THESIS DEFENSE

Institute for Software Research

Software Engineering

Informing the Design and Refinement of Privacy and Security Controls

Daniel Smullen

Friday, September 24th, 2021
Online via Zoom
10:00 AM - 1:00 PM

Amid increasing privacy and security risks, managing one's privacy and security settings is becoming ever more important. Yet, the proliferation of security and privacy controls is making this task overwhelmingly complex. Are they the right controls? Are they effective? This dissertation's objective is to study how effective existing settings are, assess whether they give users the awareness and control they need, and to inform ways to improve them.

We begin by examining how people interact with browsers' privacy and security settings. This is followed by a study designed to inform the development of more effective settings and defaults. Finally, we explore machine learning techniques with the aim of helping users configure their settings and further reduce user burden. Our results form the basis for our recommendations to improve privacy and security controls, the discussion of public policy implications, and generalizability to other domains.

This is the first dissertation to explore a broad cross section of privacy and security decisions, systematically exploring their effectiveness and manageability. We reveal that existing privacy and security controls may not be effectively addressing people's concerns or expectations. However, the problem is fundamentally about having the right settings, not necessarily the most options, as this fails to consider the limits of what people are realistically capable of configuring. To avoid redundancy and confusion, the settings also need to align with people's mental models. Moreover, people's diverse preferences and concerns can align across categories of apps and websites, data practices, purposes, and many other factors -- these can form the basis for consolidation and standardization. Yet standardized settings, such as mobile app permissions, can still be misaligned with people's mental models. Simply adding more expressive settings is a tempting solution but improving control and effectiveness by proliferating settings can trade-off manageability and increase user burden. Machine learning can simplify the task of managing one's settings, which can help to overcome this trade-off. Privacy and security controls can be redesigned to be more effective -- without exceeding users' ability to configure them.

Committee: Dr. Norman Sadeh (Chair), Dr. Lorrie Faith Cranor, Dr. Alessandro Acquisti, Dr. Rebecca Weiss (Mozilla), Dr. Yaxing Yao (University of Maryland, Baltimore County)

institute for SOFTWARE RESEARCH | Carnegie Mellon University School of Computer Science