Out of sight, out of mind: Effects of displaying access-control information near the item it controls

Kami Vaniea*, Lujo Bauer*, Lorrie Faith Cranor*, and Michael K. Reiter†
*Carnegie Mellon University, Pittsburgh, PA, USA
{kami,lbauer,lorrie}@cmu.edu
†University of North Carolina at Chapel Hill, Chapel Hill, NC, USA
reiter@cs.unc.edu

***Abstract*—We take a detailed look at how users, while focusing on non-permission tasks, notice and fix access-control permission errors depending on where the access-control policy is spatially located on a photo-sharing website. The access-control policy was placed on an online photo-sharing website under the photo or album, on the sidebar, or on a separate settings page. We find that placing the access-control policy directly under photos and album thumbnails improves participants' ability to notice errors in their access-control settings without negatively impacting non-access-control tasks.**

***Keywords*-access control, human factors, visualization**

I. Introduction

End users find it challenging to stay aware of and manage sharing preferences for content that they publish on social networks and photo-sharing sites [1], [2], [3], [4]. This problem is becoming even more difficult as sites become more dynamic, with constant uploading of content and shifting friend groups.

Online users rarely interact with access-control policy as their primary task [5]. They log onto Facebook, Flickr, or Google+ to share content, catch up on news, and interact with friends—not to "do security." Access control typically remains in the background until some event, such as an embarrassing experience, brings it to a user's attention [2].

An empirical study of Facebook users compared participant's sharing intentions to the implemented privacy policy, and found that every participant had at least one mismatch [1]. Another survey, of Facebook users' understanding of applications, found that only one out of 516 surveyed users was able to accurately answer what parts of their Facebook profile the survey application could access [6].

In this paper, we investigate interfaces intended to help users stay abreast of their access-control policy even when they are engaged in another activity as their primary task. More specifically, in the context of a photo-sharing site, we investigate whether making access-control policy visible to users while they are engaged in a non-security-related primary task can improve the users' understanding of and ability to correctly set a desired access-control policy. We test two kinds of interfaces: a sidebar interface, in which access-control policy information is embedded in the sidebar of the main album-management interface; and an under-photo interface, in which access-control information is shown under a photo or album when a user moves the mouse cursor over it. We call these interfaces *proximity interfaces*, because the access-control information is in close spacial proximity to the item it describes.

To understand the effect of these interfaces, we carried out a 34-person laboratory study. Comparing sidebar and under-photo proximity interfaces to a control condition in which a user has to actively seek out access-control information, we tested whether proximity interfaces (1) help users notice and fix access-control errors; (2) help users become more aware of their access-control settings; and (3) interfere with users' ability to execute non-access-control related tasks.

We found that participants who were shown access-control information under the photo or album they were working with were statistically significantly more likely to notice and fix policy errors. On the other hand, participants in the sidebar condition performed similarly to those in the control condition. Neither the sidebar nor the under-photo proximity interfaces appeared to interfere with users' ability to execute non-access-control-related tasks. We also collect eye-tracker data, which we use to understand in more detail when and how participants notice information presented via proximity displays.

Overall, our results bolster the case for using proximity interfaces for displaying access-control policy, but also highlight the importance of integrating them as tightly as possible with users' primary tasks.

II. Proximity Access-Control Displays

Proximity displays for access control put access-control information in close spatial proximity to the item that the information describes. In this manner, even users who are not pursuing an access-control-related task will be exposed to the access-control policy for the album they are working with and can obtain detailed information with little effort.

A. Design

The proximity displays that we explore in this paper show access-control policy in grid form, with each row of the

978-1-4673-2326-0/12/$31.00 ©2012 IEEE

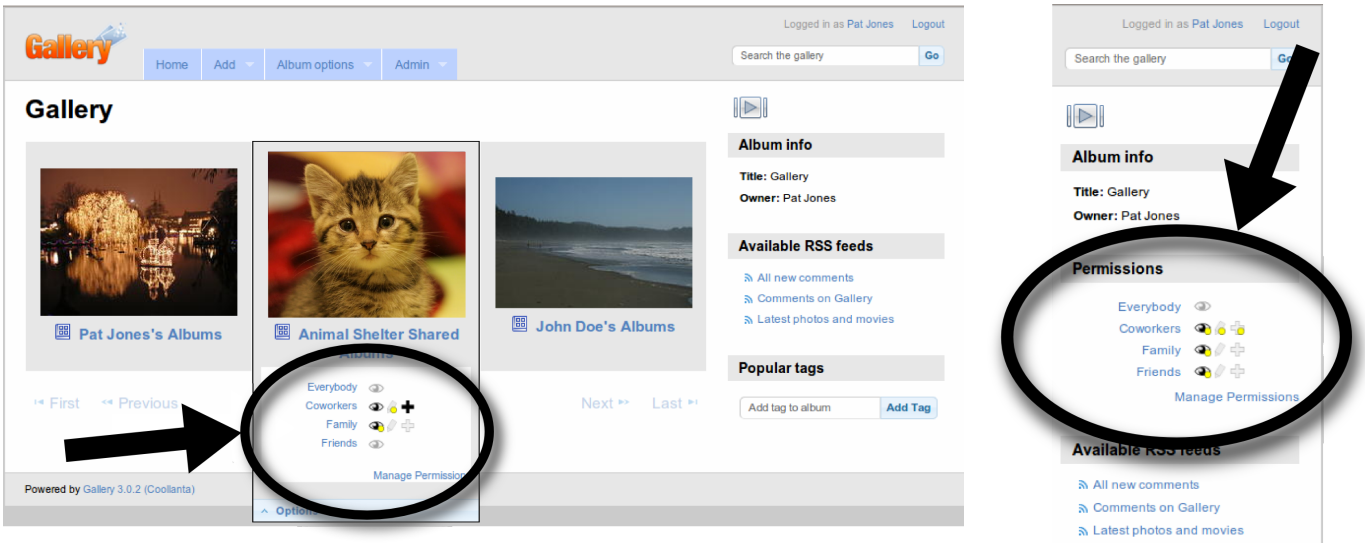

(a) Under-photo condition

(b) Sidebar condition

Figure 1. Examples of Gallery interface with proximity displays: (a) in the under-photo condition, and (b) the sidebar condition. Proximity display in the under-photo condition. The proximity display in (a) shows that the group Everybody has no permission; Coworkers can view and add to this album and all subalbums, but can only edit some subalbums; Family can view some this album and all subalbums; and Friends cannot view anything.

grid showing the permissions a particular group has to the album in question. Mousing over the group name will reveal the group members, and the permissions are indicated by icons (view 👁, edit ✎, and add photo ✚). Grayed-out or missing icons indicate lack of permission; icons with a yellow dot indicate that subalbums or photos do not have consistent permissions (e.g., the group may have a specific permission on some subalbums but not on others). If a group cannot view an album, then all other permissions are also unavailable. Figure 1(a) is an example of such a display taken from our under-photo condition. Mousing over any icon on the proximity display results in a tool tip with an explanation of the permission in its current context. In Figure 1(a), for example, mousing over the icon next to "Everybody" would display "The group Everybody cannot view Animal Shelter Shared Albums."

The proximity information display design is based on several existing works which we will briefly mention here. Further details can be found in Related Work (Section VI).

The idea of using close spatial proximity to link concepts is well known and part of Gestalt principles [7]. We use it here to bring access-control information into the immediate context of the user's work-flow. We want checking and changing access-control settings to be as natural as checking other spatially linked features such as titles.

The grid design is based on work by Reeder et al., who successfully used a combination of grids and effective permissions (discussed below) to make it easier for users to manage file permission settings [8], [9]. Participants were able to use the grid to get a quick sense of permissions settings and focus on important components easily.

The decision to use icons in the grid is based on work by Tam et al., who tested multiple permission display layouts against participant comprehension speed [10]. They found that displays that used visual icons allowed users to find data quicker and were preferred by the users. They also found that participants performed better when permissions were organized by action icons.

B. Implementation

We implemented the proximity displays on Gallery, an open source photo-sharing website system [11]. We chose Gallery because it has a rich API for interacting with access-control permissions. Excluding the changes described below, we used a default Gallery 3.1 installation and theme.

1) Effective permissions and hierarchies: Gallery allows albums to contain subalbums. Permissions on parent and child albums can be different, making them harder to visualize when looking at only a single level of an album hierarchy. Similarly, Gallery has a built-in group "Everybody," which includes users from all other groups, making possible policy conflicts between Everybody and other groups. Prior work shows that end users find it much easier to understand access-control policy when they are shown *effective permissions* (the result of evaluating all relevant policy rules) rather than the sets of policy rules that induce them [8]. Accordingly, we designed our proximity interfaces to show effective permissions.

2) Permission-modification interface: Gallery does not have a holistic interface for viewing and modifying permissions; instead, permissions need to be viewed or adjusted independently for each album. As this makes understanding

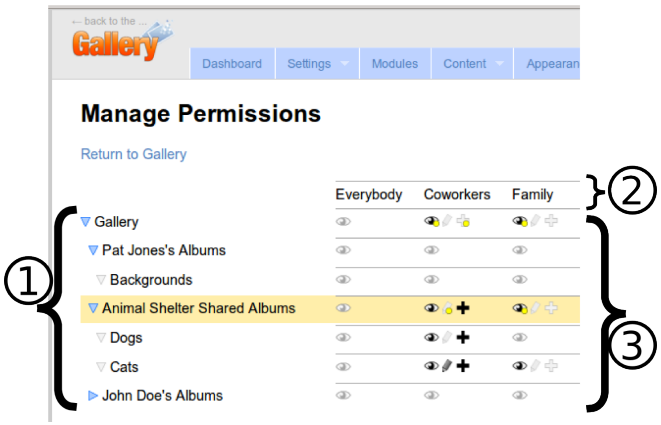

Figure 2. Policy-modification interface used by all participants to make changes to the access-control policy. All the albums are listed along the left (1), user groups are listed along the top of the grid (2), and view, edit, and add permissions are shown as icons in the central grid (3). This interface also contains a legend (not shown) in the bottom left explaining the meaning of all the symbols.

Task	Area	Permission subtask	Album state	Prompted
	Work Information Page			
1-5	Warm-up	Read, Add	Existing	Prompt
6	Coworkers	None	Existing	None
7	Coworkers	Add	New	None
8	Coworkers	Remove	Existing	None
9	Coworkers	Read	Changed	Prompt
	Friends Information Page			
10	Friends	Remove	New	Prompt
11	Friends	Read	Existing	Prompt
12	Friends	None	Existing	None
13	Friends	Add	Changed	Prompt
	Family Information Page			
14	Family	Add	Existing	Prompt
15	Family	None	Existing	None
16	Family	Read	New	None
17	Family	Remove	Changed	Prompt

Table I
TASKS AND INFORMATION GIVEN TO PARTICIPANTS IN CHRONOLOGICAL ORDER.

the overall policy unreasonably difficult, we implemented a grid-style policy interface that allowed viewing and editing of all permissions on one screen. This interface was heavily inspired by recent work on Expandable Grids [8], [9], and was available to users in all our conditions. An example of this interface is shown in Figure 2.

3) Proximity displays in Gallery: We modified Gallery to show proximity displays either directly below the album thumbnail when a user positions the mouse cursor on the thumbnail (under-photo condition) or along the side of the screen (sidebar condition). The default Gallery interface reacts to a user mousing over an album thumbnail by showing the title, owner, and number of visits under the album thumbnail. To prevent access-control related information from appearing in several places on the screen, we removed the Gallery-supplied information about album ownership and number of views from all conditions. In the under-photo condition, we replaced this information with our proximity display. In the sidebar condition we placed a proximity display as the second item on the sidebar, between the album info and RSS feed information that is in the sidebar by default.

III. METHODOLOGY

We designed a 1.5-hour laboratory study in which 34 participants were divided into three conditions: two proximity-display conditions and a control condition. In the study, users took part in a role-playing scenario in which they performed a variety of tasks, including various permissions-management tasks on a set of albums. We arrived at the final design for the study after a 4-person pilot and a 26-participant pre-study.

A. Protocol

The study was a between-participants design with a round-robin assignment to experimental conditions. A think-aloud protocol was used. Participants in all conditions performed the same tasks, and the only variable between conditions was the Gallery interface participants were exposed to. The control condition displayed the default interface, which always included a link to the holistic policy-visualization-and-editing grid described in the previous section. The sidebar condition included a proximity display in the sidebar, and the under-photo condition included a display that appeared under each photo or album when the mouse cursor was over the photo/album. The tutorial used to familiarize the participant with the Gallery interface also differed slightly by condition.

Participants were asked to role play the part of Pat Jones, who manages several online photo albums using Gallery. During the course of the study, participants received information about events in Pat's life, including emails from coworkers, family, and friends. These emails, delivered to participants in printed-out form by the researcher administering the study, included requests from Pat's coworkers, family members, and friends to perform various tasks with the online albums.

As Pat Jones, participants started with a tutorial that asked them to walk through manipulating photos using Gallery which had been previously set up with seven albums in hierarchies and simplistic permissions. When the participant completed the tutorial, the researcher had them open a new Gallery site that had many more albums and more complex permissions. These albums did not overlap the tutorial albums.

After the tutorial, the participant was first asked to perform five clearly defined and progressively more complex

warm-up tasks (rows 1–5 in Table I): rotate a photo, read a permission, delete a photo, change a permission, and change some titles. If any tasks were not successfully completed, the researcher prompted the participant with an email that pointed out the error; if the participant still could not complete the task, they were instructed by the researcher how to do so. This was done to ensure that all participants knew how to operate Gallery and to help them get acclimated to working with the albums.

The bulk of the study consisted of tasks 6–17, summarized in Table I. Each task was composed of a set of *subtasks*, individual permission, rotation, deletion, spelling, or renaming errors that needed to be corrected. Each task had a primary subtask directly expressed by the email sender and several additional subtasks implied by errors such as rotated photos, or incorrect permission errors. The tasks were divided into three sets based on whether the albums the participant would manipulate contained photos of coworkers, friends, or family (shown in the second column of Table I). Before each set of tasks, the participant was given an information sheet explaining their normal interactions with this group of people. Half the tasks required adding or removing a permission (shown in the third column of Table I). A quarter conveyed to the participants desired permissions, but no permissions needed to be changed. The final quarter had no access-control component. All tasks contained at least one title, rotate, delete, or organize subtask intended to distract the participant. Each task was performed on albums in one of three states (shown in the fourth column of Table I). *Existing* albums were already set up in Gallery when the participant started. *New* albums were created by the participant. *Changed* albums were those for which the participant had previously read or changed a permission, but, unknown to the participant, some part of the album had been altered by the researcher after the participant had last seen the permissions. Tasks for which failure to complete a permission subtask resulted in an email calling this out were called *prompted*; all others were *unprompted* (rightmost column, Table I). When a participant failed to complete a prompted task they received an email from one of Pat's coworkers, friends, or family members pointing out the error and requesting that it be fixed.

In addition to the task-related albums, there were four albums which the participant was never directed to interact with. Two of these albums had correct permissions and two albums had incorrect permissions.

At the end of the study, participants filled out a survey that asked them to recall the view and add permissions for every album they worked with, the two albums which had incorrect permissions but were not part of a task, and two non-task albums with correct permissions. For each suggested combination of album, group, and permission the participant could answer *True*, *False*, or *Not Sure*. For each set of questions about an album the participant was asked how confident they were of their answers.

B. Recruitment and demographics

We recruited 34 participants using a university-run electronic bulletin board for advertising research studies. Participants ranged in age from 18 to 41 with a mean age of 23.9. Twenty two of the participants were students. One participant was excluded due to an inability to complete even half the study in the allotted 1.5 hours. After this exclusion, we were left with 11 participants per condition.

C. Data collection and analysis

We collected and coded data derived from a combination of in-session notes, screen-capture video, audio, exported information from an eye tracker, a snapshot of the resulting permission state of the photo website, and the survey. All data was loaded into a database so information from different sources could be correlated. Unless otherwise noted, all statistical tests are one-way ANOVAs using study condition as the explanatory variable. We adjusted the alpha values in our statistical tests to account for multiple testing.

1) Notes and task accuracy: During the study the researcher kept detailed in-session notes and collected timestamps every time a participant was given an email or information page.

Each subtask was coded as *correct* or *error* based on the state of the subtask when the participant declared themselves to be done with the task. We also collected and coded the state of subtasks after the participant had received all email prompts, but this data was used only to verify that participants were able to complete the subtasks. Unless otherwise noted, all references to subtask accuracy refer to the subtask state before the participant was prompted.

Read-permission subtasks were considered correctly completed if participants read the permissions (as defined below) and did not change them. Control participants were judged to have *read the permission* if they opened the holistic permission-management interface and the permission was visible on the screen. Non-control participants were judged to have *read a permission* if they (1) opened the permission-management interface; or (2) read the permission aloud; or (3) indicated through mouse behavior that they were reading the permission display; or (4) pointed at the permission display with their hand while clearly reading the screen.

2) Eye tracker: We used an SMI eye tracker to record video of events occurring on the screen, audio of the participant, and the time and screen coordinates of fixations and user events (e.g., mouse clicks).

In the under-photo condition, proximity displays appeared below photos and tended to be visible for only short times. To determine when and where displays appeared for each user we used a custom Matlab script that scanned each video frame for a unique static part of the proximity display and recorded the time and location of each display.

3) Permission settings at end of each session: At the end of every study session, we archived the state of the entire photo-sharing website, and its final permission settings were automatically extracted into a database. The website was then reset to its initial state.

To categorize changes to permission settings, we introduce some terminology. Here, a "permission" refers to the triple of group, album, and action, and can be set to *allow* or *deny*. We call a permission *specified* if the participant was informed what the permission should be (i.e., allow or deny), and *unspecified* otherwise. All permissions related to an album mentioned in a task are *explicitly part of a task*. Permissions are *implicitly part of a task* if a participant was told what the permissions should be, but not told to manipulate the album. All other permissions are *not part of a task*.

Each permission at the end of a session was compared to its initial state and marked as "changed" or "unchanged." Permissions that were *specified* were compared to their intended state (i.e., the state to which the participant was instructed to set them), and marked as *correct*, *too permissive*, or *too restrictive*. Permissions that were *unspecified* and *changed* were compared to initial permission settings and marked as *more permissive* or *more restrictive*. We refer to all unspecified changed permissions as *uninstructed changes*.

IV. Hypothesis testing

In this section we test our three hypotheses: that participants who saw proximity displays (1) notice and fix access-control errors; (2) remember permissions when asked to recall them; and (3) experience no negative effects on non-permission tasks. In the next section we will use the other data we collected to analyze why our participants behaved the way they did.

A. Notice and fix access-control errors

As can be seen in Figure 3, participants in the under-photo condition outperform both sidebar and control participants in noticing and fixing permission errors study-wide. Participants in the under-photo condition also perform better on tasks where permissions need to be removed. We found limited difference in participants' ability to notice and fix issues with albums that were not directly part of a task.

Participants in the under-photo condition were significantly better than sidebar ($p<0.005$) and control ($p<0.006$) at reading and fixing errors without prompting. Under-photo participants correctly completed 52% of read and change permission subtasks while sidebar completed 30% and control completed 24%. There was no significant difference between control and sidebar participants' ability to read permissions and fix errors without prompting. To analyze the effect of condition on participants' ability to notice and fix errors, we used a linear mixed model where user and task were treated as random effects and where condition, album state, and permission subtask were fixed effects.

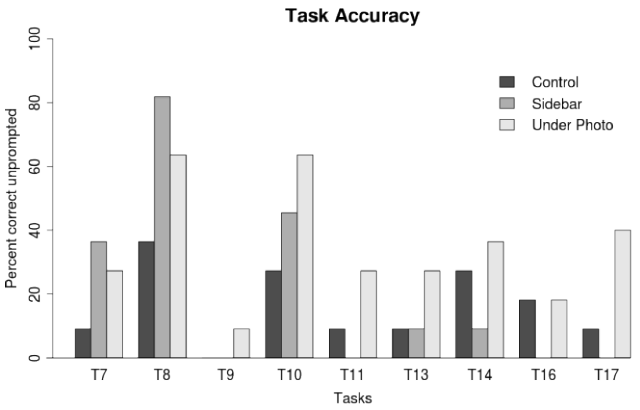

Figure 3. Percent of participants who either correctly noticed and fixed an access-control errors before they were prompted to do so (tasks 7,8,10,13,14, and 17) or read a permission before they were prompted to do so (tasks 9, 11, and 16) as a function of task and study group.

As expected, we observed that very few control participants proactively read permissions (as defined in Section III-C1). Based on in-session observation, control participants proactively read permissions only 11 times out of the possible 66 participant-subtask pairs, and 8 of those instances were by two participants. Seven of the eleven control participants never read any permission before being prompted and no participant read all permissions unprompted. Compare this to the under-photo condition where permissions were proactively read in 27 of the 66 participant-subtask pairs. Every under-photo participant proactively read at least one permission and one participant read every single permission.

If we further break down the tasks by permission subtasks and album states, we find that when looking at permission subtasks, under-photo users performed significantly better than control ($p<0.0004$) on tasks where participants were asked to remove permissions (8, 10, and 17). However, there was no significant difference for tasks where participants needed to add permissions (7, 13, and 14) or where permissions did not need to be changed (9, 11, and 16). We found no significant difference between conditions when working with different album states. All participants did significantly worse when working with changed albums (8, 13 and 17) than existing (8, 11, and 14) ($p<0.0002$) or new (7, 10, and 16) ($p<0.003$) albums.

B. Remember their permissions when asked to recall them

Each participant was asked 128 questions about 13 albums, 4 groups, and 2 actions (view and add). Participants' ability to recall permissions was not significantly different across conditions. Participants were most accurate when answering questions about albums they had worked with and least accurate on albums they had never worked with.

On average, participants answered 58% of the questions, and 72% of those they answered correctly. When asked about

albums that had specified permissions, participants answered 75% of the questions, and were correct on 76% of those. When asked about albums that were not part of a task, participants answered 43% of the questions and were correct only 66% of the time.

C. Experience no negative effects on non-permission tasks

The sidebar and under-photo participants experienced no negative effects from the introduction of proximity displays on their screens. We looked at the impact the displays might have on their ability to complete permission subtasks after prompting, complete non-permission subtasks before prompting, and the average time required to complete tasks.

Participants in all conditions were able to complete permission tasks equally well once they were prompted to do so. This shows that our display did not negatively impact participants' ability to complete the task when they knew it needed to be completed. With the exception of tasks 7, 8, and 16, if a participant did not successfully complete a permission task they were prompted, once, to do so. One participant from each condition experienced one instance when they did not complete a permission subtask after trying in response to prompting. One additional sidebar participant did not attempt to complete one task after prompting.

We also tested participants' ability to complete the non-permission subtasks. Each participant was asked to create 3 new albums, upload 14 photos into those albums, rotate 13 photos, delete 5 photos, move 23 photos to other albums, and reorder 18 photos in one album. We ran a one-sided ANOVA using participant as a random variable and concluded there was no significant difference in non-permission task accuracy between conditions ($p>0.1$). Participants were able to complete an average of 95% of rotate, delete, organize, and create tasks without needing to be prompted. Title tasks were the most challenging, with an average of 75% completed correctly without prompting. In the pre-test we had problems with many participants completing all non-permission tasks so we made the tasks more challenging to increase the likelihood that we would notice any impact of extra cognitive load caused by proximity displays.

Participants took similar amounts of time to complete each task regardless of their study condition. Because this study used think aloud, which causes greater variation in timing, it is difficult to make a claim of no difference with great certainty. The pre-study had identical conditions and very similar tasks but did not use think aloud. In the pre-study we observed no statistically significant difference in task times between conditions. We therefore conclude that it is unlikely that proximity displays cause any difference in the time required to complete a task.

V. Qualitative Results and Discussion

We collected detailed observational and eye-tracking data from each participant as they interacted with the interface.

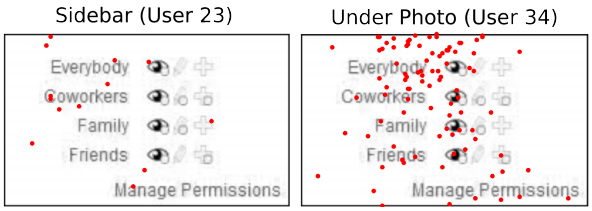

Figure 4. Location of all fixations on proximity displays for the median sidebar and under-photo participants.

These observations allow us to get a more complete picture of why some conditions outperform others and the general permission-error noticing behavior of participants. In this section examine how proximity displays affect users' interactions with access-control permissions.

A. Noticing permission errors

The first step in fixing a permission error is to notice that it is there. Websites like Facebook and Flickr continuously face the problem that their users are not aware of their privacy permissions are and that those permissions differ from what the users want [6], [1].

We examined when our participants noticed proximity displays and how that related to checking and fixing permission errors in the full permissions interface. We observed that under-photo participants periodically checked permissions while engaged in different tasks. Control participants seemed to either check permissions on every task or completely forget about permissions.

1) Permissions are changed at the beginning and end of tasks: Basing our analysis on observations made by the researcher during the study, we notice that study participants often changed or read a permission either right after reading the email or at the end of the task described by the email. When a participant focused on permissions—either by verbalizing this or by clearly focusing on a proximity display—they immediately either dismissed the permissions as correct or corrected them.

This behavior can be best seen on task 10, where the participant is first given an information sheet that explains that her mother panics easily about things like jumping out of airplanes, and so Pat does not mention these things. Then the participant gets an email asking the participant to upload a set of photos of people (including Pat) jumping off of buildings. Many participants immediately change permissions so Family cannot see any albums and then complete the rest of the task. Other participants complete the task, including titling one of the photos "Pat Jones," get to the end of the task, re-read the email, glance at the album and then either check a proximity display or verbalize a concern about permissions.

We hypothesize that permissions were changed at the beginning or end of tasks because it was necessary to go to

a secondary page to make changes. Other subtasks appear to have been performed in arbitrary order.

2) Under-photo participants look at the proximity displays frequently and do not immediately change permissions: Examining eye-tracker data enables a more nuanced understanding of participants' use of proximity displays in our test conditions, since it allows us to observe participants fixating on proximity displays even when this is not otherwise apparent to the researcher.

The sidebar participants saw a single proximity display per page and a display was almost always visible to them. Under-photo participants only saw a display when their mouse was over an album or photo. As can be seen in Figure 1(a), when a user places their mouse over the album, an "options" link appears under the thumbnail. In the under-photo condition, the proximity display appears between the user's current mouse position and the newly appeared "options" link, forcing participants to move their gaze across the proximity display to focus on the "options" link.

The eye-tracker output distinguishes between a fixation, where the participant's gaze rests on a point, and instances when a participant's gaze rapidly passes over a point. This allowed us to determine when a participant looked at the display as opposed to passing their gaze over it. Under-photo participants fixated on significantly more displays than participants in the sidebar condition ($p<0.002$). However, there was no statistical difference in the total amount of time spent fixating on the displays ($p>0.044$).

The sidebar participants primarily look at the proximity display just before navigating to the permission-management interface; under-photo participants look at the displays earlier and then later look just before navigating. We examined all pages participants worked with before navigating to the permissions modification interface and used eye-tracker data to determine when a participant was fixating on proximity displays. Figure 5 shows how often participants in both experimental conditions fixated on a proximity display, normalized over the time they spent on the page. Sidebar participants do exactly as expected: they look at proximity displays rarely while viewing the page, but when they do notice them they tend to then navigate to the permission-management page (hence the peak between 80 and 100% on the first graph in Figure 5). Under-photo participants look at displays at various points while working with a page and then look again just before navigating to the permission-management interface (hence the peak just before 100% on the second graph in Figure 5).

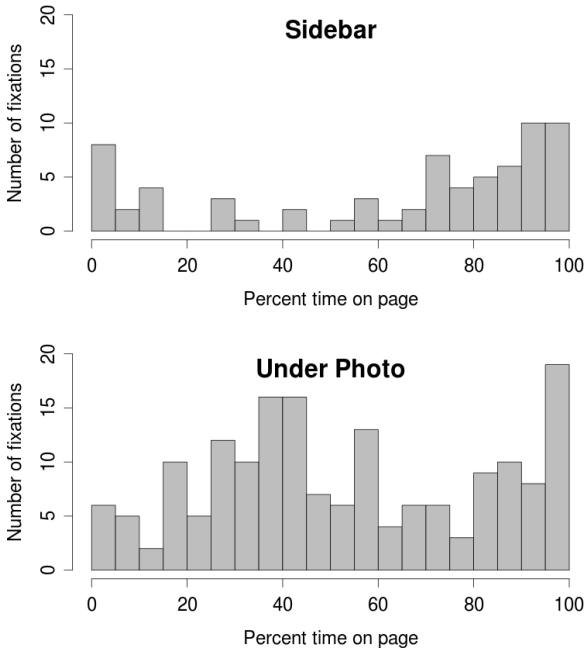

Figure 5. The combined number of times participants fixated on a proximity display on the page before they visited the management interface. The graphs are a combination of all sidebar and under-photo participants. A histogram of the number of fixations as a function of the amount of time the user had been viewing the page normalized based on the total view time on the page.

B. Noticing changed permissions

Three pairs of tasks (6 and 9, 10 and 13, 14 and 17) were designed to investigate whether participants notice permission changes in an album they had previously worked with. In each of these pairs, the participant interacts with album permissions in the first task, and the experimenter changes the album's permissions behind the scenes before the second task. The participant then returns to the modified album to performs the second task. When performing these second tasks, we found that participants rarely proactively read permissions, and instead appeared, based on think-aloud comments, to consult their memory.

Two control participants correctly read and fixed the permissions on one of the three tasks. Four sidebar participants read permissions on these tasks but only one of them realized there was an error. Seven under-photo participants looked at the permissions on one of the permission-change tasks before they were prompted and one of them managed to accomplish all task goals.

In our pre-study we found that some participants did not recognize that it was possible for albums to change, behind the scenes, between tasks. To combat this, we manipulated, between tasks 6 and 9, the titles of photos that participants interact with during those tasks. Task 9 explicitly called out this change to participants. Nevertheless, from comments made by participants, it appeared that participants typically consulted their memory to determine the state of permissions in an album they were visiting for a second time, rather than referring to the proximity displays or full permission interface. We hypothesize that participants often did not appear to consider the possibility that permissions had changed because very few of today's online environments allow multiple people to modify permissions.

These tasks on changed albums simulate real-world situations where a user has interacted with their albums and has a mental model of the permissions, but the permissions do not match the model. Their mental model could be inaccurate for many reasons, such as the user making an error while manipulating permission and the website changing defaults or permission options. That users do not check permissions on albums they have previously worked with, even when checking is easy, means that in this type of real-world situation they are unlikely to successfully identify an error.

C. Permission-modification interface

Many websites place all permission settings on one settings page. This means that whenever a user visits this page they have an opportunity to see and evaluate permissions for the whole site. In this study we altered the default Gallery permission-modification experience to put all the permissions on one page to better emulate the current state of the art and to give our control participants a better chance to notice errors in other albums.

All participants, not just control, used this page as a way to notice and fix permission errors in both the album they were currently working with and others. Participants tended to open this page, start fixing the issue that brought them there, and then notice other permission issues and go fix those, too. Some of these other "errors" were specified permissions that should be changed but many were uninstructed but deliberate changes.

For example: after proactively preventing the group Family from seeing the Building Jumping photos, User 37 said, "And she probably should not see any of these," and turned off Family's ability to view the Sky Diving photos.

We hypothesized that under-photo participants would use the proximity display to notice errors in albums they were not directly working with. Because most participants noticed errors on the full permission-modification interface, it is impossible to know if under-photo or sidebar participants would have noticed errors outside their primary task in the absence of the full interface.

Reeder et al. showed that by displaying effective permissions in a grid format participants were better able to identify and fix errors than in a non-grid rule-based system [8]. It is likely that we are seeing a similar effect in this study.

D. Memory

One of our premises was that users who understand their permission settings will be better able to manage the security of their information. However, in Section IV-B we reported that there was no significant difference in memory between conditions. There are two potential reasons for this: (1) the full-sized policy-modification interface confounded results and (2) the specified policy is artificial.

As mentioned in the prior section, participants in all conditions performed multiple permission-modification actions using the permission-modification page. Due to prompting, even control condition participants visited this page several times during the study and actively engaged with it. This extended exposure may have caused all participants to remember permissions equally.

Each album, or set of albums, had its policy expressed in the emails or written instructions given to participants. We conveyed the policy in this way so that participants could easily reference the description of policy if they were uncertain about it. The artificiality of the policy may have made it more challenging to internalize and remember. We observed that every participant made at least one permission change outside the policies expressed in the emails, suggesting that participants were inclined to make up or interpret policies beyond what was specified.

VI. Related Work

While our study scenarios and ideal access-control policies were synthetic, we endeavored to base them on realistic photo-sharing behavior and concerns. Peahen et al. found that sharing decisions are often related to the people in photos and the environment in which they are taken [12]. Besmer and Lipford also report that concern over "impression management" is a major factor driving concerns about photo sharing [4].

Proximity displays had not been explored in the access-control domain but have been shown to be effective in related contexts. Lieberman and Miller used a proximity display in an email-composition window to show photos of the people to whom the email was addressed [13]. They found that this improved users' ability to notice when the wrong person was being emailed. Wang built a privacy control panel that showed participants the available privacy options and consequences on a book-recommendation site [14]. Participants were not primed to look at the privacy-control interface, but researchers found that 66% claimed to have made a change to their privacy options and 83% claimed to have paid attention to the interface.

An alternative to passive information displays are active warnings. These have been shown to work best in situations where users need to be aware of a potentially dangerous situation and computers can detect the problem with high confidence. Egelman et al. conducted a lab study where they asked participants to make complex changes to their privacy settings [15]. They found that unless participants were given actionable warnings about potentially incorrect settings the participants either never realized there was an error or failed to correct the error. Sunshine et al. showed that passive SSL security indicators were insufficient to notify users and that active indicators worked much better [16]. However, too many security indicators can cause overload, and users then start rejecting the information [17], [18]. When Sotirakopoulos et al. re-ran the Sunshine SSL study several years later, they found that users had adapted and

active security indicators provided less benefit over passive indicators than had previously been found [19].

Tam et al. examined different ways to present access-control information so that users who saw permission settings for only a few seconds could accurately answer questions about them [10]. They found that displays with visual icons were preferred and enabled users to find data more quickly. They also found that performance improved when permissions were organized by action icons.

VII. Conclusion

We examined the effect of positioning proximity access-control displays near photo albums on participants' ability to notice and correct errors with their access-control permissions. We asked participants to complete several tasks with permission and non-permission subtasks. We observed that participants in the under-photo condition, where access-control information was located under each photo and album, performed statistically significantly better at noticing and fixing errors in albums associated with tasks. We also observed that participants in this condition looked at more proximity displays than participants in the sidebar condition.

We believe our study has implications for website interface design for sites where participants' permission preferences are likely to change over time. It is already the case that empowering end users to effectively manage the privacy of the content they put online is a major issue. New social-networking sites such as Google+ emphasize access control as a way of differentiating themselves from competitors. Our study provides guidance to such sites as to effective means of keeping users more in tune with their policies.

The primary limitation of our study is, we believe, that our participants were challenged to configure policies that were not of their own making and for content that was not their own; this artificiality might have influenced our outcomes. Moreover, we found that designing a study to test a secondary task, such as permission management, presents inherent difficulties. Notably, participants had to be made aware of what the ideal policy should be while at the same time not biasing them towards focusing too much on permissions. It would be interesting to reevaluate our findings on users' own content and policies and in longer-term studies involving repeated user exposure to permissions and the effects of time on their memory.

Acknowledgments

This work was supported in part by Carnegie Mellon CyLab under Army Research Office grants DAAD19-02-1-0389 and W911NF-09-1-0273; by ONR grants N000141010155 and N000141010343; by NSF grants DGE-0903659 and CNS-0831428; and by a gift from Cisco Systems, Inc.

References

[1] M. Madejski, M. Johnson, and S. M. Bellovin, "The failure of online social network privacy settings," Department of Computer Science, Columbia University, Tech. Rep. CUCS-010-11, 2011.

[2] Y. Wang, S. Komanduri, P. G. Leon, G. Norcie, A. Acquisti, and L. F. Cranor, ""I regretted the minute I pressed share": A Qualitative Study of Regrets on Facebook," in *Proc. SOUPS*, 2011.

[3] T. Whalen, D. Smetters, and E. F. Churchill, "User experiences with sharing and access control," in *Proc. CHI*, 2006.

[4] A. Besmer and H. Richter Lipford, "Moving beyond untagging: Photo privacy in a tagged world," in *Proc. CHI*, 2010.

[5] L. F. Cranor, "Privacy policies and privacy preferences," in *Privacy and Usability*, L. F. Cranor and S. Garfinkel, Eds. O'Reilly, 2005.

[6] J. King, A. Lampinen, and A. Smolen, "Privacy: Is there an app for that?" in *Proc. SOUPS*, 2011.

[7] W. D. Ellis, *A Source Book of Gestalt Psychology*. Psychology Press, 1999.

[8] R. W. Reeder, L. Bauer, L. F. Cranor, M. K. Reiter, K. Bacon, K. How, and H. Strong, "Expandable grids for visualizing and authoring computer security policies," in *Proc. CHI*, 2008.

[9] R. W. Reeder, L. Bauer, L. F. Cranor, M. K. Reiter, and K. Vaniea, "More than skin deep: Measuring effects of the underlying model on access-control system usability," in *Proc. CHI*, 2011.

[10] J. Tam, R. W. Reeder, and S. Schechter, "I'm allowing what? disclosing the authority applications demand of users as a condition of installation." Microsoft, Tech. Rep. MSR-TR-2010-54, 2010.

[11] "Gallery 3," Website, Mar. 2012, http://gallery.menalto.com/.

[12] S. Ahern, D. Eckles, N. S. Good, S. King, M. Naaman, and R. Nair, "Over-exposed?: Privacy patterns and considerations in online and mobile photo sharing," in *Proc. CHI*, 2007.

[13] E. Lieberman and R. Miller, "Facemail: showing faces of recipients to prevent misdirected email," in *Proc. SOUPS*, 2007.

[14] Y. Wang, "A framework for Privacy-Enhanced personalization," Ph.D. Dissertation, University of California, Irvine, 2010.

[15] S. Egelman, A. Oates, and S. Krishnamurthi, "Oops, i did it again: Mitigating repeated access control errors on Facebook," in *Proc. CHI*, 2011.

[16] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. F. Cranor, "Crying wolf: An empirical study of SSL warning effectiveness," in *Proc. USENIX Security*, 2009.

[17] A. Adams and M. A. Sasse, "Users are not the enemy," in *Communications of the ACM*, 1999.

[18] C. Herley, "So long, and no thanks for the externalities: The rational rejection of security advice by users," in *Proc. NSPW'09*, 2009.

[19] A. Sotirakopoulos, K. Hawkey, and K. Beznosov, "On the challenges in usable security lab studies: Lessons learned from replicating a study on SSL warnings," in *Proc. SOUPS*, 2011.