to Close Window	25.53%
HDTI-2	While Typing	22.11%
HDTI-1	During Video	20.62%
	HDTI average	**22.75%**

Table 4. Percentage adherence to the CCT security message for low- and high-DTI conditions.

Although fMRI has limitations, it provides insights that are not possible using behavioral experiments alone. Namely, the fMRI results provide strong neural evidence that DTI does inhibit adherence to security messages, and explain why adherence is better at low-DTI versus high-DTI times. This contribution is unique in the security warning and interruption literatures. Thus, Experiment 1 provides a neural explanation for why low-DTI timings are effective, and Experiment 2 tests this premise in a naturalistic setting.

In Experiment 2, we applied our fMRI findings to the context of a web browser, specifically the CCT message in Google Chrome for Windows version 45, which is an important example of a security message that can be delayed. Currently, this message is displayed without regard to the activity of the user and has a less-than-desirable level of adherence. To improve adherence to the CCT, we collaborated with members of the Google Chrome security team to identify three low-DTI times at which to display the CCT. We found in our online experiment that, on average, adherence to the CCT was 36% higher at low-DTI times compared to high-DTI times. Further, these examples point the way to finding similar low-DTI times in the browsing experience and other contexts.

Together, our results in Experiments 1 and 2 suggest that for those security messages that can be delayed, there is a considerable benefit in displaying a security message at a low-DTI time. For those security messages that cannot be delayed, our results warn of the substantial negative impact that DTI can have on security message adherence. Practitioners and researchers should therefore explore ways to present interrupting security messages in such a way as to make them more resistant to DTI. For example, because our results show that activation in the MTL brain regions is lower when a message interrupts a primary task, security messages should be designed to minimize users' reliance on declarative or long-term memory.

7.2 Limitations

This research is subject to certain limitations. First, the MRI method imposes constraints that may hinder the realism of the task. For example, participants must lie supine and still in a narrow tube for the duration of the experiment. However, these constraints are a necessary trade-off in order to obtain insights regarding neural correlates of DTI that are unavailable using other methods. However, Experiment 2 at least partially compensated for this limitation in that it used a realistic

task to provide greater ecological validity. Future research may apply field methodologies to further improve ecological validity.

Second, we did not measure neural activity in Experiment 2. This was by design so that we could apply our findings from Experiment 1 in a more realistic context, for which fMRI is not well-suited. Because the nature of the tasks in Experiments 1 and 2 were different, the exact neural systems involved could also differ. For example, because Experiment 1 used a memory-based working task, the MTL was strongly implicated, which is involved in memory encoding and retrieval. It is possible that other brain regions would be activated for the task of Experiment 2. However, our main finding from Experiment 1 was that the timing of the interruption mattered, with the high-DTI condition having negative effects both neurally and behaviorally compared to the low-DTI condition. The behavioral results from Experiment 2 show this same pattern, consistent with DTI theory. Given the matched behavior between the two tasks, we would expect the overall timing-dependent pattern of activation to remain constant were we to have participants perform the second task in the MRI scanner.

Third, our design for Experiment 1 required participants to become very familiar with the list of risky permission warnings. We believe that we may have trained them more thoroughly than is typical of corporate training regarding security messages. We intentionally did not want the participants to use their own security judgment because doing so might bias the data. Similarly, the experiment exposed participants to 54 security warnings over the course of Experiment 1. This number was necessary to ensure a sufficient signal-to-noise ratio for the fMRI analysis. Most users will not encounter anything close to that many security messages of the same type within an hour. However, in Experiment 2, each participant responded to only one security message. This was another way that Experiment 2 enhanced the overall ecological validity of our results.

Fourth, Experiment 2 displayed the CCT prompt randomly at low- and high-DTI times. However, the tool did not actually check for malware on the users' computers. Thus, some participants may have ignored it if they were certain that their computer was free of malware. Future research should therefore cross validate our results in scenarios when the prompt is shown in response to actual detected malware.

Finally, we examined only two types of security messages. The range of security messages is broad and has great variability. We leave it to future research to investigate how other security messages affect user behavior.

8 Conclusion

Our key message is that although security messages are ubiquitous in personal computing, they should be bounded in their presentation. We demonstrate using fMRI that the timing of a security message's interruption substantially affects the occurrence of DTI in the brain as well as security message adherence. Further, using the context of the CCT message in a realistic online experiment, we demonstrate a practical way to mitigate the DTI effect by finessing the timing of when the security message is displayed. Our findings provide to practice an easy-to-implement, cost-effective approach to increase security message adherence that can be applied to a wide variety of security messages that are not tied to a specific event or user action and can therefore be delayed. For those security messages that cannot be delayed, our results warn of the substantial negative impact that DTI can have on adherence.

9 Appendix: fMRI technical details

9.1 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 = 3.0 mm (no skip), matrix size = 64 × 64, and voxel size = 3.44 mm × 3.44 mm × 3 mm.

9.2 Protocol

Participants were given a verbal briefing about the MRI procedures and the task, and they were then situated supine in the scanner. Participants viewed visual stimuli 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 the images appeared normal when viewed through the mirror. We first performed a 10-second localizer scan, followed by a 7-minute structural scan. Following these scans, we started the experimental task. Total time in the scanner was about 50 minutes. To avoid issues of autocorrelation and achieve reasonable efficiency we employed a jittered (+/- 3 sec) interstimulus interval.

9.3 Analysis

We analyzed the MRI data with the Analysis of Functional Images (AFNI) software suite [7]. Briefly, functional data was slice-time corrected to account for differences in acquisition time for different slices of each volume.

To control for head motion, we registered each volume with the middle volume of each run. We aligned data from each run to the run nearest in time to the acquisition of the structural scan. We then co-registered the structural scan to the functional scans. Spatial normalization was accomplished by first warping the structural scan to the Talairach Atlas [42], followed by warping to a template brain with Advanced Neuroimaging Tools [ANTs; 2].

Single subject regression analyses were carried out by creating regressors for each event type: memory code display, high-DTI rehearsal, low-DTI rehearsal, memory retrieval, low-DTI warning, and warning-only warning. Periods without explicit task demands were included in the model as an implicit baseline. We modeled stimulus durations as described above. We blurred statistical parameter maps (beta values) from the single-subject regression analysis using an 8-mm FWHM Gaussian kernel. We then entered beta values for the conditions of interest into group-level analyses, such as *t*-tests, which were used to determine significant clusters of activation.

To correct for Type I errors (i.e., control for multiple comparisons), all reported clusters of activation were thresholded at a voxel-wise p-value < .02 and a spatial extent threshold of $k > 40$ contiguous voxels (1080 mm^3), yielding an overall p-value < .05, as determined through Monte Carlo simulations.

As the system randomized the order of presentation, we paired the MRI data with the behavioral data and interaction with the web-based program using a signal generated by the scanner at the beginning of each task.

10 REFERENCES

[1] Adamczyk, P.D. and Bailey, B.P. If not now, when? The effects of interruption at different moments within task execution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2004).

[2] Advants. ANTs version 1.9. (2011). http://sourceforge.net/projects/advants/

[3] Anderson, B., Kirwan, B., Eargle, D., Howard, S. and Vance, A. How polymorphic warnings reduce habituation in the brain: Insights from an fMRI study. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)* (2015).

[4] Borst, J.P., Taatgen, N.A. and van Rijn, H. What makes interruptions disruptive? A process-model account of the effects of the problem state bottleneck on task interruption and resumption. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2015).

[5] 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* (2011).

[6] ByReputation. Average web page load times by industry. http://www.byreputation.com/Average-Web-Page-Load-Times_a/452.htm

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

[8] Cutrell, E., Czerwinski, M. and Horvitz, E. Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. In *Human-Computer Interaction* (2001).

[9] Dhamija, R., Tygar, J.D. and Hearst, M. Why phishing works. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2006).

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

[11] Egelman, S. and Schechter, S. The importance of being earnest [in security warnings]. In *Financial cryptography and data security*, A.-R. SADEGHI Ed. Springer Berlin Heidelberg, 52-59, 2013.

[12] Felt, A.P., Ainslie, A., Reeder, R.W., Consolvo, S., Thyagaraja, S., Bettes, A., Harris, H. and Grimes, J. Improving SSL warnings: Comprehension and adherence. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2015).

[13] Felt, A.P., Ha, E., Egelman, S., Haney, A., Chin, E. and Wagner, D. 2012. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security* ACM, 3:1-3:14.

[14] Fogarty, J., Hudson, S.E., Atkeson, C.G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J.C. and Yang, J. Predicting human interruptibility with sensors. *ACM Transactions on Computer-Human Interaction (TOCHI) 12*, 1 (2005), 119-146.

[15] Friedman, H. and Goldman-Rakic, P. Activation of the hippocampus and dentate gyrus by working-memory: A 2- deoxyglucose study of behaving rhesus monkeys. *The Journal of Neuroscience 8*, 12 (1988), 4693-4706.

[16] Google 2015. Elisabeth morant, product manager for chrome usable security, anti-malware, and assist features.

[17] Halbeisen, G. and Walther, E. 2015. Dual-task interference in evaluative conditioning: Similarity matters! In *The Quarterly Journal of Experimental Psychology*, 1-33.

[18] Horvitz, E., Breese, J., Heckerman, D., Hovel, D. and Rommelse, K. The lumière project: Bayesian user modeling for inferring the goals and needs of software users. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence* (1998).

[19] Iqbal, S.T., Adamczyk, P.D., Zheng, X.S. and Bailey, B.P. Towards an index of opportunity: Understanding changes in mental workload during task execution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2005).

[20] Iqbal, S.T. and Bailey, B.P. Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (2005).

[21] Iqbal, S.T. and Horvitz, E. Disruption and recovery of computing tasks: Field study, analysis, and directions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2007).

[22] Jett, Q.R. and George, J.M. Work interrupted: A closer look at the role of interruptions in organizational life. *Academy of Management Review 28*, 3 (2003), 494-507.

[23] Logan, G.D. Attention in character-classification tasks: Evidence for the automaticity of component stages. *Journal of Experimental Psychology 107*, 1 (1978), 32-63.

[24] Mark, G., Gudith, D. and Klocke, U. The cost of interrupted work: More speed and stress. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2008).

[25] Mark, G., Voida, S. and Cardello, A. A pace not dictated by electrons: An empirical study of work without email. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2012).

[26] Mark, G., Wang, Y. and Niiya, M. Stress and multitasking in everyday college life: An empirical study of online activity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2014).

[27] McFarlane, D. Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction 17*, 1 (2002), 63-139.

[28] Navon, D. and Miller, J. Role of outcome conflict in dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance 13*, 3 (1987), 435.

[29] Neupane, A., Rahman, M.L., Saxena, N. and Hirshfield, L. 2015. A multi-modal neuro-physiological study of phishing detection and malware warnings. In *ACM Conference on Computer and Communications Security (CCS)*, Denver, CO.

[30] Neupane, A., Saxena, N., Kuruvilla, K., Georgescu, M. and Kana, R. Neural signatures of user-centered security: An fMRI study of phishing, and malware warnings. In *The Network and Distributed System Security Symposium (NDSS)* (2014).

[31] NPR. Interruption science: Costly distractions at work. (2005). http://www.npr.org/templates/story/story.php?storyId=4958831

[32] Oliver, N., Horvitz, E. and Garg, A. 2002. Layered representations for recognizing office activity. In *Proceedings of the International Conference on Multimodal Interaction (ICMI 2002)*, Pittsburgh, PA, 3-8.

[33] Pashler, H. Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin 116*, 2 (1994), 220-244.

[34] Plummer, P., Apple, S., Dowd, C. and Keith, E. Texting and walking: Effect of environmental setting and task prioritization on dual-task interference in healthy young adults. *Gait & Posture 41*, 1 (2015), 46-51.

[35] Rémy, F., Wenderoth, N., Lipkens, K. and Swinnen, S.P. Dual-task interference during initial learning of a new motor task results from competition for the same brain areas. *Neuropsychologia 48*, 9 (2010), 2517-2527.

[36] Schon, K., Newmark, R.E., Ross, R.S. and Stern, C.E. A working memory buffer in parahippocampal regions: Evidence from a load effect during the delay period. *Cerebral Cortex*, Forthcoming (2015).

[37] Speier, C., Vessey, I. and Valacich, J.S. The effects of interruptions, task complexity, and information presentation on computer‐supported decision‐making performance. *Decision Sciences 34*, 4 (2003), 771-797.

[38] StatCounter. Statcounter global stats. (2015). http://gs.statcounter.com/#desktop-browser-ww-monthly-201408-201508

[39] Steelman, Z.R., Hammer, B.I. and Limayem, M. Data collection in the digital age: Innovative alternatives to student samples. *MIS Quarterly 38*, 2 (2014), 355-378.

[40] Strayer, D.L. and Johnston, W.A. Driven to distraction: Dual-task studies of simulated driving and conversing on a cellular telephone. *Psychological Science 12*, 6 (2001), 462-466.

[41] Szameitat, A.J., Schubert, T. and Muller, H.J. How to test for dual-task-specific effects in brain imaging

studies: An evaluation of potential analysis methods. *NeuroImage 54*, 3 (2011), 1765-1773.

[42] TALAIRACH, J. AND TOURNOUX, P. *Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system: An approach to cerebral imaging*. Thieme, Stuttgart, 1988.

[43] TANI, T. AND YAMADA, S. Estimating user interruptibility by measuring table-top pressure. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems* (2013).

[44] TOMBU, M. AND JOLICŒUR, P. A central capacity sharing model of dual-task performance. *Journal of Experimental Psychology: Human Perception and Performance 29*, 1 (2003), 3-18.

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

[46] WEST, R. The psychology of security. *Communications of the ACM 51*, 4 (2008), 34-40.

[47] YEE, K.-P. Aligning security and usability. *Security & Privacy, IEEE 2*, 5 (2004), 48-55.

[48] ZÜGER, M. AND FRITZ, T. Interruptibility of software developers and its prediction using psycho-physiological sensors. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015).