RESEARCH STATEMENT • *Patrick Gage Kelley*

Users are increasingly expected to manage complex privacy settings in their normal online interactions. From shopping to social networks, users make decisions about sharing their personal information with corporations and contacts, frequently with little assistance. Current solutions require consumers to read long documents or go out of their way to manage complex settings buried deep in management interfaces, all of which lead to little or no actual control.

My goal is to help people cope with the shifting privacy landscape. While my work looks at many aspects of how users make decisions regarding their privacy, my dissertation focuses on two specific areas: the current state of web privacy policies and mobile phone application permissions. I explored consumers' current understandings of privacy in these domains, and then used that knowledge to iteratively design and test more comprehensible information displays.

Online Privacy Policies

Website privacy policies should *help* consumers decide which companies they are willing to share their information with. Privacy policies are supplied by companies as a form of self-regulation to meet the FTC requirement of "notification." By having notified consumers of what information will be collected, how it will be used, and with whom it will be shared, consumers are, in theory, able to make informed decisions. These policies are also meant to inform consumers of the choices they have in managing their information: whether use of their information or sharing with third parties can be limited and if it is possible to request modification or removal of their information.

However, privacy policies are commonly long, textual explanations of data practices, frequently written by lawyers to protect companies against legal action. Our research, and that of others, indicates users believe locating information in today's online privacy policies is difficult and time-consuming; thus, privacy policies are rarely read.

I began this work with the goal of designing a privacy policy format that would actually benefit consumers, as several prior efforts had failed. Our goal was a format that would help users accurately and quickly find information, make comparisons between policies easier, and provide a more enjoyable experience [1]. Our design approach allowed us to explore other efforts in standardizations, labeling, and designing privacy policy formats, while quickly and iteratively building a library of testable privacy label formats.

Our final label design allows for information to be found in the same place every time. It removes leeway and complicated terminology by using four standard symbols that can be easily compared. It provides quick high-level visual feedback when looking at the overall intensity of the page, can be printed to a single sheet, fits in a browser window, and has a glossary of useful terms attached. People who used it to find privacy information rated it not only better than the text policies, but actually enjoyable to use.

A snapshot of our iterative design process for our privacy label.

Our design performed significantly better across a variety of measures than the full-text and layered-text policies that currently exist online today [2]. The large amount of text in full-text policies and the necessity of drilling down through a layered policy to understand specific practices lengthens the amount of time and effort currently required to understand a policy. Additionally, more complex questions about data practices frequently require reading multiple sections of these text policies and understanding the way different clauses interact, which is not an easy task. Our format simplifies these interactions, removes the complicated terminology and dense language, provides standardized descriptions of privacy practices, and provides tested definitions of those terms to assist consumers. Our standardized policy left no room for erroneous, wavering, or unclear text, serving as a concise text alternative to tabular formats.

For this work, I was awarded First Place at the ACM SIGCHI 2009 Student Research Competition and and then First Place in the ACM Grand Finals in June 2010. The work was also selected for inclusion in 2010's *Papers for Policy Makers*, produced by the Future of Privacy Forum.

Smartphone Application Permissions

Since the launch of the first Android phone in October 2008, both consumers and phone manufacturers have adopted the platform in record numbers. Android phones accounted for more than half of all smartphone sales as of Q3 2011. With each smartphone sold, more users are downloading applications from the Android Market. Applications are not pre-screened for inclusion in the market; instead, the market requires users to make two choices when reviewing potential applications for their device:

1. Do I believe this application will compromise the security of my phone if I install it?
2. Do I trust this developer with access to my personal information?

Users must make these decisions based on advertising, word-of-mouth information, market reviews and ratings, and the Android permissions display. I sought to examine how the Android permissions display is understood and how it could be better designed to help users

make these decisions as the only required display and trusted information source they are provided.

To reach a deeper and more nuanced understanding of how people navigate the current Android ecosystem, I conducted semi-structured interviews in summer 2011 with 20 participants in Seattle and Pittsburgh [3]. The interviews were exploratory in nature, seeking broad understanding of participants' interactions with their smartphones as well as diving deeply into issues surrounding the display of permissions, the safety of the Android Market, and possible harms of information sharing.

Most importantly, I found that while users recall viewing the permissions display, they do not understand Android permissions.

Specifically, the human-readable terms displayed before installing an application are at best vague and at worst confusing, misleading, jargon-filled, and poorly grouped. This lack of understanding makes it difficult for people to make informed decisions when installing new software on their phones. For the most part, the permissions are ignored, with participants instead trusting word of mouth, ratings, and Android market reviews.

Users are also largely uninformed about the potential for malware or malicious applications to exist in the Android Market. They have difficulty describing the possible harm that could be caused by applications collecting and sharing their personal information. More alarmingly, I find that people have substantial misconceptions about their safety when installing applications from the Android Market. While participants said they try to find good applications in the market, they believe they are protected by oversight processes which do not exist.

Overall, users are not currently well prepared to make informed privacy and security decisions involving installing applications from the Android market.

To begin to address this, I have launched an Android permissions display redesign project. I am considering two main design threads, one which simply rearranges the information and uses clearer terminology and examples, and one which adds additional information on data sharing, advertising packages, and why applications are making permissions requests. This work is currently underway.

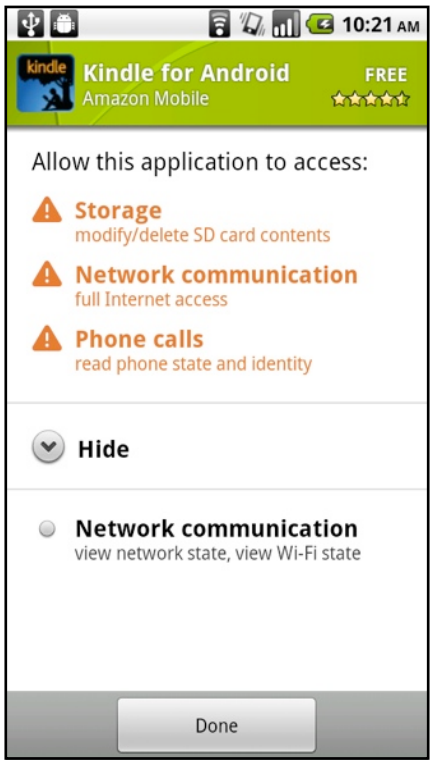

Example Android permissions display for Amazon's Kindle application.

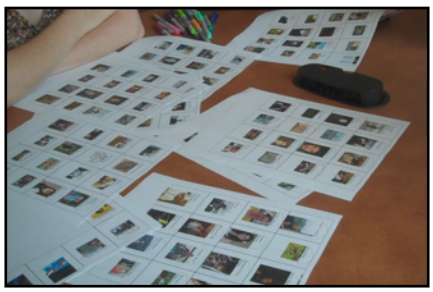

Participants grouping their Facebook friends through physical "interfaces."

RESEARCH AGENDA

Although I am presently completing my dissertation research, I have already begun to plan and lay the foundation for my future research agenda. While I intend to continue my focus on usability and privacy, I hope to cultivate a robust library of design patterns and apply these to other domains I am interested in.

Social Networks

Social networks have been a hotbed of privacy activity over the last several years. Facebook is ripe for better information displays to explain the numerous privacy options on the site. I have begun to investigate social networks, specifically looking at "friend grouping," such as Facebook Lists or Google+ Circles [4]. Group-based permissions remain at the core of fine-grained privacy control in most social

networks. In order to use these fine-grained controls, users must be able to accurately and usefully specify friend groups. This work explored how the interface design impacted the groups users made, how and why people might choose to group their online contacts, and the resistance users have to trusting these groups. This study was recently awarded an Honorable Mention at INTERACT 2011. I have also investigated on Twitter, working to document private tweets that are publicly leaked [5] and exploring the types of personal regret that occur on Twitter [6].

As both status update systems and social networking sites proliferate across the Internet, users must form a mental model for these new modes of content sharing, and this is where our continued work in privacy information displays can benefit users. I hope to continue this line of work and provide users with tools that help them understand their sharing behavior and act in a more privacy-protecting manner [6].

General Design Principles

The design principles identified in my work on Facebook grouping, online privacy policies, and android smartphones should be able to be generalized across further privacy and security domains. To do this, I will continue refining the iterative methods I have used in specific domains and begin testing specific design principles in isolation to measure the impact of each. For example, creating controlled experiments where standardization, or simplified design are used in isolation. As a broader goal of my work, one that I expect will span many years, I hope to describe, test, and refine a series of design principles for privacy interfaces that help consumers better understand data practices and take more active control of their information, and compel them to behave in a more privacy-protecting manner. The design principles I intend to explore include simplified design, standardization, explanation, automation, nudging, and holistic views (see below).

Proposed Privacy Design Principles

1. *Standardization* – Standardized terms, layouts, interface design patterns, and user options will be used to simplify and clarify both the information presented and the methods for users to interact with the interfaces we present.

2. *Extended explanation* – Where possible definitions, additional explanations, and potential outcomes/ impacts will be detailed so users seeking a deeper understanding can learn more about the terms and concepts used.

3. *Decision removal* – Where possible, taking decisions away from users (or moving them to advanced setting screens) will simplify the choices users must make, and focus their attention on the most important tasks. This is a necessity in Cranor's Human in the Loop model.

4. *Automation* – Repetitive tasks should be made easier through helping users quickly learn how to take advantage of an interface, and if applicable allowing an automated computer system to learn and repeat a user's preferred behaviors.

5. *Interface nudging* – Where a preferred behavior is recognized, the interface will leverage graphic design principles to make this action more likely (e.g., increased size, emphasized text, color, prominent placement, etc.).

6. *Holistic views* – As Reeder explored in his work on access control, holistic views are essential to provide a big picture view of a complex system where users are changing single, often unrelated, settings in isolation. Providing high-level, summarizing overviews helps users deal with complex systems.

7. *Simplified design* – Good communication/information design principles should always be applied: simplified interfaces, removed clutter (text and graphic), high level overviews (with additional screens for advanced users), clear labels, clear demarcations between sections, few colors, few text styles, and repetition.

Design Principles in Additional Domains

I would like to apply these design principles in other fields. Much of the original background for the privacy label came from non-privacy, non-online labels such as nutritional and pharmaceutical information. I believe that many of the principles I have since explored can also be applied to other domains. Specifically, I have interest in exploring the application of these principles in educational materials and online journalism. I have been involved with the CUPS anti-phishing project, which is now led by a spin-off company, Wombat Security Technologies, with which I have worked as a consultant. I would leverage my work developing games and other educational materials, with the addition of the above principles. Also, for the entirety of my Carnegie Mellon career I have been involved in various roles with the university newspaper, *The Tartan*, and before that worked at RIT's weekly publication, *Reporter*. I devoutly care about the future of journalism and how the digital era is, quite quickly, changing every aspect of the journalistic process and publication. I believe these same information design principles could be automatically applied to the deluge of articles that fill the Google News collection.

User Controllable Privacy/Policy Learning

Finally, while I believe that consumer education is severely lacking in online privacy, I am interested in assisting users in managing some of these decisions through automation. With my colleagues, I developed the concept of User Controllable Privacy/Policy Learning (UCPL). UCPL is a replacement for standard, black-box machine learning, in which the user has direct access into the ML model. As data is added, and features and outcomes are updated by a given algorithm, the user always has direct oversight and can modify the model directly. While I began to pursue this idea in the location sharing domain [8], I am interested in further applications of user-controlled machine learning.

REFERENCES

[1] Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. **Patrick Gage Kelley,** Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. CHI 2010.

[2] Design of A Privacy Label. **Patrick Gage Kelley,** Joanna Bresee, Robert W. Reeder, and Lorrie Faith Cranor. SOUPS 2009.

[3] A Conundrum of Permissions: Installing Applications on an Android Smartphone. **Patrick Gage Kelley,** Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. Under review.

[4] An Investigation into Facebook Friend Grouping. **Patrick Gage Kelley,** Robin Brewer, Yael Mayer, Lorrie Faith Cranor, and Norman Sadeh. INTERACT 2011.

[5] RT @IWantPrivacy: Widespread Violation of Privacy Settings in the Twitter Social Network. Brendan Meeder, Jennifer Tam, **Patrick Gage Kelley,** and Lorrie Faith Cranor. W2SP 2010.

[6] Feelings, Relationships, Sex and Alcohol: Understanding Regrets on Twitter. **Patrick Gage Kelley,** Manya Sleeper, Justin Cranshaw, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh. Under review.

[7] Nudging Users Towards Privacy on Mobile Devices. Rebecca Balebako, Pedro Leon, Hazim Almuhimedi, **Patrick Gage Kelley,** Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. CHI PINC 2011.

[8] User Controllable Learning of Security and Privacy Policies. **Patrick Gage Kelley,** Paul Hankes Drielsma, Norman Sadeh, and Lorrie Faith Cranor. AISec 2008.