

Malicious Interfaces and Personalization's Uninviting Future

Being online isn't as fun as it used to be. Online advertising increasingly intrudes, distracts, and interferes with accomplishing even simple tasks; purposely misleading menus make navigation confusing and unintuitive; and far too often,

GREGORY
CONTI
AND EDWARD
SOBIESK
*US Military
Academy at
West Point*

the cost of getting to desired content involves divulging personal information that isn't relevant to the task at hand. In short, malicious interface designs are propagating across the Web.

Today's malicious interfaces are largely untargeted—crudely attempting to trick, coerce, or otherwise manipulate users into taking some undesired action, such as disclosing private information, making a purchase, clicking on an advertising link, registering a rarely used software application, failing to read an end user license agreement, or quietly signing the user up for undesired mailing lists through pre-checked boxes.

The future promises a bleaker scenario. Online companies don't hide their desire to "personalize" the user experience and conduct targeted advertising. Widespread data collection, retention, and mining are now occurring, and targeted use of this data to personalize malicious interfaces will follow shortly. With the dramatic increase in strategies such as cross-site tracking, which involves massive data collection, retention, and mining of multi-site visits, online

companies are building the information stockpiles necessary to create the personalized experiences promised in the next decade. Given malicious interfaces' prevalence, we argue that we can expect much of our personalized future to be tailored malicious interfaces in which the interface designer is an adversary who is also armed with profiling information to make his techniques even more effective.

In this article, we explore the issue of malicious interface design and describe why these problems are only going to get worse. We also present a taxonomy of malicious interface techniques and discuss currently available countermeasures to address this challenge. We consider this work a call to the security, privacy, interface design, and online advertising communities to start addressing these issues now.

The Interface Designer as an Adversary

In an ideal world, interfaces help users accomplish tasks quickly, easily, and efficiently. However, in the real world, the opposite often occurs. Interfaces are frequently designed to manipulate users into

taking action or revealing information that they didn't intend to. We define a malicious interface as an interface design that consciously attempts to achieve the designer's objectives ahead of the user's by employing techniques that negatively impact user experience. Malicious interfaces abound when a designer sacrifices users' time, attention, and personal data by exploiting the user's perception to achieve desired ends. Thus, the interface designer becomes an adversary who doesn't attack through traditional network vulnerabilities or malware but instead targets human cognition.

Malicious interfaces shouldn't be confused with bad design. Rather, they're deliberate attempts to exploit users and don't occur by accident. Human-based attacks aren't new to computing. Active research is ongoing in social engineering,^{1,2} phishing,³ denial of information,^{4,5} and attacks on information visualization systems.⁶ With online malicious interfaces, we face a new and very potent adversary, an interface designer who aggressively employs quasi-legitimate technical means to accomplish an end motivated by an intensely competitive and immature Web-based advertising model. Malicious interfaces employ trickery, browsing misdirection, forced advertisement viewing, and spoofed interface elements, among many other techniques. We see malicious interfaces both on and off the desktop, virtually anywhere

profit is at stake, including operating system distributions littered with “trial” commercial software, gas pumps designed to subtly convince users to choose premium gasoline or buy a car wash, retail checkout devices that hide our ability to use a debit card as a credit card, and even toothpaste dispensers that dispense more than the necessary amount of toothpaste.⁷

On the Web, we've all maneuvered through misleading links, disabled back buttons, browsers with “sponsored” default bookmarks, unexpected and unnecessary forms, blinking advertisements, and pop-ups covering desired content. It's likely you've been coerced into disclosing information on HTML forms requiring functionally unnecessary but mandatory fields. Perhaps you've heard about the attack against the non-profit Epilepsy Foundation's online forums, in which attackers embedded flashing animations designed to induce epileptic seizures.⁸

Figure 1 shows an example in which a user must fight the interface. Although it's difficult to tell, this is actually a newspaper Web site. The large automotive advertisement at the top expands automatically as the mouse pointer rolls over a hot region. The drop-down menu, rather than cleanly helping users navigate, also contains a large advertisement. Users visiting the site to learn about car dealers are in luck, but those interested in the news will have to struggle with the interface.

By no means are all Web designers and advertisers constructing malicious interfaces. The dramatic increase, though, in the development and use of techniques clearly designed to trick, mislead, consume attention, elicit sensitive information, or otherwise subvert users to achieve the designer's objectives have become so prevalent that all users must at a minimum become far more sensitive to these dangers; we as professional tech-

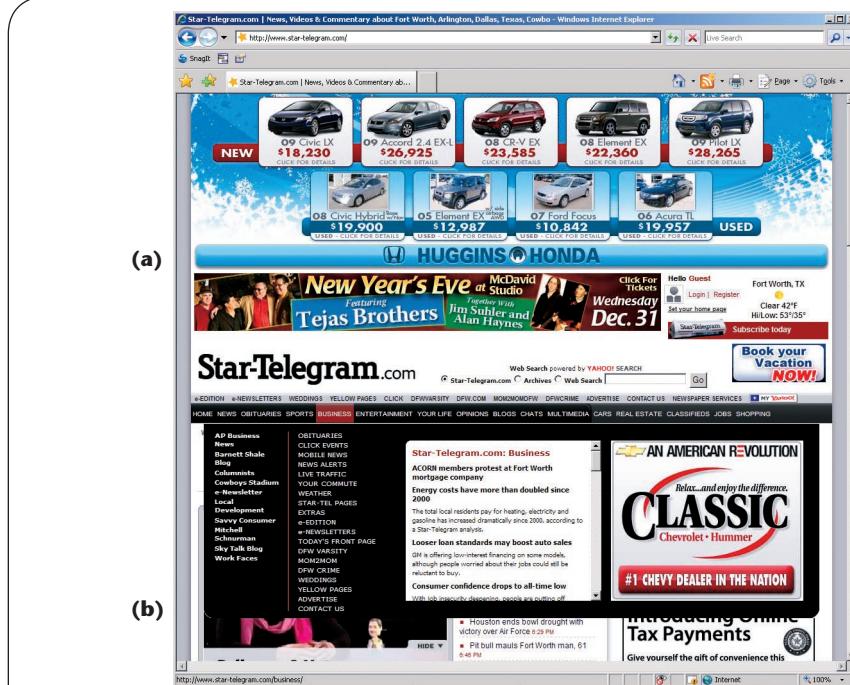


Figure 1. A news Web site demonstrating a malicious interface. (a) An errant mouse movement can trigger an expanding advertisement at the top, and (b) the drop-down menu includes a second large advertisement rather than a clean navigational interface.

nologists must aggressively fight against these downward trends in interface design.

The Gathering Storm of a “Personalized” Future

The problem of personalized interfaces is compounded by the ease with which online companies can conduct real-time experiments to determine a new advertising or interface design technique's success or failure and immediately make production changes. Large online companies, such as Google, Yahoo!, and Microsoft, frequently laud the value of online experiments and personalization's future on the Web. User profiling and the ease of online experimentation combined portend a future in which large percentages of the Web user population will be subject to extremely sophisticated manipulation.⁹

If you're reading this article, chances are you already have the

skills to help protect yourself from these threats, but for most of the population, this isn't the case. One of the saddest aspects of malicious interfaces is that the most vulnerable users are at greatest risk. The young and the old as well as the physically, cognitively, and perceptually challenged regularly encounter deliberately designed malicious interfaces, can't defend themselves, and perhaps don't even realize they've been attacked and exploited. Crafting a successful interface for these user populations is difficult in the best case, but malicious interface techniques degrade or even render unusable large swaths of the Web.

What motivates interface designers to deliberately violate design best practices and attack the user? Clearly, financial motivation is the most common reason—if a subtle, or not so subtle, change to an interface makes an observable improvement in the financial bottom line or corporate information assets, employing malicious

techniques is a real possibility. The situation worsens when the Web's prevalent business model uses advertising to pay for nominally "free" online tools. In the past, we argued that these tools aren't actually free—we pay for them with micropayments of personal information.¹⁰ To this sum, we now add our time, attention, and frustration expended by dealing with malicious interfaces.

A second, more subtle, reason for the increase in such interfaces is a culture of one-upmanship currently infecting the online advertising domain. If Company A has a blinking advertisement, then Company B needs an advertisement that not only blinks but also enlarges when a user inadvertently hovers the mouse pointer nearby; so, the momentum of malicious techniques has quickly spiraled into the state it is today.

Attack Techniques

Malicious interfaces fall into several broad categories, but all share common underpinnings—the attack techniques exploit human weaknesses and limitations, including cognition, perception, memory retention, and dexterity. In addition, malicious interfaces also consume users' time, attention, and personal information. The attacks deliberately frustrate user task accomplishment in order to facilitate achievement of the interface designer's objectives, even when these objectives fall in direct opposition to users' desired goals. For example, a user might wish to read a news story, but the interface designer wants to maximize advertisement click-throughs, so he or she designs an advertisement to look like a link to a news article. Table 1 lists a taxonomy of many current attack techniques.

These categories are examples of what we've detected in the wild, but this list is just a starting point. Other techniques exist, and new techniques are being developed on

a regular basis. A malicious interface designer need only invert current (or future) design best practices to create new techniques. Adding to the frustration level is that although each malicious technique can be used in isolation, they're often combined to increase the malicious design's effectiveness while significantly decreasing the value of the user's online experience.

Countermeasures

Countering malicious interfaces is an open problem that requires emphasis and research from both the security and human-computer interaction communities. However, the problem is tractable, and we can approach it through legal, political, economic, social, and technological means.

Many malicious interfaces are in direct opposition to accessibility laws and guidelines such as the Americans with Disability Act and the UK's Disability Discrimination Act. More aggressive enforcement of existing laws and continued regulatory and political scrutiny will help reduce malicious interfaces' prevalence. Economic and social pressures are also viable complements to political pressure. Because malicious interfaces are often profit-driven, affected communities can generate public opposition and encourage consumers to avoid companies using malicious techniques, thereby threatening profit margins and inducing positive change.

In cases in which we can mitigate malicious interfaces via technological means, the positive results can be dramatic. In some cases, users have no direct control over the interface, such as at a gas pump, but with one of the worst offenders—the Web—users have significant control because their browsers render the interface. Countermeasures, such as Greasemonkey (www.greasespot.net), which uses small snippets of JavaScript to customize the Web

interface, ad-blocking software such as Adblock Plus (adblockplus.org), NoScript (noscript.net), which disables untrusted active content, and personal proxies such as Privoxy (www.privoxy.org), which allows inline Web traffic modification and filtering, provide resources that reduce malicious interfaces' impact on the Web. However, only advanced users typically employ many of these tools. We must conduct continued research to create usable combinations of these techniques that popular browser distributions can widely deploy by default. The successful integration of pop-up blocking, which popular browsers currently enable by default, is an important precedent that future countermeasures should emulate.

Of course, online companies argue that to provide you with perfect service, they need to have as close to perfect knowledge about you as possible. In fact, one of Google's stated goals is to understand you so well that you can ask Google whether you should take a particular job.¹¹ Although we concur that knowledge about users can improve service, we believe, and have previously argued, that such knowledge is a tremendous corporate asset that's ripe for abuse and exploitation, and the creation of personalized malicious interfaces using such knowledge is an ethical and legal issue that will arise over the next few years. Experts in computer security, computer-human interaction, privacy, online advertising, and online services must unite to seriously address this issue and seek solutions. Every year without such action dramatically increases the challenge's severity. □

Acknowledgments

We thank Gillian "Gus" Andrews as well as the BlackHat, DEFCON, and HOPE communities for their support and insights regarding this research.

Table 1. Classes of malicious interface techniques.

ATTACK CLASS	DESCRIPTION	EXAMPLE
Coercion	Forcing the user to take an undesired action	Requiring disclosure of a user's home and work phone number or other personal information on a Web form
Confusion	Placing the user in a situation in which progressing toward task accomplishment is uncertain	Providing deliberately vague or contradictory instructions on a Web site
Distraction	Exploiting perception to attract users toward the designer's goal	Employing blinking advertisements on a Web page while users are trying to read a news story
Error inducement and exploitation	Creating interfaces designed to induce user errors or exploiting errors to facilitate designers' desired goals	Redirecting users to an advertisement-laden error page when they mistype a URL
Forced work	Increasing user workload to "punish" the user for noncompliance with the designer's goals	Forcing the user to manually uncheck numerous checkboxes to avoid undesired mailing lists
Interruption	Interrupting users' task flow to divert them toward the designer's goal	Frequently interrupting users and asking if they would like to register a software application
Manipulated navigation	Creating navigation systems designed to frustrate user tasks and divert users toward designer-driven task accomplishment	Making free versions of software difficult to find on a site while guiding users toward paid versions
Obfuscation	Making desired content and functionality difficult to identify or use	Placing a large end user license agreement in a small scrolling text box
Restricted functionality	Creating interfaces such that desired functionality is difficult or impossible	Using JavaScript to disable a browser's back button to "trap" users on a Web page
Trick	Deliberately attempting to deceive the user through malicious content or interface design elements	Designing advertisements to spoof legitimate interface elements

The views expressed here are the authors' and do not reflect the official policy or position of the US Military Academy, the Department of the Army, the Department of Defense, or the US Government.

References

1. K. Mitnick and W. Simon, *The Art of Deception: Controlling the Human Element of Security*, Wiley, 2003.
2. J. Long and J. Wiles, *No Tech Hacking: A Guide to Social Engineering, Dumpster Diving, and Shoulder Surfing*, Syngress, 2008.
3. T. Jagatic et al., "Social Phishing," *Comm. ACM*, vol. 50, no. 10, 2007, pp. 94–100.
4. G. Conti and M. Ahamad, "A Framework for Countering Denial of Information Attacks," *IEEE Security & Privacy*, vol. 3, no. 6, 2005, pp. 50–56.
5. M. Ahamad et al., "Guarding the Next Internet Frontier: Counter-
- ing Denial of Information Attacks," *Proc. New Security Paradigms Workshop*, ACM Press, 2002, pp. 136–143.
6. G. Conti, M. Ahamad, and J. Stasko, "Attacking Information Visualization System Usability: Overloading and Deceiving the Human," *Proc. Symp. Usable Privacy and Security (SOUPS 05)*, ACM Press, 2005, pp. 89–100.
7. G. Conti, "Evil Interfaces: Violating the User," Hackers on Planet Earth (HOPE), July 2008; www.thelasthope.org/talks.html.
8. K. Poulsen, "Hackers Assault Epilepsy Patients via Computer," *Wired*, 28 Mar. 2008; www.wired.com/politics/security/news/2008/03/epilepsy.
9. C. Doctorow, "Future Tense: Pester Power," *Comm. ACM*, vol. 51, no. 12, 2008, pp. 119–120.
10. E. Sobiesk and G. Conti, "The Cost of Free Web Tools," *IEEE Security & Privacy*, vol. 5, no. 3, 2007, pp. 66–68.
11. C. Daniel and M. Palmer, "Google's Goal: To Organise Your Daily Life," *Financial Times Online*, 22 May 2007; www.ft.com/cms/s/2/c3e49548-088e-11dc-b11e-000b5df10621.html.

Gregory Conti is an assistant professor in the Department of Electrical Engineering and Computer Science at the US Military Academy at West Point. He has a PhD in computer science from the Georgia Institute of Technology. Contact him at gregory.conti@usma.edu.

Edward Sobiesk is an associate professor in the Department of Electrical Engineering and Computer Science at the US Military Academy at West Point. He has a PhD in computer and information sciences from the University of Minnesota. Contact him at edward.sobiesk@usma.edu.