

Usable Privacy and Security: A Grand Challenge for HCI

Jason I. Hong

Carnegie Mellon University

5000 Forbes Ave, Pittsburgh PA 15213

jasonh@cs.cmu.edu

ABSTRACT

In this position paper, we argue that usable privacy and security is a grand challenge that needs more attention from the HCI community. We also discuss benefits to and new challenges for HCI, and use our research experiences to provide a critique of HCI.

Usable Privacy and Security: A Grand Challenge for HCI

Jason I. Hong

Carnegie Mellon University

5000 Forbes Ave, Pittsburgh PA 15213

jasonh@cs.cmu.edu

ABSTRACT

In this position paper, we argue that usable privacy and security is a grand challenge that needs more attention from the HCI community. We also discuss benefits to and new challenges for HCI, and use our research experiences to provide a critique of HCI.

INTRODUCTION

Information and communication technologies are pervasive in all aspects of every day life, including transportation, manufacturing, utilities, finance, and entertainment. It would not be an understatement to say that modern society depends on these systems being highly reliable. However, we have been witnessing an increasing number of privacy and security failures in these systems. What is interesting is that many of these failures happen not because of breakdowns in algorithms or hardware, but because of failures in the user interface.

As we become more reliant on computer infrastructures, the consequences of security breaches are becoming more severe. In 2003, the CRA issued a grand challenge for computer security and privacy: "Give end-users security controls they can understand and privacy they can control for the dynamic, pervasive computing environments of the future." Similarly, the National Academy of Engineering (NAE) included "secure cyberspace" in their 2008 Grand Challenges for Engineering, arguing that more research is needed on the psychology of computer users, how people interact with their computers, and how "cultural and social influences can affect how people use computers and electronic information in ways that increase the risk of cybersecurity breaches."

The emerging field of usable privacy and security draws on ideas from HCI, computer security, and many other fields, to develop human-centered systems for managing security and privacy that are effective in practice. In this position paper, we discuss three things. The first is a case study of our research and experiences with usable privacy and security, in the context of protecting people from online phishing scams. The second is a discussion of challenges for usable privacy and security, discussing opportunities for how HCI can help and how HCI can also benefit. The third is a critique of perceived problems with the CHI community, using examples from our experiences to illustrate the problems as well as ways of addressing these problems. Many parts of our discussion are opinionated and

speculative, with the intention to stir up discussion of where our field should be going.

USER INTERFACES AND ALGORITHMS FOR ANTI-PHISHING

Phishing is a semantic attack that targets the users of a system rather than the hardware or software. The most common semantic attack today is phishing, where criminals impersonate legitimate people or organizations and trick people into giving up sensitive information or installing malware on their computers. Perhaps the most common form of phishing are fake "please update your account" emails that direct people to sites that appear like real sites. Gartner group estimated that phishing caused \$3.2 billion dollars in direct damages in the United States in 2007. Note that this does not include indirect damage to an organization's reputation or loss of potential sales. This figure also does not include spear-phishing attacks, where criminals target specific individuals using a great deal of contextual information about them. Spear-phishing attacks have been used to illicitly obtain sensitive information from corporations, governments, and military.

Our group is comprised of 5 different faculty, with expertise in decision sciences, economics, machine learning, computer science, public policy, and human-computer interaction. We have identified three general strategies to protect people from phishing, which generalize into three basic strategies for usable privacy and security: (1) make it invisible, (2) make it understandable, through better awareness, usability, and metaphors, and (3) train users. Our group has been pursuing all three strategies for anti-phishing. Our overarching philosophy is that we should automate where possible, but since no algorithm will always be 100% accurate, we must also support end-users with better interfaces and better training where necessary.

In terms of automation, we have made use of sophisticated machine learning algorithms and information retrieval algorithms to find phishing emails and web pages. We have also conducted a series of studies evaluating the effectiveness of commercial anti-phishing toolbars that detect phishing web sites. In terms of better user interfaces, we have conducted several studies to understand whether existing browser warnings notifying end-users of danger are effective or not. In terms of training, we have developed two novel training mechanisms: PhishGuru, where system administrators might pro-actively sending fake phishing

emails to train people in their organization, and Anti-Phishing Phil, a game that teaches people about phishing.

Thus far, our work has generated a great deal of interest and collaboration from a number of partners. Our automated email filter is undergoing a field trial at Carnegie Mellon University's main email servers, where it will filter several million emails per day. Our research evaluating anti-phishing toolbars has been cited by several companies, with ongoing evaluations being presented to the Anti-Phishing Working Group, a consortium of companies "committed to wiping out Internet scams and fraud." Design suggestions from our studies to understand browser warnings have been incorporated into the latest version of Microsoft's Internet Explorer 8. PhishGuru's methodology of sending fake phishing emails to train individuals has undergone field trials at three different companies, and been cited by two different companies trying to commercialize the work. PhishGuru's training materials have also been adopted by APWG on their landing page, a page that ISPs and web sites can show after taking down a phishing web site. Anti-Phishing Phil has been played by over 80,000 people, licensed by two companies, demoed at many security days meant to teach people about good security practices, and translated into Portuguese with several more translations underway. Finally, our group is commercializing all of this work through a startup we have founded, named Wombat Security Technologies.

CHALLENGES FOR HCI IN USABLE PRIVACY AND SECURITY

Although the research community has made a great deal of progress, there are still many open challenges that remain to be addressed, and HCI has much to offer in this regard. Here, we outline three areas of work. First, there is still a great need for better methods, tools, and design patterns that align privacy and security with usability. Too often, critics see this as an either/or decision, that increasing security may reduce usability, and vice versa. Second, there is also a pressing need for better methods and tools for running user studies that are rigorous, realistic, and ethical. Third, we need the equivalent of discount usability techniques for privacy and security. It is simply too expensive to user test every use case to ensure that there is an acceptable level of success. Usable privacy and security needs new evaluation methods akin to heuristic evaluation and cognitive walkthrough, as well as new models analogous to GOMS, to help drive the field forward.

USABLE PRIVACY AND SECURITY IS GOOD FOR HCI

Here, we draw on our experiences with usable privacy and security at Carnegie Mellon University, using it as a lens for critiquing the CHI community. From an educational perspective, our course in usable privacy and security has been successful in teaching people about basic HCI principles, in particular people who would probably not normally take HCI courses. This course has also been useful as outreach towards a community that has not

traditionally linked to HCI, creating another advocate for HCI within our university.

From a relevance standpoint, privacy and security issues appear in the media practically every day. As noted in the introduction, usable privacy and security has been declared a grand challenge by the National Academy of Engineering and by the CRA. Michael Chertoff, the US Secretary of Homeland Security, has called for a Manhattan Project to secure the cyberdefenses of the national infrastructure. As one can see, it is fairly easy to argue that privacy and security are of paramount importance to companies and to national security. However, one problem all of us in the HCI community have faced is that it is much harder for us to make this argument about our discipline. An alternative path may be for HCI to partner more strongly with other disciplines, which will lead to more communities that can advocate for us.

From an impact perspective, to some extent, our group was also surprised by the amount of interest by industry in our research, so much so that we formed a startup company to commercialize the work. One business professor explained it this way: is what you are building a feature, a product, or a business? This simple insight explains why a non-trivial amount of HCI research, while being rigorous and influential within the research community, has such a difficult time impacting the design of actual products. A non-trivial amount of research focuses on improving features or improving efficiency by a marginal amount, but unless the research solves a big enough pain that people are willing to pay for, it is simply unlikely to be adopted. In contrast, research that focuses on new kinds of products or even new businesses may find it easier to commercialize, but this kind of research is more difficult, and has potentially more holes and thus harder to publish.

CONCLUDING REMARKS

While on the surface security and HCI may seem far afield from each other, these two communities actually share much in common. Like HCI, effective security requires a holistic view. Just as a single interface design flaw can prevent users from completing their tasks, a single vulnerability can compromise an entire system. Like HCI, many security practitioners lament that security is often treated as an afterthought, and that they are all too often expected to slap on security rather than being part of the design from the beginning.

We hope that this position paper has opened up new possibilities for usable privacy and security and offered food for thought as to directions for our community.

ACKNOWLEDGMENTS

Special thanks to Alessandro Acquisti, Lorrie Cranor, Julie Downs, Norman Sadeh, and the students in our research group. Also, thanks to all of the people over the years who have helped refine the arguments in this position paper. The opinions in this paper are solely those of the author.