A Tool for Policy Comprehension

Alessandra Mazzia
Computer Science & Engineering
University of Michigan
2260 Hayward Ave.
Ann Arbor, MI 48109 USA
amazzia@umich.edu

Kristen LeFevre
Computer Science & Engineering
University of Michigan
2260 Hayward Ave.
Ann Arbor, MI 48109 USA
klefevre@umich.edu

Eytan Adar
School of Information
University of Michigan
105 South State St.
Ann Arbor, MI 48109 USA
eadar@umich.edu

This work was supported by NSF grant CNS-1017149.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.
ACM 978-1-4503-0268-5/11/05.

Abstract
The problem of privacy *policy comprehension* in social network systems (SNS) is exacerbated by ever changing policies and system features, making it difficult for users to reconcile their mental model with the system state. In this paper, we describe the PViz prototype, a tool designed for community-oriented comprehension of SNS privacy settings.

Keywords
Privacy, Social Networks, Visualization

ACM Classification Keywords
H5.m. Information interfaces and presentation

General Terms
Human Factors, Security

Introduction
Though online social networking systems have existed for many years, the changing features of these systems, coupled with mass adoption, have exacerbated the problem of privacy management. These rapid changes have created a situation in which boundary regulation may be difficult to achieve, and it is often difficult to construct an accurate mental model of who can access what. Tools for managing privacy settings frequently couple *control* (specifying who can access what) with *awareness* and *comprehension* (understanding who can access what, given existing settings). However, they do not necessarily account for the types of "queries" users would like to make to

reconcile their mental model of the system state (or desired state) with the policy defaults of the system, the limitations of the system's privacy management features, and individually-enacted settings.

Currently, sites like Facebook allow users to specify fine-grained policies controlling the visibility of their personal data. For example, Facebook's "Custom Settings" page allows users to specify which pieces of profile data (e.g., *Political Views* or *Status Updates*) are visible and to whom. Unfortunately, studies and experience have consistently shown that average users struggle to create, evaluate, and maintain such policies [1,5]. In this paper, we consider the *policy comprehension* problem. Our goal is to assist the user in understanding the visibility of her data in a natural way. Recent work has observed that users' mental models of privacy and visibility in social networks often involve natural communities, or social circles, within their local networks of friends [3,4]. However, existing policy comprehension tools, such as the *Audience View* [5], which allows the user to view her profile as it appears to each of her friends, are not naturally aligned with this mental model.

In this work, we further draw a distinction between *single tasks*, in which the user seeks to understand whether a data item is visible to a specific friend, and *group tasks*, in which the user would like to determine whether a data item is visible to a natural sub-group of friends.

Example 1. Consider Margaret, a user who is evaluating her privacy settings on a popular social network. Margaret would like to stay in touch with John Self, a high school boyfriend. Margaret was also a member of the cross-country team and has remained friends with a number of ex-teammates.

In this example, a single task would be to identify Margaret's phone number visibility to John. Notice that single tasks are easily resolved using an "audience view" interface; Margaret can simply view her profile as it appears to John. In contrast, group tasks prove more challenging. An example of a group task would be for Margaret to ask whether her phone number is visible to all of her cross country friends. To answer this question using the audience view is more difficult; it requires Margaret to enumerate every member of the cross country team, and to view her profile as it appears to each of them.

To address the policy comprehension problem, we have designed and built a tool, called PViz, which corresponds more directly with users' mental models of privacy. PViz allows the user to understand the visibility of her profile at multiple levels of granularity, and according to natural subgroupings of friends.

PViz Overview

The PViz policy comprehension tool is centered on a graphical display, which shows the user's social network. Each node in the display represents a semantically meaningful sub-group of the active user's friends (a community) or an individual friend. Figure 1(a) shows a screenshot of PViz displaying Margaret's social network. Inspecting the display shows that PViz has found five main communities of friends.

To quickly interpret privacy settings in PViz, the user simply looks at the color of the node representing the community of interest. Node color is assigned based on which profile item is currently selected and the user's privacy settings. For a community, the number of friends for whom the profile item is visible is encoded as a gradient from 0% (light) and 100% (dark) as well as an explicit numerical cue when hovering the mouse.

To the left of the graphical display, PViz shows a list of profile items for which the user can configure privacy settings. To view privacy settings for a specific item, the user must select the item from the list.

PViz includes a search tool, which facilitates the completion of single tasks. The user may enter a friend's name into the search box, and PViz will respond by centering the graphical display on and highlighting the node containing that friend.

PViz includes a text box that displays the names of all members of the currently selected node (community).

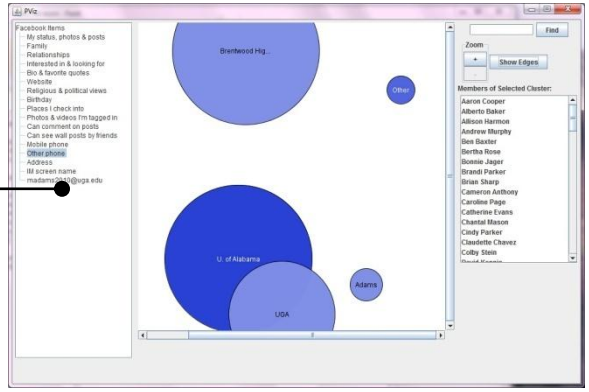

(a). Coarse granularity view

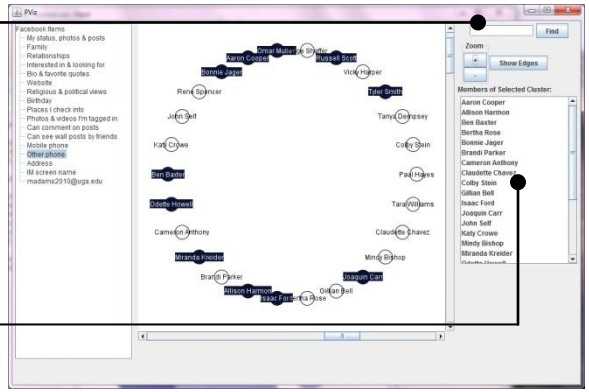

(b). Fine granularity view

figure 1. PViz allows the user to understand privacy settings at different levels of granularity.

In Figure 1a, notice that the node labeled "U. of Alabama" is darker than the node labeled "UGA," indicating that a larger percentage of friends in the "U. of Alabama" community can see Margaret's "Other Phone" than in the "UGA" community.

By zooming in on a community PViz reveals constituent sub-communities by visually breaking large "community" nodes into sub-communities at different zoom levels. In addition to the graphical display, PViz also provides the user with several ways of interacting with the social network graph (Figure 1).

Consider again the single and group tasks from Example 1 which are both easily completed using PViz. To check whether her phone number is visible to John Self (single task), Margaret first selects the "Phone" profile item, and then uses the search box to find the node containing John. At the individual level (Figure 1b), Margaret finds that the node representing John is white, indicating that John cannot see her phone number. To check whether her phone number is visible to her cross country friends (group task), Margaret starts at the coarsest level of granularity. She recognizes that her high school cross country friends should be a subset of her high school friends, so she zooms in on the node labeled "Brentwood High School," which reveals a node labeled "BHS Cross Country." To ensure that the node contains the appropriate friends, Margaret may select the node, and inspect the list of friends who belong to the community. After locating the node of interest, Margaret can interpret her privacy settings for that community. In this example, the "BHS Cross Country" node indicates that only a portion of this group (42%) can see her phone number.

Implementation and Technical Challenges

We have implemented a prototype of PViz in the context of Facebook. To partition social network graphs into communities, we applied a multi-level variation of a common community-finding algorithm [7]. However, other community-finding algorithms, as well as explicit user groupings, can be integrated into PViz.

A key technical challenge we noted in creating PViz was the importance of proper labels for the hierarchically-structured data visualization that would support quick targeting of relevant (sub)communities. We identified two main criteria for a good label: (1) It should be concise, and (2) It should uniquely distinguish the users in the community from others in the network. At present, we are experimenting with techniques from Information Retrieval (finding "tags" based on profile information that maximize a combination of precision and recall for communities), and inferring a propositional rule that distinguishes members of the community from other users in the network based on a logical combination of tags.

PViz User Study

To test the effectiveness of PViz we are in the process of conducting a user study comparing it to two alternative tools, Facebook's Custom Settings Page (CS) and Facebook's Audience View (AV), which are representative of the state of the art in comprehension tools for fine-grained social network privacy policies.

While an obvious study design would have users answering questions (single and group tasks) about their own profiles, most potential subjects have not configured their Facebook privacy settings (e.g. 25% of households with a Facebook account either did not use or were not aware of Facebook's privacy settings [2]). To deal with this, we have constructed a simulated Facebook environment (same interface design but limited functions) that utilizes an artificial, yet realistic, standard environment in which to conduct the study.

Early results indicate that users are able to complete more tasks, and complete them more accurately, than existing tools. Despite the novelty of the interface, for those tasks possible through present tools, PViz does not increase the completion time [6].

Conclusions and Future Work

As average users of social network sites continue to struggle with the creation and maintenance of effective privacy policies, new tools are necessary to help individuals reconcile individual privacy mental models with system state. In this paper, we presented a tool for policy comprehension that corresponds more naturally to users' mental models of privacy.

We are currently evaluating the system and integrating new features to expand functionality beyond comprehension to the effective control of privacy settings directly from the interface.

Bibliography

[1] Acquisti, A., and Gross, R. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Privacy Enhancing Technologies Workshop* (2006).

[2] Consumer Reports. Social insecurity: What millions of online users don't know can hurt them. http://www.consumerreports.org/cro/magazine-archive/2010/june/electronics-computers/social-insecurity/overview/index.htm.

[3] Fang, L., and LeFevre, K. Privacy wizards for social networking sites. In *Proc. WWW* (2010).

[4] Jones, S., and O'Neill, E. Feasibility of structural network clustering for group-based privacy control in social networks. In *Proc. SOUPS* (2010).

[5] Lipford, H., Besmer, A., and Watson, J. Understanding privacy settings in Facebook with an audience view. In *Proc. of the 1st Conference on Usability, Psychology, and Security* (2008).

[6] Mazzia, A., LeFevre, K., and Adar, E. The PViz comprehension tool for social network privacy settings. University of Michigan CSE Technical Report CSE-TR-570-11 (2011).

[7] Newman, M., and Girvan, M. Finding and evaluating community structure in networks. *Physical Review* 69(2) (2004).