Oops, I Did It Again:

Mitigating Repeated Access Control Errors on Facebook

Serge Egelman National Institute of Standards and Technology Gaithersburg, MD serge.egelman@nist.gov Andrew Oates Google, Inc. Cambridge, MA andrewoates@gmail.com Shriram Krishnamurthi Brown University Providence, RI sk@cs.brown.edu

ABSTRACT

We performed a study of Facebook users to examine how they coped with limitations of the Facebook privacy settings interface. Students graduating and joining the workforce create significant problems for all but the most basic privacy settings on social networking websites. We therefore created realistic scenarios exploiting work/play boundaries that required users to specify access control policies that were impossible due to various limitations. We examined whether users were aware of these problems without being prompted, and once given feedback, what their coping strategies were. Overall, we found that simply alerting participants to potential errors was ineffective, but when choices were also presented, participants introduced significantly fewer errors. Based on our findings, we designed a privacy settings interface based on Venn diagrams, which we validated with a usability study. We conclude that this interface may be more effective than the current privacy settings interface.

Author Keywords

Privacy, Social Networks, Access Control, Policy Authoring

ACM Classification Keywords

D.4.6 Security and protection, H.5.2 User Interfaces

General Terms

Security, Human Factors, Experimentation, Design

INTRODUCTION

Once the realm of system administrators, even novice users are now expected to author access control policies. When using social networking websites, users may configure their profiles so that only certain people can view private information. As users share content, erroneous access control policies can have profound privacy consequences, especially when there are disconnects between the policy author's mental model of what the policy *should* say and how the system *actually* evaluates the policy. These "semantic errors" are

CHI 2011, May 7-12, 2011, Vancouver, BC, Canada.

much harder to detect than syntactic errors because they result in technically valid policies. Research suggests that policy authors are least likely to create semantic policy errors when they use natural language to specify the policy [3]. However, due to various factors, allowing policy authors to use natural language whenever they specify access control policies is simply not practical. While access control mechanisms have proliferated over the course of forty years, research on the interfaces that a human uses to author policies has only just begun. The Computing Research Association (CRA) has gone so far as to list giving "end-users security controls they can understand and privacy they can control for the dynamic, pervasive computing environments of the future" as one of the top four grand challenges in information security and assurance [4].

We present a study of how Facebook¹ users create, detect, and resolve semantic access control errors. The current interface creates opportunities for policy ambiguities—situations where a policy could be interpreted in different ways based on a parser's predefined semantics. When left undetected, policy ambiguities can transform into semantic errors. We contribute to the literature on access control interfaces by empirically evaluating the effectiveness of corrective feedback during policy authoring; we show when feedback can and cannot minimize policy errors. We identify rudimentarybut-useful scenarios where it would be impossible to correctly configure Facebook's privacy settings. Finally, we present and evaluate a new interface to help users visualize effective permissions on social networking websites.

We first present related work on privacy and online social networks. Next, we present an experiment where participants altered their privacy settings to fit realistic scenarios that exploited work-play boundaries. We examined whether participants noticed ambiguities and whether they corrected those ambiguities using two different types of feedback. Finally, we designed a privacy settings interface using Venn diagrams so that participants could better visualize how their work networks and social networks overlapped. In our second experiment, participants created significantly lower error rates in half of the scenarios, and equal error rates in the other scenarios. Finally, we discuss participants' Facebook privacy behaviors and the limitations of our experiments.

Copyright 2011 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the U.S. Government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

¹Products are identified to specify the experimental procedure adequately. This is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology.

BACKGROUND

Social networking websites allow users to post personal information to share with their friends, while the users rely on access control settings to prevent the data from being shared with unintended parties. In 2005, Gross and Acquisti mined the Carnegie Mellon network on Facebook. They found that "90.8% of profiles contain an image, 87.8% of users reveal their birth date, 39.9% list a phone number (including 28.8% of profiles that contain a cellphone number), and 50.8% list their current residence [9]." Young and Quan-Haase performed a study in 2009 and found that 64% of study participants had restricted their profiles to friends [22]. Because social networking websites are primarily used to stay in touch with existing friends, rather than to form new online relationships [12, 11], we would expect users to limit the private information they expose to complete strangers. Indeed, boyd and Hargittai found that in 2010, 98% of their study participants had previously modified privacy settings [2].

Strater and Lipford performed a study of 18 Facebook users and found that over 72% of their participants took an "all or nothing" approach to privacy: they made their profiles either completely open or restricted them to only their friends. Only five participants used fine-grained controls to restrict access based on relationships and the type of information that was to be accessed [16].

As online communities evolve, controlling information on social networking websites has become a lot more complicated than restricting access to friends. Binder et al. performed a survey of Facebook users to examine how differing social spheres interact and found that privacy concerns were directly correlated with the number of family members a user had friended [1]. Stutzman and Kramer-Duffield found that some Facebook users went so far as to create a second profile to share with their "real" friends [17]. Skeels and Grudin conducted a survey of corporate Facebook users to examine how they managed privacy settings when social networks consisted of both friends and coworkers. They found that many respondents had so many problems changing privacy settings that many resorted to censoring their content rather than configuring privacy settings. In fact, many respondents complained about not being able to divide their Facebook friends into groups, oblivious to the fact that such a feature existed [14]. Access control concerns are only increasing as social networking tools become more and more accepted within the workplace [19, 6, 13, 5, 15, 21].

Researchers have proposed alternative privacy settings interfaces for social networking websites. For instance, Watson et al. proposed AudienceView, which allows users to configure settings while viewing effective permissions [20]. Tootoonchian et al. proposed Lockr, which is based on access control lists [18]. We are unaware of proposed interfaces designed around the interplay between overlapping social networks. Likewise, tools have been proposed to help users understand policy semantics through change-impact analysis [8], as well as to minimize semantic errors by asking users clarifying questions [7], we are unaware of usability studies that have been performed on these techniques.

FEEDBACK AND GUIDANCE

In our first experiment, we examined how often Facebook users detected semantic errors when specifying access control policies (i.e., privacy settings). Additionally, we examined how they went about correcting these errors and whether certain types of feedback would help them specify the correct policies. We performed a laboratory study on 33 participants who performed tasks on Facebook that were likely to lead to semantic errors in their access control policies.

Methodology

From May 26 to June 1, 2010, we conducted a laboratory experiment where we recruited Facebook users in the Brown University community. Since over 25% of Facebook users are in the 18-24 age group [10], we believe our sample is representative of an important demographic of Facebook users.

We recruited participants using flyers posted around campus as well as ads on Facebook. We directed potential participants to an online screening survey designed to establish Facebook usage and their affiliation with Brown. So as to not prime participants, we asked several subterfuge questions about their usage of several other social networking websites. Our experimental tasks required participants to have an existing *brown.edu* email address so that they could join the Brown network on Facebook if they were not already members. We invited only those participants who had been using Facebook for six months or more to schedule a time to visit our laboratory on campus. Because we wanted this study to be realistic with regard to how Facebook users might behave outside of the laboratory, at no time did we instruct participants on using Facebook's privacy settings.

When participants arrived at our laboratory, we first asked them questions about their Facebook usage. This survey contained questions about how often they posted various types of content (e.g., status updates, photos, links, etc.) and the size of their social networks. Next, participants performed four different tasks using their real Facebook profiles. They were instructed to role-play the following scenario:

You graduated from Brown and are now working at a local startup called "Conglomi." We will ask you to perform tasks using Facebook that you might do while you are at work. As a first step, the experimenter will help you install a Facebook app in order to help us collect data during the experiment. No personal data will be used and the app will be removed at the end of this study. This app will add you to the Conglomi network on Facebook, so that you can interact with coworkers. In addition to this, the app will also send you friend invitations from a few of your coworkers.

Next, we gave participants a handout listing their fictitious coworkers and indicating which Facebook networks each coworker was a member of, as well as which of the coworkers were friends with them on Facebook. The Facebook app that participants installed was used solely for data collection purposes. Unbeknownst to participants, their connections were being intercepted by our proxy server so that we could

| | Album | Custom Privacy |
|--------------------------------------------|--------------------|----------------------------------------------------------------------------------------|
| Album Name: | Locatio Descrip | |
| Location: | | These people: Only Friends |
| Privacy: | Privacy | Only friends and selected networks can see this. And these networks: Brown Conglomi |
| Friends and Networks Friends of Friends | | X Hide this from |
| Only Friends 1 Customize | | These people: |
| | l | B Save Setting Cancel |

Figure 1. The default Facebook privacy settings interface. Users can use the interface on the left to share content with all Facebook users, only friends, friends and networks, or friends of friends (A). If users choose to "customize" their privacy settings, the interface on the right appears (B).

automatically add the coworkers and the Conglomi network to their profiles without actually modifying any data stored on Facebook. We also used the proxy to alter the privacy settings interface based on our experimental conditions.

Once participants added the app, we instructed them to check their email to confirm their friend requests. Five friend requests were generated by our software and appeared to come from Facebook, although in reality participants' friend lists and networks were never altered. After this, we gave participants a set of four tasks to complete using Facebook:

- 1. **Tag A Photo** Browse through the photos that either you or one of your friends has uploaded to Facebook. Tag someone in one of these photos.
- 2. **Party Scenario** Last Friday all of your coworkers stayed at work until midnight because of a deadline. However, some of your friends from Brown University were throwing a party at the same time. You told your coworkers that you were feeling sick so that you could leave work early in order to attend the party. You could be fired if anyone at work finds out that you were at the party. One of your friends from the party has just asked you to upload your photos. Please do so, while being aware of your concerns. You want to share the photos with as many friends and other people who may have attended the party, but with none of your coworkers.
- Comment on A Photo— Browse through the photos that either you or one of your friends has uploaded to Facebook. Leave a comment on any one photo.
- 4. **Recruitment Scenario** As a small Providence-based business, your manager understands that its important to recruit Brown alumni. You have recently been put in charge of coordinating a recruiting event with other Brown alumni who also work at Conglomi. At the last planning meeting, you took several photos that you want to share with the other planners. However, these photos should only be

shared with the other planners (i.e., Conglomi employees who are Brown alumni). If these photos were disclosed outside of the planning group, you could get fired.

The tagging and commenting tasks were chosen solely as a buffer between the photo uploading tasks. In the photo uploading tasks, we provided participants with photos in a directory on the laboratory computer. These two tasks were intentionally designed to exploit ambiguities in Facebook's existing privacy settings interface (Figure 1). We randomized the order in which participants performed these two tasks.

We asked participants to allow friends and classmates to view the party photos, but to deny access to users in the Conglomi network, regardless of their potential Brown network membership or friend status. We created this task to examine participants' reactions to an ambiguity in the Facebook privacy settings interface: Facebook defaults to allowing rather than denying access. Specifically, when a network is not selected for sharing, members of that network may still be granted access due to other network memberships. For instance, in Figure 1B, while the Brown and Conglomi networks are not selected, members of these networks who are friends with the user will be granted access.

As written, the party scenario was impossible to complete using Facebook's existing interface because an entire network cannot be explicitly denied access. While Facebook allows the user to deny access to a specific list of people (Figure 1B), the user can only name existing friends (i.e., a user cannot enter non-friends or entire networks in this box). If a study participant granted access to friends plus the entire Brown network, there would be no way of denying access to non-friends in the Brown network. Thus, we considered it "correct" if participants only granted access to friends, and then explicitly hid the photos from the friends who were also affiliated with Conglomi. While these photos would be hidden from members of the Brown network who were not friends with the participant, this would be the only way of guaranteeing the photos would be hidden from all members of the Conglomi network. In order to do this, participants had to perform the following steps:

- 1. Select "Customize" (Figure 1A).
- 2. Leave default of "Only Friends" (Figure 1B).
- 3. Do not select any networks (Figure 1B).
- 4. Name the five coworkers to hide from (Figure 1B).

In the recruitment scenario, we asked participants to grant access to only those people in both the Brown and Conglomi networks—the intersection. Facebook's existing interface may also be ambiguous to some users because while they can select multiple networks to which access should be granted, Facebook will parse this as users in *either* network—the union. For instance, in Figure 1B, if both the Brown and Conglomi networks are selected, users may not understand whether this represents the union or the intersection of the networks. We wanted to determine whether this limitation of the interface was clear to participants. For the participants who understood that the task was impossible, we were curious how they would attempt to complete it.

The only way around this limitation is to manually specify the name of every desired individual who should have access. However, Facebook only allows friends to be specified in this manner. Thus, we considered it correct if participants granted access to only the two friends who were members of both the Brown and Conglomi networks. This errs on the side of denying access to Facebook users who are not at the intersection of the Brown and Conglomi networks, but comes at a cost of also denying access to some users who should have access—and do not because they are not friends with the participant. Participants needed to perform the following steps to accomplish this:

- 1. Select "Customize" (Figure 1A).
- 2. Change "Only Friends" to "Specific People" (Figure 1B).
- 3. Name the two coworkers in box that appears (Figure 1B).
- 4. Do not select any networks (Figure 1B).

We consider these two limitations "policy ambiguities," since without having prior knowledge of the system's semantics, the user has no way of knowing how Facebook will parse these policies. We designed two experimental conditions to automatically detect the latter type of ambiguity: the difference between "default allow" and "default deny" policies. That is, when access is granted to one of two networks, did the participant intend to explicitly deny access to members of the second network, or should access be granted solely based on membership in the first network? We decided not to include logic to detect the second ambiguity, the difference between the union and intersection, because we were concerned with overwhelming participants by asking for too much input regarding their policy ambiguities. We included the recruitment scenario to examine whether the automatic detection of *some* ambiguities would prime participants into manually locating additional-more relevant-ambiguities. The study conditions were as follows:

| You se Should users | lected Conglon s in both netwo | ni but not Brow rks be includer | vn d or excluded? | |
|---------------------------------------|----------------------------------------|---------------------------------------------------|--------------------------------------------------|--|
| Michael Tobey Brown Conglomi | Antony Deyesso Brown Conglomi | Caitlyn Anders Friends Brown Conglomi | Emily Portica Friends Brown Conglomi | |
| Do you wish | to continue, or | go back and o | hange your settings? | |

Figure 2. The privacy warnings in the two experimental conditions. In the *Warn* condition (A), participants were warned of an ambiguity without given any guidance. In the *Choose* condition (B), participants were given a choice to help them disambiguate.

| Warning | | |
|----------------------------------------|----------------------|--------|
| Are you sure you want to use the defau | It privacy settings? | |
| | Continue | Cancel |

Figure 3. All participants saw this prompt if they neglected to change the privacy settings from the default of sharing with everyone. Clicking "cancel" allowed them to return to the settings interface and specify a custom access control policy, whereas clicking "continue" confirmed that they wanted to share the content with everyone.

- **Control** Participants used a privacy settings interface provided by our proxy that exactly mimicked the existing Facebook interface (Figure 1).
- **Warn** Participants used the same initial interface as those in the *Control* condition. However, when a policy ambiguity was detected at the time that participants tried to save their changes, a warning was displayed (Figure 2A).
- **Choose** Participants used the same interface as those in the *Warn* condition. However, when an ambiguity was detected, participants were given options to help clarify the ambiguity (Figure 2B). When participants chose to either allow or deny access to the users depicted in the dialog box: their pictures were either superimposed with red Xs or green boxes.

Regardless of condition, we required all participants to alter settings using the "customize" interface in Facebook (Figure 1B). Under the defaults, if a participant neglected to alter the privacy settings, she would grant access to all Facebook users—and we would not yield data on the usability of the privacy settings interface. Thus, we created a prompt so

that if participants attempted to share a photo album without changing the default settings, they were asked if this was their intention (Figure 3). Thus, we collected data on how many participants would have used the defaults to grant access to all of Facebook, as well as how they changed the settings after being explicitly primed to do so.

In the experimental conditions, when an ambiguity was detected, a warning showed a list of Facebook users who would be affected because they were members of both a network that was selected and a network that was unselected.² If the participant did nothing, these users would be granted access. For those in the *Warn* condition, this warning offered participants the option to go back and edit their settings before saving, but did not offer any other guidance. On the other hand, those in the *Choose* condition were explicitly asked whether to allow or deny access to these specific users.

For instance, if a participant selected Brown but left Conglomi unselected, it may be unclear to the participant what permissions should apply to members of both networks. That is, by not selecting Conglomi, did the participant intend to explicitly deny access to members of Conglomi? Or did she intend to grant access to all members of the Brown network regardless of their other networks?

We directed participants to an exit survey after they completed the experiment. We asked them if they had previously modified their Facebook privacy settings, if they had previously encountered scenarios similar to those in this study, how they overcame Facebook's limitations in those scenarios, with whom they intended to share photos during this study, how confident they were that those people had access, and how confident they were that no one else had access.

Analysis

We collected information about our 33 participants' ages, genders, education levels, and nationalities. We found no significant differences between the randomly assigned conditions with regard to demographics. Overall, our youngest participant was 18, while our oldest was 24 ($\mu = 20.39$, $\sigma = 1.54$). Ten of our participants were female, and the remaining 23 were male. Six of our participants were graduates of Brown, while 27 were students. Finally, thirty of our participants were from the US, two were European, and one was Indian.

In the rest of this section we present participants' task performance during the party scenario and the recruiting scenario. We explain how participants coped with Facebook's limitations, how providing automatic feedback helped participants clear up access control ambiguities, and how in certain cases, providing feedback worked to participants' detriment.

Party Scenario

We asked participants to modify permissions so that friends and classmates could view the uploaded photos, so long as these individuals were not members of the Conglomi network. Before participants in the experimental conditions had

| Party Scenario | Control | | Warn | | Choose | |
|------------------|---------|-----|------|-----|--------|-----|
| Ν | 13 | | 11 | | 9 | |
| Correct at start | 5 | 38% | 3 | 27% | 3 | 33% |
| Saw feedback | | | 7 | 64% | 8 | 89% |
| Correct at end | 5 | 38% | 2 | 18% | 8 | 89% |

Table 1. The number of participants who correctly set permissions during the party scenario before any feedback was displayed, who saw feedback in the experimental conditions based on ambiguity detection, and who correctly set permissions after seeing the feedback.

the opportunity to see any feedback regarding policy ambiguities, we observed that a third of participants across all of the conditions (11 of 33) correctly set the permissions. Of the 22 who initially made errors, those errors fell into two categories: "over-sharing" such that unintended users were granted access (17 of 22), and "under-sharing" such that intended users were denied access (5 of 22).

Without feedback, participants erred on the side of data leakage; seventeen participants authored policies that would have resulted in unintentionally sharing with members of the Conglomi network. In all but two cases, these errors were due to ambiguities in the interface that we predicted and detected in our experimental conditions. Specifically, fifteen participants chose to restrict the photos to their friends and members of the Brown network, making sure that the Conglomi network was not selected. Unbeknownst to them, members of the Conglomi network who were either friends or members of the Brown network still had access. Eight of these participants went so far as to explicitly type the names of their Conglomi friends who should be denied access. This, however, had no effect on the members of the Conglomi network who were also members of the Brown network, but not friends. Thus, this specific ambiguity accounted for 68% of the errors we observed during the party scenario.

In our experimental conditions, participants were prompted when this ambiguity was detected and they were given the opportunity to change their settings. Using Fisher's exact test with the Holm-Bonferroni correction, we found that the feedback in the *Choose* condition was extremely helpful during this task (Table 1): participants used it to remove significantly more errors than those in the other conditions (p < 0.0247 vs. *Control*; p < 0.0055 vs. *Warn*). When alerted to the presence of the ambiguity and given immediate feedback on how it could be clarified, all five participants who would have specified the wrong policy were able to understand and correct their errors before the policy took effect.

The interface in the *Choose* condition detected ambiguities and suggested ways of clarifying them, whereas the interface in the *Warn* condition only detected the ambiguities and left it up to the user to determine how to correct them. The purpose of this was to separate the effect of the feedback from the effect of the guidance. A total of seven participants in this condition were prompted because their policies contained predictable ambiguities. Four of these participants chose to continue without making any changes to their policies, whereas two others made changes that did not correct the errors. The seventh participant initially set

²We considered the list of "friends" to be a network.

| Recruitment | Control | | Warn | | Choose | |
|------------------|---------|-----|------|-----|--------|-----|
| Ν | 13 | | 11 | | 9 | |
| Correct at start | 7 | 54% | 4 | 36% | 5 | 56% |
| Saw feedback | | | 4 | 36% | 2 | 22% |
| Correct at end | 7 | 54% | 5 | 45% | 6 | 67% |

Table 2. The number of participants who correctly set permissions during the recruitment scenario before any feedback was displayed, who saw feedback in the experimental conditions based on ambiguity detection, and who correctly set permissions after seeing the feedback.

the correct permissions, but after seeing the feedback about an ambiguity—the fact that friends were selected but both networks were unselected—he opted to erroneously select both the Brown and Conglomi networks. Thus, providing feedback without guidance did nothing to help six participants and prompted one participant to introduce an error into a previously correct policy. When comparing the performance of those who viewed feedback between the two experimental conditions, those in the *Choose* condition introduced significantly fewer errors than those in the *Warn* condition (p < 0.002 for Fisher's exact test).

Recruitment Scenario

In the recruitment scenario, we asked participants to alter permissions such that only members of both the Brown and Conglomi networks—the intersection—could view the photos. As explained in the Methodology Section, this task was impossible to complete using Facebook's existing interface because a user cannot specify the intersection of two networks, only the union. In the party scenario, when participants in the experimental conditions selected one network but not the other, they were prompted about users in both networks. In the recruitment scenario, this feedback was less helpful because it did not necessarily help them accomplish the task; they needed to grant access to *only* these users.

Before participants in the experimental conditions saw any feedback, we found that a total of sixteen participants across all three conditions (48% of 33) specified the correct permissions (Table 2). Of the seventeen incorrect policies, two of them were due to under-sharing (12% of 17), while the remaining fifteen were due to over-sharing (88% of 17). We observed that this data leakage tendency was likely attributable to the ambiguous interface: ten of the total errors (59% of 17) were due to participants selecting both the Brown and Conglomi networks. Indeed, we predicted that participants might be confused as to whether the system would interpret multiple selected networks as the intersection instead of the union.

We discovered that the presence of feedback in the experimental conditions did little to help participants correct their policies: one participant in each of the two conditions used the feedback to correctly remove policy errors. In this scenario, because the ambiguity detection focused on overlapping networks, feedback was only displayed when participants were on the verge of specifying an incorrect policy; participants who correctly specified individual friends, or incorrectly specified all of their networks, never saw feedback. Thus, for the six participants who saw the feedback, four in the *Warn* condition and two in the *Choose* condition, it could only help them. Of course, this was not a statistically significant improvement over the *Control* condition.

Lessons

We found that without relevant feedback, participants would have introduced errors into their privacy policies a majority of the time. When participants introduced errors, these errors were more likely to result in sharing with unintended parties (data leakage) rather than denying access to authorized users. Our data show that when ambiguities were detected and relevant guidance was offered (e.g., the *Choose* condition), users took the time to edit their privacy settings. However, when relevant actionable guidance could not be offered (e.g., the *Warn* condition), users were just as likely to transform a correct-ambiguous policy into an incorrectunambiguous policy. This leads us to conclude that providing post hoc information about potential policy ambiguities is only likely to be effective when concrete guidance can be offered.

After the experiment, we asked participants how Facebook's privacy settings could be improved. Twenty participants (61% of 33) said something about wanting more fine-grained controls; examples included:

- Not enough options
- There should be options for specific intersections and unions of groups
- It was not specific enough to keep certain people from having access to pictures and info
- Ability to restrict access to individuals with multiple characteristics

The second most common complaint regarding Facebook's privacy settings interface was that it was overly complicated and offered too many options (24% of 33):

- It's a little confusing
- It is sometimes unclear exactly who has access
- It's not great not to know for sure if there are other hidden options that I might have missed
- Looking at who to share the photos was confusing

In short, users wanted more control with fewer options. These complaints hint at a much larger problem: the interface that Facebook provides for modifying privacy settings is simply inappropriate for all but the most basic privacy settings.

VENN DIAGRAMS

In our second experiment, we designed a new interface for specifying access control settings in Facebook. We created this interface as an attempt to completely eliminate the two types of ambiguities we discussed earlier. We reasoned that the act of specifying an access control policy on Facebook is a matter of indicating how the intersections of a user's networks should be granted access. We based our design on a visual metaphor with which study participants would likely be familiar: Venn diagrams. In this section we describe our usability study of this new privacy settings interface.



Figure 4. Venn diagram permissions interface for Facebook.

Methodology

In our second experiment, we examined whether a Venn diagram may help users to better visualize effective permissions for overlapping networks (Figure 4). Each network friends, Brown, and Conglomi—was depicted as a set. For each subset, participants could select "allow" or "deny" from a drop-down box, which caused the selected subset and all the nested subsets to change permissions. The color of each subset also changed to reflect the effective permissions: red for deny, green for allow.

We conducted a usability study to evaluate this interface from June 2-15, 2010. We followed the methodology and recruiting procedures from the first experiment, with two notable changes. First, we created two between-group conditions:

- **Control** Participants used a privacy settings interface identical to the current Facebook interface (Figure 1).
- Venn— Participants used the Venn diagram interface to "customize" settings (i.e., Figure 4 replaced Figure 1B).

Second, we used the following four tasks:

- 1. **Reunion Scenario** Last week you returned from a reunion event at Brown. You had a great time seeing a lot of old friends and took lots of pictures. You decide to upload some of these photos to Facebook in order to share them with your friends and other people affiliated with Brown.
- 2. **Recruitment Scenario** As a small Providence-based business, your manager understands that its important to recruit Brown alumni. You have recently been put in charge

of coordinating a recruiting event with other Brown alumni who also work at Conglomi. At the last planning meeting, you took several photos that you want to share with the other planners. However, these photos should only be shared with the other planners (i.e., Conglomi employees who are Brown alumni). If these photos were disclosed outside of the planning group, you could get fired.

- 3. Work Scenario— You just got a new office at work and are anxious to post pictures of it to show all your coworkers. You decide to upload some of these photos to Facebook in order to share them with people affiliated with Conglomi.
- 4. **Party Scenario** Last Friday all of your coworkers stayed at work until midnight because of a deadline. However, some of your friends from Brown were throwing a party at the same time. You told your coworkers that you were feeling sick so that you could leave work early in order to attend the party. You could be fired if anyone at work finds out that you were at the party. One of your friends from the party has just asked you to upload your photos. Please do so, while being aware of your concerns. You want to share the photos with as many friends and other people who may have attended the party, but with none of your coworkers.

The recruitment and party scenarios were identical to the ones in the first experiment. We designed the reunion and work scenarios in order to easily be completed using Facebook's existing interface and to not mention consequences for incorrect permissions. We did this in order to confirm that our new Venn diagram interface was at least as effective as the existing Facebook interface. The order in which participants performed the tasks was pseudo-randomized using the Latin squares method. After performing these tasks, participants completed an exit survey identical to the one used in the first experiment.

Analysis

We found no significant differences between the randomly assigned conditions with regard to demographics. Overall, our youngest participant was 18, while our oldest was 42 ($\mu = 20.42, \sigma = 3.74$). Our 40 participants were evenly split in terms of gender. Seven of our participants were graduates of Brown, while the remaining 33 were students. Finally, 80% of our participants were from the US, 10% were from Asia, and the rest were scattered.

We found that when using our Venn diagram interface, participants made equal or fewer errors in all four scenarios. In this section we first present our results for the more trivial tasks—the reunion and work scenarios—and then we present our results for the tasks that required more complex settings the recruitment and party scenarios.

Trivial Tasks

In the reunion scenario, we asked participants to grant access to "friends and other people affiliated with Brown." We intended them to allow access only for people who were either friends or members of the Brown network. However,

| | Control | | Venn | |
|-------------|---------|-----|------|-----|
| Reunion | 9 | 45% | 9 | 45% |
| Work | 4 | 20% | 14 | 70% |
| Party | 8 | 40% | 18 | 90% |
| Recruitment | 10 | 50% | 12 | 60% |

Table 3. The number of participants who correctly set permissions during the four scenarios in the experiment. Significantly more participants in the *Venn* condition set the correct permissions for the work and party scenarios.

during the experiment we realized that this wording could also be interpreted to mean "friends from Brown as well as others from Brown." Thus, we also considered it correct if participants granted access to only the Brown network. To correctly perform this task using the existing Facebook interface, participants needed to do the following steps:

- 1. Select "Customize" (Figure 1A).
- 2. Select the Brown network (Figure 1B).

We found that only nine participants in each condition (45%) of 20) correctly set permissions during the reunion scenario (Table 3). In the Control condition, five people erred by over-sharing-granting access to unintended people-while six people denied access to the Brown network, limiting access to friends. In the Venn condition, six participants overshared the photos, while the remaining five erroneously denied access to intended recipients. Of the 22 people across both conditions who set incorrect permissions, 18 of them did so because they chose a preset option (e.g., "friends," "everyone," etc.) rather than launching the custom settings interface. Overall, we were surprised that a minority of participants in each condition set the correct permissions. This may have been due to participants not paying attention to the specific task; the lack of stated consequences may have caused some to over-share, while those who under-shared may be habituated to only sharing content with friends.

In the work scenario, we asked participants to simply share content with the entire Conglomi network. To correctly perform this task using the existing Facebook interface, participants needed to perform the following steps:

- 1. Select "Customize" (Figure 1A).
- 2. Change "Only Friends" to "Only Me" (Figure 1B).
- 3. Select the Conglomi network (Figure 1B).

Participants in the *Venn* condition were significantly more likely to perform this task correctly than those in the *Control* condition (p < 0.0036 for Fisher's exact test). Only six participants in the *Venn* condition made mistakes, whereas fourteen participants in the *Control* condition produced erroneous policies. The errors in the *Control* condition were the direct result of the interface's default settings: ten participants over-shared by neglecting to change the default in Step 2 from "Only Friends" to "Only Me." This had the effect of sharing the photos with the union of the Conglomi and friend networks, rather than only the Conglomi network. Participants in the *Venn* condition did not make this mistake because the interface defaulted to denying everyone access.

Complex Tasks

In our first experiment, we examined two impossible scenarios: denying access to an entire network and specifying the intersection of two networks. We designed the interface in the *Venn* condition directly around these limitations, so that users could visualize exactly which subsets of their friends and networks would be granted or denied access.

In the party scenario, we asked participants to grant access to friends and members of the Brown network, while denying access to anyone affiliated with Conglomi. As written, this task was impossible for participants in the *Control* condition. Instead, these participants could grant access to friends and then deny access to coworker-friends using the steps outlined in the methodology for the first experiment.

We found that participants in the *Venn* condition made significantly fewer errors than those in the *Control* condition (p < 0.002 for Fisher's exact test); twelve participants made errors in the *Control* condition compared to only two in the *Venn* condition. Only one of these erroneous policies resulted in under-sharing—the participant opted to share with no one. No participants in the *Control* condition explicitly granted access to the Conglomi network, though members of this network were inadvertently granted access through their memberships in the Brown network. This over-sharing was the most common mistake in the *Control* condition, and accounted for 83% of the errors.

Of the two participants who made errors in the *Venn* condition, one never saw the custom interface because he shared with "Only Friends" (Figure 1A). The second participant mentioned that some of her "work friends are in Brown" and so she denied access to the entire Brown network, rather than just the intersection with the Conglomi network.

In the recruitment scenario, we asked participants to grant access to the intersection of the Conglomi and Brown networks. Like the first experiment, this task was impossible using Facebook's existing interface. Instead, users in the *Control* condition were able to complete this task by manually granting access to people at the intersection of their three networks: friends, Brown, and Conglomi.

Once participants in the Venn condition viewed the custom interface, they could correctly set permissions with a single mouse click. Despite this, we were surprised to find no significant differences between the two conditions. Eight people made mistakes in the Venn condition, whereas ten people made mistakes in the *Control* condition. The mistakes in the *Venn* condition stemmed from five participants over-sharing by simply allowing the entire Conglomi network (63% of 8), and three other participants under-sharing by denying access to the intersection of the three networks (37% of 8). In the Control condition, 90% of the errors were due to oversharing. A third of these were from participants selecting both the Brown and Conglomi networks, believing that this would represent the intersection and not the union. We repeatedly observed participants making this same mistake in the first experiment due to the ambiguous interface.

DISCUSSION

We examined two techniques to help users minimize access control errors on Facebook. In our first experiment, we observed that providing users with feedback about potential policy ambiguities was only effective when specific guidance could be offered. Without guidance, users were significantly less likely to prevent errors, and in some cases, introduced errors where none previously existed. In our second experiment, we introduced a Venn diagram interface to help users see how their networks overlapped. Overall, users of this interface introduced 55% of the errors that those using the existing Facebook interface introduced. In this section we present data on participants' privacy behaviors outside of this study, as well as limitations and future work.

Privacy Behaviors

During the six tasks spanning the two experiments, we found that participants were more than 3.5x as likely to over-share than to under-share when they created access control policy errors. In order to gauge how pervasive this problem is outside of the laboratory, we asked all 73 participants in both experiments several questions about their regular Facebook usage and privacy settings.

All of our participants had used Facebook for over a year, and 90% of them had used it for more than two years. In order to examine the potential for information leakage, participants showed us their profiles so that we could examine the number of photos in which participants were tagged ($\mu = 526.32$, $\sigma = 353.11$), networks they had joined ($\mu =$ 1.77, $\sigma = 0.51$), and applications that they had installed ($\mu = 50.17$, $\sigma = 90.77$).³

We asked participants how and why they modified their privacy settings and found that twenty of them (27% of 73) specifically mentioned changing their settings in response to media reports of Facebook changing the defaults. These results are similar to findings by boyd and Hargittai [2]:

- Infrequently, usually only when my friends inform me that Facebook has altered the default privacy settings
- When Facebook is updated I double check that it is still set on only friends
- Every time I hear a privacy change has been made

Despite the high potential for data leakage and perceptions of constantly changing privacy settings, six participants stated that they had never modified their privacy settings prior to this study (8% of 73). Twenty-eight participants (38% of 73) mentioned that they simply restricted their profiles to friends. The remaining 46 participants used other strategies to keep their information private, several mentioned going out of their way to block coworkers or family members:

- Everyone could see everything except my dad
- I blocked my employer from viewing pictures of me
- Blocked students I work with from viewing my photos
- Limited profile for adults/relatives who had friended me

Limitations and Future Work

Detecting ambiguities and providing users with guidance on how to prevent semantic errors was effective in certain cases, though it can be counterproductive when specific guidance cannot be offered or it is not helpful. In response to these findings, we designed and tested a new interface that allowed users to visualize how privacy settings applied to their overlapping networks. When we asked participants in our second experiment to rate the interface they used on a 5-point Likert scale, they rated the Venn diagram interface significantly higher than the control, both overall ($t_{38} = 3.08, p <$ 0.004), and in terms of ease of use $(t_{38} = 3.10, p < 0.004)$. One might be concerned that the ease of use imparted by the Venn diagram solution may only apply to computer science majors. However, 92% of our participants were from outside the computer science department. At the same time, further study is needed to determine if Venn diagrams are intuitive to Facebook users without a college education.

A Venn diagram interface is also only usable if participants have three or fewer overlapping sets (i.e., two networks plus a list of friends). Of the 73 participants across both experiments, we observed that 20 (27%) belonged to only one Facebook network, 50 belonged to two networks (69%), and three belonged to three networks (4%). Taking the 95% Confidence Interval (CI), this implies our solution is usable by at least 88% of our target demographic. Gross and Acquisti analyzed data from a 2008 Facebook dump of the entire Carnegie Mellon network (following their 2005 study [9]).⁴ Of the 6,404 viewable profiles, 5,591 were members of two or fewer networks. This indicates a 95% CI of over 86%.

We chose to not directly compare the techniques in the two experiments because the experimental procedures differed. Further study is needed on the effectiveness of a privacy settings interface that offers both a Venn diagram visualization and post hoc guidance. Finally, the ambiguities that we examined were far from a complete list of the shortcomings in Facebook's privacy settings interface. We also believe that the lack of realism and consequences in our study tasks impacted the effort participants exerted in setting the correct permissions (likely indicating an upper bound for the error rate). A longitudinal study is needed to determine how participants cope with complex privacy scenarios in real life, how those situations can be automatically detected, and what types of feedback are most effective at minimizing errors.

ACKNOWLEDGMENTS

This work was supported by NSF grants CT-0830945 and CNS-0627310. The first two authors were affiliated with Brown University when this research was performed.

REFERENCES

 J. Binder, A. Howes, and A. Sutcliffe. The problem of conflicting social spheres: Effects of network structure on experienced tension in social network sites. In *CHI* '09: Proceedings of the 27th International Conference on Human Factors in Computing Systems, pages 965–974, New York, NY, USA, 2009. ACM.

⁴Private communication on November 17, 2010.

³We did not collect data on applications during the first experiment.

- 2. d. boyd and E. Hargittai. Facebook privacy settings: Who cares? *First Monday*, 15(8), August 2010.
- C. A. Brodie, C.-M. Karat, and J. Karat. An empirical study of natural language parsing of privacy policy rules using the SPARCLE policy workbench. In SOUPS '06: Proceedings of the Second Symposium on Usable Privacy and Security, pages 8–19, New York, NY, USA, 2006. ACM.
- Computer Research Association. CRA Conference on Grand Research Challenges in Information Security & Assurance, 2003. http://archive.cra.org/reports/trustworthy.computing.pdf.
- J. DiMicco, D. R. Millen, W. Geyer, C. Dugan, B. Brownholtz, and M. Muller. Motivations for social networking at work. In CSCW '08: Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work, pages 711–720, New York, NY, USA, 2008. ACM.
- J. M. DiMicco and D. R. Millen. Identity Management: Multiple presentations of self in Facebook. In *GROUP* '07: Proceedings of the 2007 International ACM Conference on Supporting Group Work, pages 383–386, New York, NY, USA, 2007. ACM.
- K. Fisler and S. Krishnamurthi. A model of triangulating environments for policy authoring. In SACMAT '10: Proceeding of the 15th ACM Symposium on Access Control Models and Technologies, pages 3–12, New York, NY, USA, 2010. ACM.
- K. Fisler, S. Krishnamurthi, L. A. Meyerovich, and M. C. Tschantz. Verification and change-impact analysis of access-control policies. In *ICSE '05: Proceedings of the 27th International Conference on Software Engineering*, pages 196–205, New York, NY, USA, 2005. ACM.
- R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In WPES '05: Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, pages 71–80, New York, NY, USA, 2005. ACM.
- iStrategy Labs. Facebook demographics and statistics report, January 2010. http://www.istrategylabs.com/2010/01/facebookdemographics-and-statistics-report-2010-145-growthin-1-year/.
- A. N. Joinson. Looking at, looking up or keeping up with people?: motives and use of facebook. In CHI '08: Proceedings of the 26th International Conference on Human Factors in Computing Systems, pages 1027–1036, New York, NY, USA, 2008. ACM.
- C. Lampe, N. Ellison, and C. Steinfield. A face(book) in the crowd: Social searching vs. social browsing. In CSCW '06: Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work, pages 167–170, New York, NY, USA, 2006. ACM.

- C. Lampe, N. B. Ellison, and C. Steinfield. Changes in use and perception of Facebook. In CSCW '08: Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work, pages 721–730, New York, NY, USA, 2008. ACM.
- M. M. Skeels and J. Grudin. When social networks cross boundaries: A case study of workplace use of Facebook and LinkedIn. In *GROUP '09: Proceedings* of the 2009 International ACM Conference on Supporting Group Work, pages 95–104, New York, NY, USA, 2009. ACM.
- 15. C. Steinfield, J. M. DiMicco, N. B. Ellison, and C. Lampe. Bowling online: social networking and social capital within the organization. In C&T '09: Proceedings of the 4th International Conference on Communities and Technologies, pages 245–254, New York, NY, USA, 2009. ACM.
- K. Strater and H. R. Lipford. Strategies and struggles with privacy in an online social networking community. In BCS-HCI '08: Proceedings of the 22nd British HCI Group Annual Conference, pages 111–119, Swinton, UK, UK, 2008. British Computer Society.
- F. Stutzman and J. Kramer-Duffield. Friends only: examining a privacy-enhancing behavior in Facebook. In CHI '10: Proceedings of the 28th International Conference on Human Factors in Computing Systems, pages 1553–1562, New York, NY, USA, 2010. ACM.
- A. Tootoonchian, K. K. Gollu, S. Saroiu, Y. Ganjali, and A. Wolman. Lockr: social access control for web 2.0. In WOSN '08: Proceedings of the 1st Workshop on Online Social Networks, pages 43–48, New York, NY, USA, 2008. ACM.
- T. Turner, P. Qvarfordt, J. T. Biehl, G. Golovchinsky, and M. Back. Exploring the workplace communication ecology. In *CHI '10: Proceedings of the 28th International Conference on Human Factors in Computing Systems*, pages 841–850, New York, NY, USA, 2010. ACM.
- 20. J. Watson, M. Whitney, and H. R. Lipford. Configuring audience-oriented privacy policies. In SafeConfig '09: Proceedings of the 2nd ACM Workshop on Assurable and Usable Security Configuration, pages 71–78, New York, NY, USA, 2009. ACM.
- 21. A. Wu, J. M. DiMicco, and D. R. Millen. Detecting professional versus personal closeness using an enterprise social network site. In *CHI '10: Proceedings* of the 28th International Conference on Human Factors in Computing Systems, pages 1955–1964, New York, NY, USA, 2010. ACM.
- 22. A. L. Young and A. Quan-Haase. Information revelation and internet privacy concerns on social network sites: a case study of Facebook. In C&T '09: Proceedings of the 4th International Conference on Communities and Technologies, pages 265–274, New York, NY, USA, 2009. ACM.

May 7-12, 2011 • Vancouver, BC, Canada