

Visual vs. Compact: A Comparison of Privacy Policy Interfaces

Heather Richter Lipford,
Jason Watson, Michael Whitney
 Software and Information Systems
 University of North Carolina, Charlotte
 Charlotte, NC 28223
 {Heather.Lipford, jwatso8,
 mwhitne6}@uncg.edu

Katherine Froiland
 Computer Science Department
 University of Minnesota, Morris
 600 East 4th St.
 Morris, MN 56267
 froi0019@morris.umn.edu

Robert W. Reeder
 Microsoft
 One Microsoft Way
 Redmond, WA 98052
 roreeder@microsoft.com

ABSTRACT

In this paper, we compare the impact of two different privacy policy representations – AudienceView and Expandable Grids – on users modifying privacy policies for a social network site. Despite the very different interfaces, there were very few differences in user performance. However, users had clear, and different, preferences and acknowledged the tradeoffs between the two representations. Our results imply that while either interface would be a usable option for policy settings, a combination may appeal to a wider audience and offer the best of both worlds.

Author Keywords

Privacy, access control policy, social network site.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Design, Human Factors, Security.

INTRODUCTION

Many applications and websites allow users to manage a large amount of information such as files and photos, friends lists, and detailed interests and activities. Users are charged with maintaining their privacy and security policies that govern how their information is shared with other users or organizations. Despite research advances in access control mechanisms and languages, relatively little research

has examined interface mechanisms that help users create and modify such policies.

Many access control mechanisms represent a policy as a set of rules governing the permissions that various groups or roles are granted to various data. Some interfaces simplify the rules into sets of checkboxes or menus, such as choosing whether certain data items are public or private. A number of studies have demonstrated that users struggle to understand and properly manage privacy and security policies [2, 4] and need simpler mechanisms and more usable interfaces.

There are several notable investigations into usable access control interfaces that explore different policy representations. Expandable Grids is a general, matrix-based visualization of a policy showing the effective combination of policy rules [1]. AudienceView represents a policy as the different views of information as seen by various audiences, or groups of users [4]. In addition, the SPARCLE workbench allows users to construct policies using a natural language interface [1]. Each representation has been shown to be an improvement over an existing rule or checkbox style interface. However, each was examined in different domains, and no one has examined the tradeoffs between these representations.

In this paper, we seek to compare two of these representations, namely Expandable Grids and AudienceView. For this comparison, we chose the domain of privacy policies on social network sites, particularly Facebook. These access control policies are available to hundreds of millions of users who maintain profiles filled with personal information.

EXPANDABLE GRIDS

The Expandable Grids interface, shown in Figure 1, was created as a general method for representing and modifying access control policies [1]. Expandable Grids show precisely what a policy allows or does not allow in a matrix with hierarchical axes that can be expanded or contracted to show more or less policy detail. For this study, the set of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.

Copyright 2010 ACM 978-1-60558-929-9/10/04...\$10.00.

principals, shown along the top axis, were the groups of people that Facebook allows for privacy policies: “All My Friends”, “Friends of Friends”, and “All My Networks” categories. Friends had subgroups of “Best Friends,” “Family,” and “Shady Friends” while specific networks were the institution and city.

The set of resources is shown along the left. These are the categories and individual information fields provided in Facebook, such as “Basic Information” which includes fields such as “Birthday” and “Hometown.” In this domain, permissions are simple access – allow or deny. Thus, a red box indicates that access is denied, a green allowed, and yellow that the category has a mixture of allow and deny permissions. The user clicks on the box to set the permission for the particular information and user group. When setting a permission on a category of resources or principals, the rule is then applied to all of the individual pieces of information or people in that group.




Figure 1. The Expandable Grids interface.

AUDIENCE VIEW

AudienceView was designed to directly incorporate the notion of audience into a privacy settings interface. A privacy policy is represented as the view that a particular audience has of the information [4]. The prototype, shown in Figure 2, was implemented to resemble the information and look of Facebook profiles. Each page shows all of the information for the profile available to a particular audience, with users changing audiences using the tabs at the top of the interface.

Users click on lock buttons to show or hide the information on each page. For protected information, the field or category is hidden and the field title is grayed out. Similar to the semantics of Expandable Grids, when individual audiences have different settings, the lock button on the “All My Friends” and “All My Networks” audience is shown in yellow to signify a mixed setting.

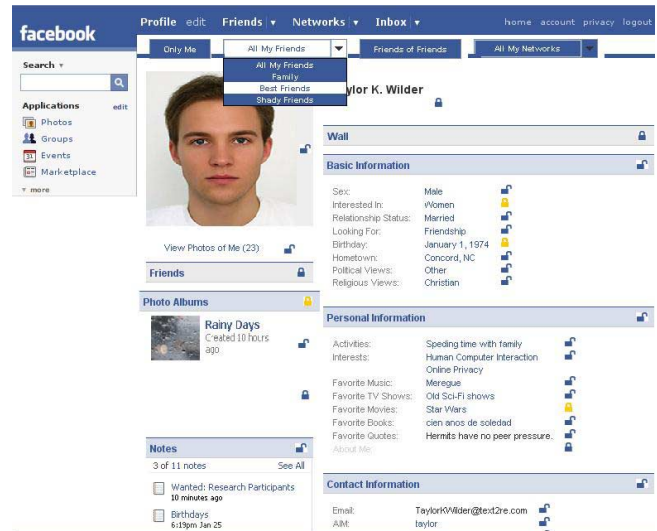


Figure 2. The AudienceView interface.

A COMPARISON

Both interfaces have previously been positively evaluated against existing policy interfaces in different domains [1, 4]. The grid visualization of Expandable Grids is more general, able to represent any kind of resource and principal, with multiple levels of hierarchy, for a range of permissions. The grid can provide an overview of all settings at once, with the ability to drill down to more specific ones. In comparison, AudienceView is more limited because it is strongly tied to the visual representation of a set of information, making the interface more difficult to reuse across domains. The metaphor may also be more difficult to convey for permissions beyond simple access control. The representation inherently limits the number of potential audience groups, as too many would be difficult to display and interact with. These differences would obviously influence in which domains either interface could be used.

From a user’s point of view, Expandable Grids is more compact, which may make it easier and faster to find and modify settings. AudienceView’s settings, on the other hand, are spread across many pages. Yet, the interface provides a more concrete context by showing the specific information that is being shared, which may help users better understand the implications of a policy. Thus, the purpose of this work is to determine how these differences impact users’ performance and preferences in understanding and modifying an access control policy.

For this comparison, we chose the domain of profile privacy settings on social network sites, where the two interfaces could present nearly identical functionality.

STUDY

We designed a within-subjects study to compare the tradeoffs between the AudienceView and Expandable Grids policy representations. We recruited participants by handing out flyers on campus and through word of mouth

on Facebook. We first gathered basic demographic information. We briefly explained the functionality of one interface and gave participants unlimited time to play with an unconfigured interface. Next, a preconfigured interface was opened, and users were instructed to complete a set of tasks as described below. Participants then answered a set of Likert-scale usability questions about the interface. The same process was then repeated for the second counterbalanced interface. Finally, participants were interviewed about which interface they preferred and why, and what they liked and disliked about each prototype.

Users were asked to complete 17 individual tasks for each interface. After each, the participants were then asked to rate their confidence in their actions or responses on a scale of 1 (very unconfident) to 7 (very confident).

The first four tasks asked users to read and understand the existing policy. For example, "Which of the following friend groups can see your relationship status?" These tasks further introduced the user to the interface and gauged their understanding of existing settings.

The second four tasks were simple configurations of a single item or category and a single audience, such as "Deny your family from viewing your photo albums." The next four tasks were more complex configurations involving multiple items for a single audience. For example, "Allow your best friends to view your email, mobile phone number, and work information." The next four tasks involved a single item configured for multiple audiences.

The final task asked participants to pretend the displayed profile was their own and to take as much time as needed to configure all of the privacy settings as they deemed appropriate. For this task only, we asked them to think aloud.

We attempted to balance the audiences and information fields to be configured across the tasks. None of the tasks depended on previous tasks for successful completion. We used usability software to record screen video, and a custom application to log timing and confidence.

RESULTS

We recruited 23 participants. 16 were ages 18-24, four 25-34, one 35-44, and two over 55, 12 males and 11 females. The background of the participants widely varied. We classified 8 participants as novice users, with less than 6 months of Facebook experience (and 1 had never used Facebook). Ten participants reported frequently modifying their privacy settings on Facebook, while 3 reported never modifying. Unfortunately, the video for 2 participants was corrupted and we were only able to analyze and report the confidence and timing data for those 2 participants.

Accuracy and Confidence

There were no differences between the accuracy rates for the two interfaces. For AudienceView, over all participants

($n=21$) only 13 tasks (4%) were performed incorrectly. Eleven errors were caused as participants could not find particular data items, mostly during the first 4 tasks. In Expandable Grids, there were 18 (5.7%) incorrect responses. Twelve of which were while modifying privacy settings and 4 were performed by a single user.

While the difference in errors is not significant, the errors made reveal tradeoffs between the two interfaces. In AudienceView, users initially appeared to have a harder time finding particular pieces of information. This was easier in Expandable Grids. However, participants did occasionally mis-click on the wrong box in the grid. This could either be a motor error, or an error in reading the proper row or column. In Expandable Grids, mis-clicks could be more difficult to recover from. For example, if a user allowed or denied a category of information or group of users, that rule propagated down to all sub-fields or sub-groups. If that action was not intended, the user must then re-do all the rules for the sub-fields to recover. This did occur several times, and for one participant we even reloaded the pre-configured settings to recover. An undo function, however, would solve this problem. AudienceView was not as prone to this issue, although users could still mistakenly navigate to the wrong audience page or forget which audience page they were configuring.

For all tasks, confidence was reported as 6.48 (out of 7) for AudienceView, and 6.51 for Expandable Grids. Using Wilcoxon signed rank tests, we found no evidence of a difference between the two interfaces. Overall, participants were highly confident interacting with both interfaces. This makes sense given the high rate of accuracy.

Timing

We performed four two-way repeated measures (2 X 4) ANOVAs with each grouping of similar tasks and interface as the two within-subjects factors. For each of these tests, we performed a Tukey HSD post hoc analysis and focused only on the measures between interfaces for the same task.

For the first set of tasks (tasks 1-4), the ANOVA ($F(3,20) = 22.358, p < .001$) and subsequent post hoc Tukey HSD test revealed a significant timing difference ($p < .05$) between interfaces for task 1. In Task 1, AudienceView was slower than Expandable Grids (XG=51.61s, AV=94.70s). We believe this is because participants appeared to need time to explore the interface at first and find particular data items. However, they sped up over time. Our analysis for tasks 5-8 and tasks 9-12 revealed no differences between interfaces.

For tasks 13-16 ($F(3,20) = 8.502, p = .001$), Expandable Grids was significantly faster for tasks 15 (XG=46.35s, AV=86.83s) and 16 (XG=36.13s, AV=60.57s), which involved configuring an item for 4 or 5 different audiences. What was surprising is that there were no differences for tasks 13 and 14, which involved fewer audiences. Thus, Expandable Grids was faster as expected for configuring multiple audiences, but only for more than 3.

Usability and Preference

Fifteen participants reported preferring AudienceView, while six preferred Expandable Grids. There were no differences in the responses to the 7 usability questions between the interfaces in general. However, not surprisingly, users rated the interface that they preferred higher on each of the usability questions.

Participants liked the visual feedback in AudienceView and felt that the interface provided a more accurate depiction of their information. *“I liked the visual aspect. It made it easier to know what you are configuring without having to go back and look”* (P12). However, they did not like having to visit so many pages to manipulate all of the privacy settings: *“It was time consuming going back to check for each group”* (P13). With Expandable Grids, users liked being able to see the entire policy at once: *“it was all on one screen”* (P13), and *“easier to see the bigger picture”* (P20). We did notice two potentially interesting findings that may warrant further investigation. The 3 users over the age of 35 all preferred Expandable Grids. Additionally, 7 out of 8 users who reported frequently modifying privacy settings on Facebook preferred AudienceView.

User Behavior

Finally, we compared the behavior and resulting settings for the final task where participants were asked to create an entire policy as though it were their own profile. Once again we found few differences. Users generally took between 4 and 5 minutes to modify all of the settings. As users got tired of the task, they frequently modified only categories of settings, such as “All My Networks” and made fewer adjustments to the more detailed friend groups or data fields, to speed up completion of the task. With Expandable Grids, several users took the strategy of first setting the entire policy to either deny or allow (which can be set using one box), and then adjusted the policy from there. This did not appear to result in any efficiencies however. This capability was not available on AudienceView.

We wondered whether the added context shown in AudienceView would change the privacy policy that was configured. While we did not find quantitative differences, one participant did comment on the benefits of the context: *“Some of that stuff on [Facebook], you know, people can steal that. And when I was doing it on [Expandable Grids], I would have probably given some of that information out... So when I looked at that on [Audience View], then I realized the importance of it, ‘cause I could visually see what I was doing”* (P1). Other participants also indicated that AudienceView did improve their confidence in the resulting policy configuration. As P13 stated, with AudienceView *“I was completely confident in what I was configuring. It was right there in front of me.”* Whereas P9 indicated that Expandable Grids *“made me nervous to click on the beginning of a row. It made me question, did I do it right?”*

CONCLUSION AND IMPLICATIONS

Overall, both interfaces were highly usable. Considering the very different representations presented in the two interfaces, we were very surprised by the general lack of performance differences. In particular, we expected that Expandable Grids would be faster for more tasks, and that the added context in AudienceView may lead to higher confidence. Yet these differences did not materialize. However, users did have clear, and different, preferences. The participants in the study acknowledged the same advantages and disadvantages we had identified with each interface. Namely, they liked the visual feedback of AudienceView, but not all the page visits. They also liked the compact overview available with Expandable Grids with all settings in one location, even though it lacked the visual feedback.

Our results indicate that either interface would be a usable option for similar privacy policies, such as those on other social network and personal information sharing sites. However, different representations may appeal to different users. Many participants suggested combining these two interfaces to provide the best of both worlds. For example, one participant commented that he would use Expandable Grids to first modify all his settings, and then use AudienceView to view and tweak them. Thus, users may benefit from a combination of both representations, to provide both a clear and concise overview of a policy, while still allowing for detailed visual feedback.

More generally, our results indicate that users do value a contextual representation and a compact representation for privacy policies. Designers and researchers of policy interfaces should consider how to achieve both contextual and compact visual representations, and investigate how to balance the tradeoffs and preferences between such representations to further improve usability, and the resulting privacy and security of information.

REFERENCES

1. Brodie, C.A., Karat, C., and Karat, J. An empirical study of natural language parsing of privacy policy rules using the SPARCLE policy workbench. In the Proceedings of SOUPS 2006, pp. 8-19.
2. Cao, X., and Iverson, L. Intentional access management: making access control usable for end-users. In the Proceedings of SOUPS 2006, pp 20-31.
3. Reeder, R.W., Bauer, L., Cranor, L.F., Reiter, M.K., Bacon, K., How, K., and Strong, H. Expandable grids for visualizing and authoring computer security policies. In the Proceeding of CHI 2008, pp. 1473-1482.
4. Smetters, D.K., and Good, N. How users use access control. In the Proceedings of SOUPS 2009.
5. Watson, J., Whitney, M., and Lipford, H.R. Configuring Audience-Oriented Privacy Policies. In the Proceedings of the Workshop on Assurable and Usable Security Configuration, ACM CCS, November 2009.