

Harder to Ignore?

Revisiting Pop-Up Fatigue and Approaches to Prevent It

Cristian Bravo-Lillo
cbravo@cmu.edu

Lorrie Cranor
lorrie@cs.cmu.edu

Saranga Komanduri
sarangak@cmu.edu

Stuart Schechter
stus@microsoft.com

Manya Sleeper
msleeper@cmu.edu

ABSTRACT

At SOUPS 2013, Bravo-Lillo et al. presented an artificial experiment in which they habituated participants to the contents of a pop-up dialog by asking them to respond to it repeatedly, and then measured participants' ability to notice when a text field within the dialog changed. The experimental treatments included various *attractors*: interface elements designed to draw or force users' attention to a text field within the dialog. In all treatments, researchers exposed participants to a large number of repetitions of the dialog before introducing the change that participants were supposed to notice. As a result, Bravo-Lillo et al. could not measure how habituation affects attention, or measure the ability of attractors to counter these effects; they could only compare the performance of attractors under high levels of habituation. We replicate and improve upon Bravo-Lillo et al.'s experiment, adding the low-habituation conditions essential to measure reductions in attention that result from increasing habituation. In the absence of attractors, increasing habituation caused a three-fold decrease in the proportion of participants who responded to the change in the dialog. As with the prior study, a greater proportion of participants responded to the change in the dialog in treatments using attractors that delayed participants' ability to dismiss the dialog. We found that, like the control, increasing habituation reduced the proportion of participants who noticed the change with some attractors. However, for the two attractors that forced the user to interact with the text field containing the change, increasing the level of habituation did not decrease the proportion of participants who responded to the change. These attractors appeared resilient to habituation.

1. INTRODUCTION

Operating systems, browsers, and other software frequently interrupt user workflow with often-irrelevant security warning dialogs. This abundance has been mentioned repeatedly as a problem in usable security research [2, 4, 5, 6, 7, 10, 11]; most authors seem to agree that dialogs are overused, and that when reaching a dialog a high proportion of users will dismiss it because they are already fatigued. Computer users have also self-reported habituation to dialogs. Krol et al. conducted a lab study wherein participants were

exposed to two similar pop-up dialogs [8]. 81% of participants clicked through the dialogs; 45% of participants freely mentioned desensitization as a reason for ignoring the dialogs.

Habituation is a simple form of learning in which "repeated or prolonged exposure to a stimulus results in gradual reduction in responding" [9]. After an extensive review, Thompson and Spencer found nine distinctive features of habituation [12]. For example, a) the decrease in response is usually exponential on the number of exposures, b) if the stimulus is taken away, the original response usually reappears in time, c) if repeated series of habituation training and spontaneous recovery are given to a person, habituation becomes progressively faster, d) the weaker the stimulus, the faster and/or stronger habituation becomes (strong stimuli usually show no significant habituation effects), and e) habituation to a given stimulus has been shown to generalize to other stimuli.

In 2013, Akhawe and Felt conducted a large study on telemetry data collected from SSL, malware, and phishing warnings in Chrome and Firefox [1]. In this study, Chrome users were more than twice as likely to ignore SSL warnings as Firefox users. Unlike Firefox, Chrome does not have an exception storing mechanism for certificate errors, and the authors suggest this as one possible reason for the disparity. Chrome users see a warning on each interaction with a self-signed certificate, which could result in many false positives and produce habituation. The authors call this "warning fatigue" and provide timing data that is consistent with this hypothesis [1].

A number of studies have found that browser dialogs resembling those dialogs that participants encounter frequently are more likely to be ignored by participants in laboratory experiments than less-familiar designs [6, 10, 11]. In addition, prior studies have found evidence of habituation beginning to occur after just one or two exposures to a new dialog [2, 4, 11]. However, these studies were not specifically designed to measure the impact of habituation. They did not completely control for other factors that might have been responsible for users' behavior and did not measure the impact of varying levels of habituation.

Bravo-Lillo et al. presented three experiments designed to measure the impact of user-interface modifications created to direct users' attention in security dialogs [3]. In the first two experiments, these interface elements, termed *attractors*, were used in a software installation dialog. The researchers used attractors to direct participants' attention to a *salient field*, a text field that contained information that would allow users to differentiate between harmless and malicious scenarios. For example, the *Swipe* attractor required participants to swipe their mouse over the salient field, which contained the publisher name, to activate the option that presented more risk (i.e., installing software).

In the 2013 study, participants in treatments that used attractors

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2014, July 9–11, 2014, Menlo Park, CA.

were more likely to be able to differentiate between suspicious and harmless scenarios, and thus appeared more likely to be paying attention to the salient text field. The study’s first two experiments used a deceptive ruse of an online game evaluation study to present dialogs to users as if they represented real security decisions with real consequences, so as to maximize ecological validity. However, since participants would not have seen attractors in real-world security dialogs before, it was possible that improvements seen were the result of the novelty of attractors, and that the benefit would wear off over time with the novelty.

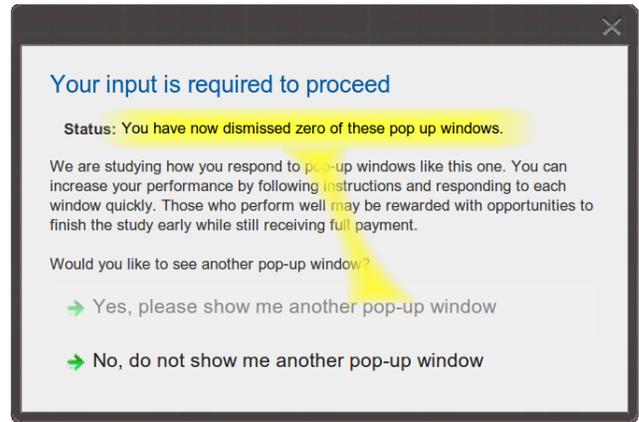
To test whether attractors were effective after habituation, the researchers also included a third experiment. Unlike the first two experiments, the researchers did not hide from participants the fact that they were studying interaction with dialogs. Rather, the researchers asked participants to respond to as many dialogs as possible in a five-minute period. The dialog asked participants if they would like to see another dialog, and the only working option was *Yes* (the *habituated* option) with the other option, *No*, having no effect. In this third experiment the researchers used attractors to direct users’ attention to a ‘Status’ field that contained information intended to be irrelevant to their choice of action—the number of dialogs they had dismissed so far. Thus, they attempted to train participants that attractors were only used to direct them to information that was of no use, habituating them to ignore the ‘Status’ field that the attractors directed them to.

In this third experiment, after a habituation period of either 2.5 minutes or 22 exposures, the researchers placed into the ‘Status’ field a *habit-breaking* instruction that was relevant to the participants’ actions. The researchers’ new instruction directed participants to choose the now-active *No* option to finish the study early. The researchers measured the performance of each attractor as the proportion of participants who chose this *alternate* option on their first opportunity. They found that, under these habituated conditions, five attractors (*Swipe*, *Type*, *Animated Connector + Swipe*, *Animated Connector + Reveal*, and *Animated Connector + Delay*) performed better than the control [3].

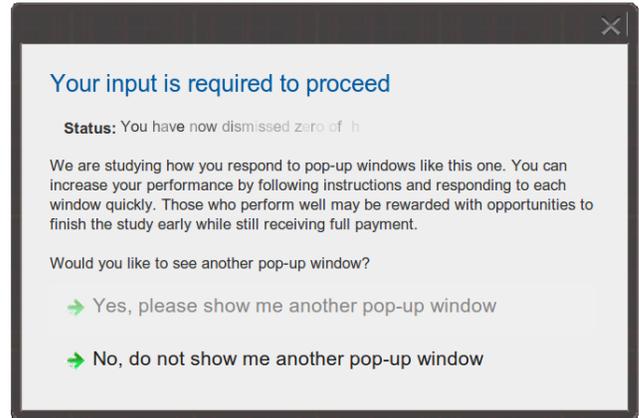
However there was a **serious limitation in this study design**. In Bravo-Lillo et al.’s prior work the researchers included only one level of habituation (changing the dialog box after 22 exposures or 2.5 minutes depending on the type of attractor). Because the researchers did not also include lower habituation conditions, it was not possible to determine whether different attractors were more *resistant* to habituation than the control. Rather, it is possible that some attractors resulted in such large initial increases in attention that they continued to perform well *despite* habituation.

In this paper we expand on Bravo-Lillo et al.’s prior work by replicating and extending the prior experiment for several of their attractors. We overcome the limitations of the previous study by including both *low-* and *high-habituation* conditions. This important extension to the study design allows us to measure attractors’ *resistance* to habituation and yields three new findings:

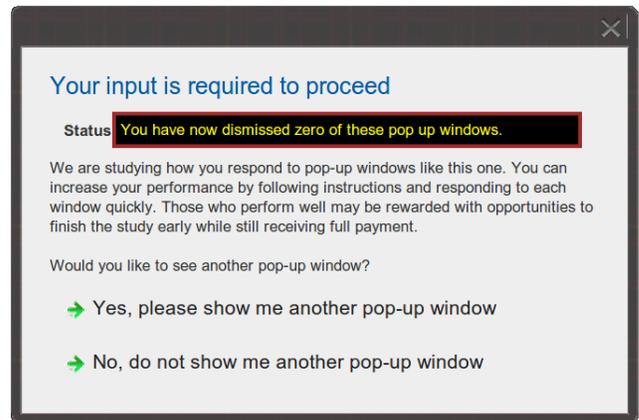
1. **Habituation reduces attention for the baseline experimental task (the control case).** For the control dialog, increasing habituation decreased the proportion of participants who would choose the alternate option (*No*) at the first opportunity by more than a factor of three. This confirms hypotheses about habituation from prior work.
2. **Some attractors failed to show resistance to habituation.** Increasing habituation negatively impacted the performance of some attractors, even though these attractors still outperformed the control in high-habituation conditions. In particular, two attractors that displayed an animation before activat-



(a) Animated Connector + Delay attractor



(b) Reveal attractor



(c) ANSI attractor

Figure 1: Dialogs that are designed to visually draw users’ attention to the salient field.

ing the habituated option (*Yes*) became less effective after habituation. This finding brings new insights to the prior finding from Bravo-Lillo et al. [3] and suggests that the forces of both novelty and habituation might explain the results of the prior study.

3. **Some attractors showed resistance to habituation.** The length of the habituation period had no measurable impact

on the performance of two attractors. The first of these two habituation-resistant attractors forced users to swipe their mouse over the salient field before choosing the *Yes* option; whereas the second, more arduous, attractor required them to retype the field’s contents. This is the first experimental evidence to demonstrate that some user-interface modifications can significantly reduce, if not entirely prevent, habituation from sapping users’ attention to warnings.

As in the prior experiment [3], we also observed that the usability cost of the swipe attractor seems to decrease with time. Once participants grew accustomed to it, they could respond to dialogs containing this attractor within three to five seconds.

2. ATTRACTORS

An attractor is an interface modification designed to draw or force attention to an information field called the *salient field*. The salient field is the part of the dialog that provides the most important information to aid the user’s decision.

We implemented five of the attractors presented by Bravo-Lillo et al. [3]. Four are *inhibitive attractors*, which prevent users from making potentially-hazardous choices until after some period of time has passed or a user performs some action. The inhibitive attractors appear only when a user moves the mouse pointer over the button representing the potentially dangerous option. In security dialogs, this *triggering option* is the option that represents a security risk (e.g., installing software). In our study, the triggering option states “*Yes, please show me another pop-up window.*” The attractor is not triggered if the user attempts to select the *No* option.

The *Animated Connector + Delay* (AC + Delay, Figure 1a) attractor is a yellow highlight that first appears behind keywords in the triggering option that relate to the salient field. Over a period of two seconds, the highlighted region progresses in the direction of the salient field, and then fills the background of the field. The attractor disables the *Yes* option for five seconds (hence, the delay).

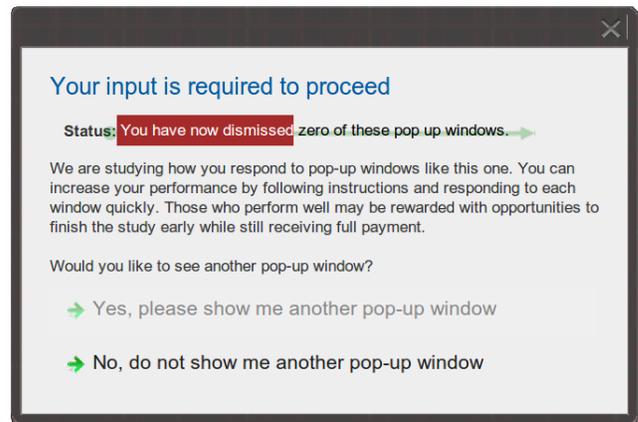
The *Reveal* attractor (Figure 1b) first hides the contents of the salient field, then progressively animates it back in a random fashion, mostly from left to right, over a period of five seconds. The motion and randomization are intended to help users notice each letter as it appears.

The *Swipe* attractor (Figure 2a) disables the *Yes* option until the user moves her mouse from left to right over the salient field. As the mouse moves over each letter, that letter becomes highlighted. If the user moves her mouse over the triggering option before swiping, a pop-up message appears that explains how to swipe and illustrates the swiping motion with an animated cursor.

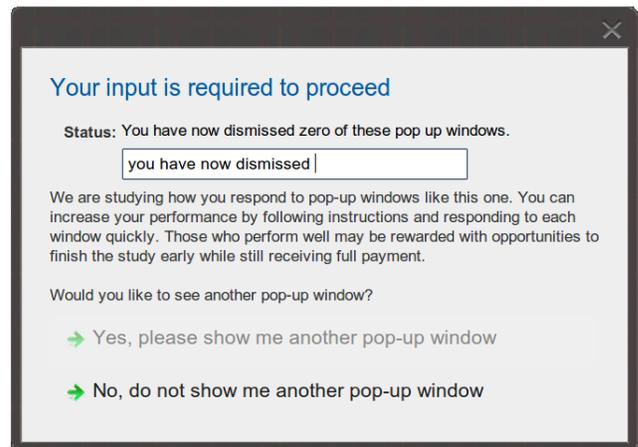
The *Type* attractor (Figure 2b) requires the user to retype the contents of the salient field (no pasting allowed). Bravo-Lillo et al. had included this treatment with the assertion that it would be quite difficult to type text without paying attention to it. We also include this treatment in part to measure effects that may confound our experiment’s ability to measure attention (such as frustration with the tedium of a task).

We also included the non-inhibitive *ANSI* attractor from the prior study (Figure 1c), which gives the salient field a black background and high-contrast yellow text to draw attention to it. This treatment helps measure the impact of novel, attention-grabbing stylistic changes that are not accompanied by time delays or required actions.

In Bravo-Lillo et al.’s prior work they also implemented several combined attractors (*Animated Connector + Swipe* and *Animated Connector + Reveal*) [3]. We tested only the uncombined versions in our study.



(a) Swipe attractor.



(b) Type attractor.

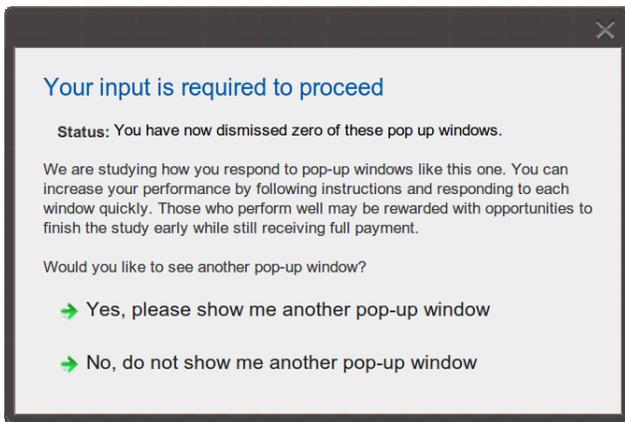
Figure 2: Dialogs that include attractors that require users to interact with the salient field.

3. STUDY DESIGN

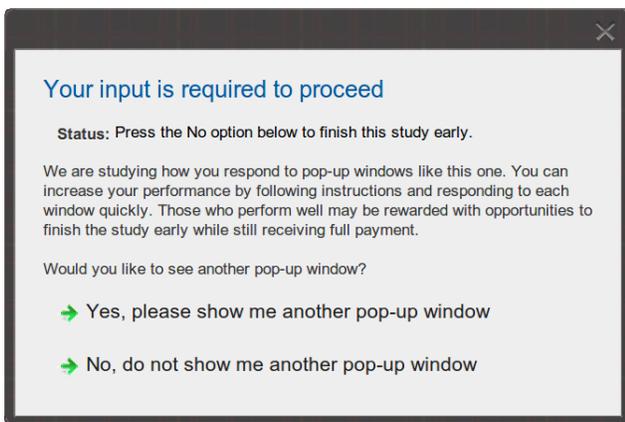
Except where noted, we replicated the experimental methodology documented in Bravo-Lillo et al. [3]. We recruited participants by advertising a human-intelligence task to workers on Amazon’s Mechanical Turk, asking them to perform a task in which they would respond to as many dialogs as possible for a fixed time period. We instructed participants that their task was to respond to questions in pop-up windows as quickly as possible over a ten-minute period. We also instructed them to look for opportunities to finish the study early.

During a *habituation period*, we displayed the dialog shown in Figure 3a. In this dialog, the contents of the information field, labeled “Status”, alternated between the message “You have now dismissed *n* of these pop up windows” and “*n* pop up windows have been dismissed so far,” where *n* was written in words (not digits). This iteration between two irrelevant messages was to ensure that attention to the ‘Status’ field was due to its content and not to its replacement by another message.

During the habituation period we prevented the *No* option from having any effect, though we did not change its appearance to indicate that we had disabled it. By removing one of the two available options we effectively forced participants to choose the *Yes* option to dismiss the dialog so as to habituate them to clicking this option.



(a) Dialog shown during the habituation phase. The Status field displayed alternately between the sentence shown in the dialog, and “N dialogs have been dismissed so far.”



(b) Dialog shown during the test phase. Note the change in the Status field.

Figure 3: Control dialogs used in the experiment.

We displayed each dialog at randomly selected coordinates within a participant’s browser. If we detected 15 seconds of inactivity we warned participants that we would exclude those who were inactive for 30 seconds or more.

The dialog we used in this experiment differs slightly from prior work in that we used the phrase “pop up windows” to describe the dialogs participants were asked to dismiss, while the prior study asked participants to dismiss “questions.” We made this change after piloting to reduce participant confusion. Several participants indicated that they expected actual questions when we asked them about problems encountered during the task.

We followed the habituation period with a *test period* during which we presented the same dialog but with the alternate (*No*) option enabled and the contents of the status field replaced with the instruction, “Press the No option below to finish this study early.” Participants who read and understood the habit-breaking instruction in the status field discovered that they should stop choosing the habituated option (*Yes*) and instead choose the alternate option (*No*). We terminated the test period when the participant chose the alternate *No* option or their ten minutes were up.

We then presented an exit survey in which we asked participants to recall the contents of the status field, instructing those with no recollection to type *None*. We paid \$1.00 to all participants who completed the experiment.

We designed six treatments: one control treatment with no attractors, and five treatments each with one of the five attractors described in the previous section. For every treatment we created a condition for each of four habituation periods, resulting in $24 (6 \times 4)$ total conditions. Each participant was assigned to a single condition.

We defined the duration of three habituation periods in terms of the number of habituation dialogs the participant would be exposed to (1, 3, and 20 exposures). These habituation periods lasted for as much time as it took participants to dismiss the dialogs they were exposed to. We did not create a zero-exposure habituation period because participants would have been entirely unfamiliar with the dialog and attractors.

We defined the duration of the fourth habituation period in units of time: 150 seconds, plus whatever additional time was required to dismiss the habituation dialog that was present at the moment the 150-second period expired. This corresponds to the 2.5 minute condition used by Bravo-Lillo et al. For this fixed-time treatment, the number of dialog exposures varies between participants, even within the same condition.

The addition of low-habituation conditions is what most differentiates our study from Bravo-Lillo et al.’s prior work. In the prior study, the researchers tested all participants at one high-habituation point (150 seconds or 22 exposures depending on the attractor). In addition, to accommodate longer habituation periods in our study, we told participants they would be spending ten minutes on our task, instead of the five minutes advertised by Bravo-Lillo et al. (primarily needed for the 20-exposure habituation period with participants in the Type condition).

As in the prior work, a participant is considered “attentive” if he or she chose the *No* option on his/her first test trial—the first trial in which he/she received the habit-breaking instruction.

To analyze the impact of habituation on each attractor, we examine the *habituation odds ratio* between two given habituation conditions. This is the ratio of participants who complied with the habit-breaking instruction in the lower of the two habituation conditions over that in the higher-habituation condition. This yields a $2 \times 2 \times 2$ contingency table: 2 treatments (attractor vs. control) \times 2 habituation conditions (lower vs. higher habituation) \times 2 outcomes (complied with the new instruction or did not). To test the null hypothesis that habituation caused the same reduction in the proportion of participants who complied with the instruction, regardless of treatment, we build a log-linear model without second-order interactions and use a likelihood-ratio test to compare this model to the observed data. If the observed data deviates significantly from the expected model, it indicates that the treatment might have an effect on the habituation odds ratio.

4. RESULTS

We ran this experiment from May 28 until June 09, 2013. We recruited a total of 3,071 participants for the study and 2,567 finished. Participants were 29.4 years old on average ($\sigma=10.1$ years), 55% male, 77% Caucasian, and the top two reported occupations were “student” (25%) and “unemployed” (15%). Based on user-agent strings, 60% of participants used Chrome, 37% used Firefox, and 3% used Internet Explorer.

For each participant, we consider the outcome a success if the participant chose the *No* option in response to the first dialog (test trial) in which they were instructed to do so, *complying* with this instruction. The *compliance rate* is the fraction of the participants in each condition who complied. We used a binomial outcome representing the result of the first test trial (complied vs. did not comply) with independent variables for the length of the habituation period

	Fixed exposure count							Fixed exposure time		
	1 exposure		3 exposures		20 exposures			150 seconds		
	med. time	R_{1e} (No·Yes)	med. time	R_{3e} (No·Yes)	med. time	R_{20e} (No·Yes)	$P \left[\frac{R_{1e}}{R_{20e}} = \frac{R_{1e}^c}{R_{20e}^c} \right]$	med. exp.	R_{150s} (No·Yes)	$P \left[\frac{R_{1e}}{R_{150s}} = \frac{R_{1e}^c}{R_{150s}^c} \right]$
Control	10 sec	50·56	3.4 sec	43·64	1.2 sec	24·90	—	192	7·99	—
ANSI	10.9 sec	57·55	3.9 sec	49·58	1 sec	15·95	= 0.1333	198	13·94	= 0.3466
AC + Delay	15.7 sec	89·18	9.8 sec	86·22	6.8 sec	65·43	= 0.9578	50	47·60	= 0.1933
Reveal	14.2 sec	84·25	8.4 sec	81·22	7 sec	57·47	= 0.6565	48	59·47	= 0.0021
Swipe	39 sec	61·45	6.9 sec	56·48	3.9 sec	59·48	= 0.0062	76.5	65·45	< 0.0001
Type	57.4 sec	79·33	16.6 sec	79·25	12.9 sec	86·13	< 0.0001	24	90·14	< 0.0001

Table 1: For each condition, we present the median of participants’ response times to their final habituation dialog (labeled *med. time*). We then present a count of the number of participants who chose the ‘no’ option on the first test trial (complying with the newly-introduced instruction) followed by the count of those who did not. Together, this is the compliance ratio R . Control group ratios are written R^c . The habituation odds ratio is the low-habituation compliance ratio over the high-habituation compliance ratio. To determine whether habituation had a greater or lesser effect in a treatment than in the control, we attempt to disprove the null hypothesis that their odds ratios are equal.

and the treatment (*attractor* or *control*). We present our results in Table 1 and graph the compliance rate as a function of the number of habituation exposures (log scale) in Figure 4. The level of habituation is measured by the number of exposures to the habituation dialog: 1, 3, and 20 exposures for the first three habituation exposures. The number of exposures varied for the fixed-time-period (150 second) conditions, so we use the median number of dialogs dismissed when plotting this point in Figure 4. Downward slopes represent a reduction in compliance.

For our *Control* dialog, which did not contain an attractor, the compliance rate starts low and declines steeply and steadily as the number of habituation exposures grows. The compliance rates of participants who saw the *ANSI* treatment were not significantly better, and were actually worse (though not significantly so) for the 20-exposure condition.

The two attractors that impose a delay but do not force the user to interact with the salient field, *Animated Connector + Delay* and *Reveal*, did best in the low-habituation conditions, but saw dramatic declines in compliance, with slopes similar to those seen for the control. We use habituation odds ratios to compare reductions in compliance in these attractor treatments with the reduction for the control. We are unable to reject the null hypothesis that the difference in reductions in compliance between the one-habituation-exposure condition and the 20-habituation-exposure condition was due to chance, suggesting that these conditions might not be more resistant to habituation than the control. The likelihood-ratio test yields a probability of the difference occurring under the null hypothesis of $p = 0.6565$ and $p = 0.9578$ (see Table 1).

In contrast, we were surprised to see the compliance rate for *Type* grow with the number of exposures. A possible explanation is that participants’ motivation to comply with the instruction, and to end the experiment early, may have increased as they grew tired of the task. This would represent a countervailing force that overpowers the minimal impact of decreased attention for this attractor. This force might be larger if users are particularly annoyed by an attractor. We use the habituation odds ratio to compare the (nonexistent) reduction in compliance due to habituation in the *Type* treatment with the threefold reduction for the control. We reject the null hypothesis that the difference in reductions in compliance between the one-habituation-exposure condition and the 20-habituation-exposure condition was due to chance, as the test yields a probability of the difference occurring under the null hypothesis of $p < 0.0001$. The same is true comparing the one-exposure condition with the 150-second condition (see Table 1).

Following *Type*, the *Swipe* attractor was second most resistant to habituation, with negligible reductions in compliance as habitua-

tion increased. Again we use the habituation odds ratio to compare this with the control. We reject the null hypothesis that the difference in reductions was due to chance with $p < 0.0062$ when looking at 1 vs. 20 exposures and $p < 0.0001$ when looking at 1 exposure vs. 150 seconds.

The lower reductions in compliance for *Animated Connector + Delay* and *Reveal* were not statistically significantly better than the control. Yet, the compliance rates for these two delay-inducing attractors were not far from that of *Swipe* under conditions of high habituation. Rather, they start from such a high initial level of compliance that they remain competitive even after significant reductions due to habituation.

However, the relative benefits of *Swipe* become more apparent when we examine the time required to interact with this attractor once participants are familiar with it. We measured the time each participant took to dismiss the last habituation dialog and calculated the median for each condition. Figure 5 shows the 25th, 50th, and 75th percentile dismissal time for the final habituation dialog in each condition. The time spent dismissing a dialog containing no useful information represents one component of the burden that attractors impose on users. Users quickly became efficient at interacting with *Swipe*. After three habituation exposures, participants learned to interact with the *Swipe* attractor as quickly as they did in the delay-based attractors. After 20 exposures, they were nearly twice as efficient in interacting with it as they were in the delay-based attractors. Nearly 75% of participants using the *Swipe* treatment were able to dismiss the 20th habituation dialog within five seconds.

4.1 Limitations

To create an experimental task that would allow us to vary habituation, we opted for a design that was necessarily artificial. Real-world habituation takes place over long periods of time. Security dialogs tend to be viewed one at a time with longer intervals between views, rather than rapidly during a ten-minute period. Also, users might use context in conjunction with the information in security dialogs to make a decision. Thus, users may behave differently when habituated in a more natural setting.

Since different attractors impose different delays, it was not possible to isolate the habituation effects of time and exposure count. Fortunately, this limitation does not appear to impact our conclusions. For the 20-exposure habituation conditions, the *Control* and *ANSI* dialogs required the least amount of time to complete, yielding the shortest habituation time periods, yet they saw the greatest reduction in compliance. In comparison, completing 20 trials took the most time for participants in the *Type* treatment, yet *Type* saw

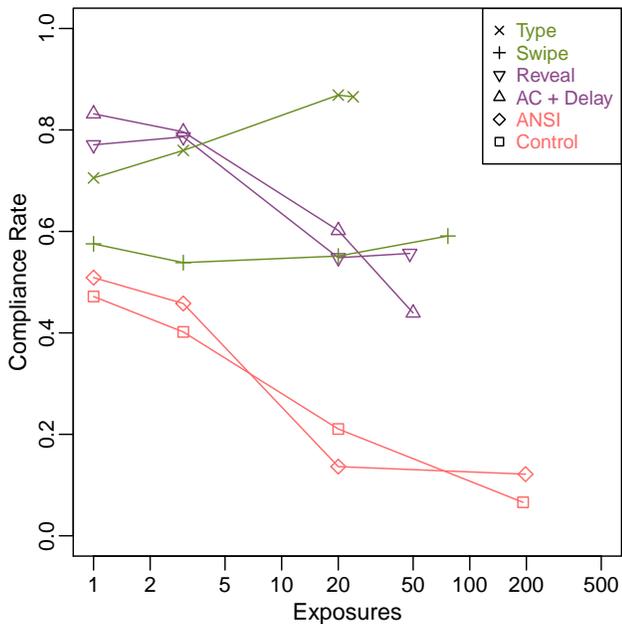


Figure 4: Compliance of participants to the instruction to click *No* in response to the first dialog in which they were asked to do so. The compliance rate is the number of participants who chose *No* over the total number of participants in that condition. The data from which this graph was generated can be found in Table 1.

an increased rate of compliance due to habituation.

Although at the beginning of the task we instructed participants to look for opportunities to finish the study early, some participants may not have paid attention to that directive or may have felt obligated to keep clicking on *yes* despite receiving the instruction in the ‘Status’ box. As our experimental design replicates that of Bravo-Lillo et al. [3], the same concerns may apply to their prior experiment.

One way to examine whether participants persevered with the experiment (clicking *yes*) despite having read the instruction to finish is to look to see if participants spent time resolving the conflict between the stated time period of the study and the instruction to finish early. If participants had noticed the change in the ‘Status’ box, they would presumably require some time to process what they had read. Searching for this decision lag, we examined the 50th, 75th, 90th, and 95th percentile response times at both the last habituation trial and the first test trial. We didn’t find any. For example, examining the Control treatment presented with the 150-seconds habituation period, there was only a 10 ms increase in response time from the last habituation trial to the first test trial, at the 50th percentile. We found a decrease in response time at the 75th, 90th, and 95th percentile. If even 5% of participants who chose ‘yes’ did so after reading the instruction and deciding to ignore it, the 95th percentile should show an increase in response time. The same was true when we examined every control and ANSI treatment. In contrast, when these participants eventually chose ‘no’ in later trials, the response time increased by a factor of 3 or more, depending on the specific percentile.

We also asked participants who clicked *yes* in the first test trial to explain their behavior. The most popular reason given was not noticing, among myriad other reasons (e.g., one participant reported understanding the instruction and deciding to follow it, but then accidentally clicking *yes* out of sheer habit). The answers reveal only

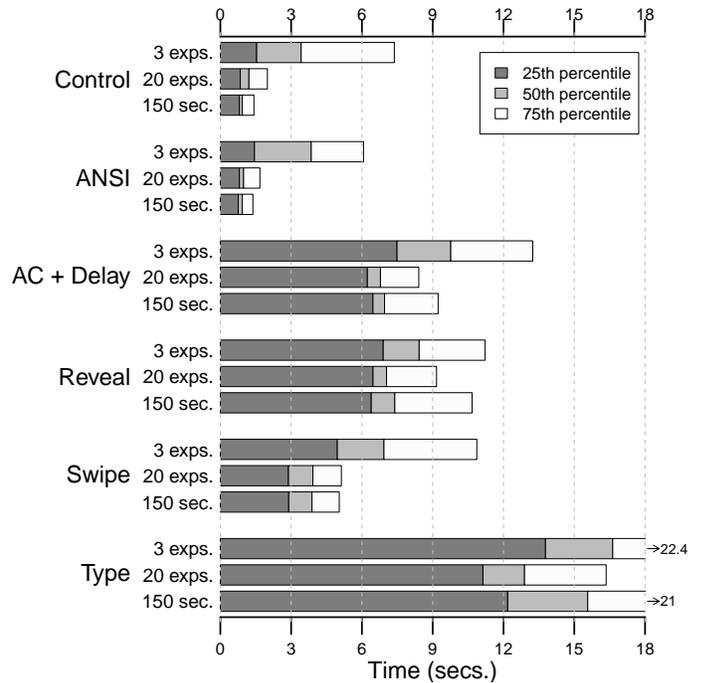


Figure 5: 25th, 50th and 75th percentiles of participants’ response time to the last habituation dialog.

a small minority who believed that we wanted them to persevere or that they would be rewarded for persevering despite our instruction that they should look for opportunities to finish the study early.

One factor that may have affected our participants’ responses was the appearance of the *No* button during the habituation period. When we disabled the functionality of the button we did not change its appearance to reflect that it was disabled, lest it would transition from a disabled to enabled appearance at the same time the habit-breaking instruction first appeared. Had the option changed appearance, we would not have been able to separate the effect of participants noticing this change from the effect of participants noticing the habit-breaking instruction. However, it’s possible that, as a result of this design choice, some participants disregarded the habit-breaking instruction because they believed clicking the *No* option would have no effect. Fortunately, we see no reason this behavior should be any more likely to occur in one condition than another, and so it should not impact cross-group comparisons.

5. DISCUSSION

While Bravo-Lillo et al.’s prior study demonstrated that some attractors performed well under conditions of heavy habituation, it was not clear whether these attractors were actually resistant to habituation. By testing four levels of exposure we found that some of the attractors that performed well in Bravo-Lillo et al.’s study, specifically *Reveal* and *AC+Delay*, become less effective with repeated exposure. As expected, in the absence of attractors, increasing habituation in the control condition caused a three-fold decrease in performance. On the other hand, two attractors, *Swipe* and *Type*, remained effective even after many exposures. These results have implications for security dialog design and also highlight the value of conducting habituation trials.

Of the attractors we tested, the *Type* attractor performed best, but also imposed the greatest usability burden. While an attractor

of this type may not be realistic for the majority of user environments where users can circumvent systems when frustrated, it may be useful in security-critical environments where lack of habituation should be prioritized over the usability burden. While the *Swipe* attractor had a lower compliance rate than *Type*, it was also resistant to habituation, and it demonstrated a reduction in usability overhead over time as users learned to use it more quickly. Thus, *Swipe* may be a good option in environments where the usability burden associated with *Type* is unacceptable.

This study also demonstrates the value of conducting habituation trials at various levels of exposure. In Bravo-Lillo et al.'s prior habituation experiment, all five inhibitive attractors outperformed the control at the one high-habituation level tested. However, our study demonstrates that some of these high-performing attractors actually become less effective with additional exposures. This suggests that when evaluating a real-world security dialog, it would be useful to test habituation at more typical exposure frequencies as well.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grants No. CNS1116934 and DGE0903659, and by the ARCS Foundation.

6. REFERENCES

- [1] D. Akhawe and A. P. Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Proceedings of the 22th USENIX Security Symposium*, 2013.
- [2] R. Böhme and S. Köpsell. Trained to Accept?: a Field Experiment on Consent Dialogs. CHI'10, pages 2403–2406. ACM, 2010.
- [3] C. Bravo-Lillo, L. Cranor, J. Downs, S. Komanduri, R. Reeder, S. Schechter, and M. Sleeper. Your Attention Please: Designing Security-Decision UIs to make Genuine Risks Harder to Ignore. SOUPS '13, pages 1–12. ACM, 2013.
- [4] J. C. Brustoloni and R. V. Salomón. Improving Security Decisions with Polymorphic and Audited Dialogs. SOUPS '07, pages 76–85. ACM, 2007.
- [5] L. F. Cranor. A Framework for Reasoning about the Human in the Loop. UPSEC'08, pages 1–15, Berkeley, CA, USA, 2008. USENIX Association.
- [6] S. Egelman, L. F. Cranor, and J. Hong. You've been Warned: an Empirical Study of the Effectiveness of Web Browser Phishing Warnings. CHI '08, pages 1065–1074. ACM, 2008.
- [7] C. Karlof, J. D. Tygar, and D. Wagner. Conditioned-safe ceremonies and a user study of an application to web authentication. In L. F. Cranor, editor, *SOUPS*, ACM International Conference Proceeding Series. ACM, 2009.
- [8] K. Krol, M. Moroz, and M. A. Sasse. Don't work. can't work? Why it's time to rethink security warnings. In *7th international conference on Risk and security of internet and systems (crisis)*, pages 1–8, 2012.
- [9] D. L. Schacter, D. T. Gilbert, and D. M. Wegner. *Psychology*. Worth Publishers, 2009.
- [10] A. Sotirakopoulos, K. Hawkey, and K. Beznosov. On the Challenges in Usable Security Lab Studies: Lessons Learned from Replicating a Study on SSL Warnings. SOUPS '11, pages 3:1–3:18. ACM, 2011.
- [11] J. Sunshine, S. Egelman, H. Almuhiemedi, N. Atri, and L. F. Cranor. Crying Wolf: an Empirical Study of SSL Warning Effectiveness. USENIX '09, 2009.
- [12] R. F. Thompson and W. A. Spencer. Habituation: a model phenomenon for the study of neuronal substrates of behavior. *Psychological review*, 73(1):16, 1966.