Using proximity information displays and audit log information to motivate users to view and maintain access-control policy

Kami Vaniea

August 2012

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Lujo Bauer, Co-chair
Lorrie Faith Cranor, Co-chair
Michael K. Reiter
Stuart Schechter
Jeannette Wing

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Abstract

Managing access to shared digital information, such as photographs and documents, is difficult for end users who are accumulating an increasingly large and diverse collection of data that they want to share with others. Current policy-management solutions require a user to proactively seek out and open a separate policy-management interface when she wants to review or change her access-control policy. However, end users treat access control as a secondary task, and rarely visit a website for the primary task of managing security. Historically, security administrators and auditors checked for access-control issues on behalf of users, but in the age of Facebook and Flickr people are responsible for their own content. Users need a way to review their access-control policies that fits into their normal workflows.

This thesis proposes the use of *proximity information displays* — small interface components spatially located near the data elements (or near a representation of data, e.g., file name in a file manager or thumbnail photo in a photo album) that contain information about who has or who could access the data. These displays are intended to help users become more aware of how their data has been used in the past and how it could be used in the future. We present empirical studies that test the hypothesis:

> Users of a system that includes proximity information displays of access control-information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface.

The focus of this thesis is understanding the impact of proximity displays on people's permission-modification behavior. The displays were conceptualized based on interviews with end users and security administrators, which highlighted the need for increased end-user awareness of their policies. Focus groups showed that people liked the idea of showing permission information in proximity to data. Finally, several evaluation studies were conducted in the lab and online using a photo-sharing website. Participants who saw proximity displays that were more comprehensive and could be easily glanced at were better able to identify access-control policy errors. Participants who saw displays that were overly coarse-grained, on the sidebar, or showed information about who had previously viewed the photos, showed no improvement over those who saw permission settings only on a secondary interface. Our studies suggest that proximity displays for access control can significantly help the majority of users, who do not normally check their access-control policies.

Contents

List of Figures

List of Tables

Chapter 1

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 [19, 68, 108, 110]. This problem is becoming even more difficult as sites become more dynamic, with constant uploading of content and shifting friend groups. In this dynamic enviornment, security policies are difficult to not only setup but also to maintain. When the current implemented policy changes due to a new group member or the user's social interactions with a group member evolve, miss-matches in the implemented policy can occur.

Having a mismatch between the currently implemented access-control policy and the policy the user believes to be enacted can place end users in dangerous or awkward situations. If we turn to the news we see numerous accounts of users who set their permissions incorrectly and experienced a loss because of it. A girl in Germany accidentally publicized her birthday party on Facebook and ended up with 15,000 RSVPs, 1,600 of which actually showed up [63]. A teaching student was denied her diploma after a photo of her drinking was shared publicly online [61].

In an ideal world the computer system would analyze user behavior and continuously maintain the access-control policy, dealing with changing environments and preferences. Unfortunately, computer policy management systems are not that accurate yet. While systems do exist that will detect and flag potential issues with access-control policies, those systems are limited by their understanding of what the user currently wants their policy to look like, a preference which can change frequently. Without this knowledge policy error detection systems are prone to both missing important errors and flagging policy components which are error free.

Because programmatically creating and maintaining access-control policy with high accuracy isn't currently feasible, it falls to the end user to periodically check and adjust their policy to meet their current needs. End users' sharing intentions change over time as their social enviornment evolves, and the content being protected changes. Additionally, sites such as Facebook periodically add and remove privacy/access-control settings, effectively altering the access-control policy on behalf of the user. The end result is that even if the user correctly implemented their intended policy using available settings, that policy would likely develop errors over time.

Online users claim that security and privacy are important to them but the reality is

that online users rarely interact with access-control policy as their primary task [29]. 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 the user's attention [34, 108]. So users are unlikely to identify errors in their current settings unless they actively decide to look for them.

Providing end users with usable privacy controls is starting to be seen as a marketable feature by websites built on user content. Social networking sites such as Google+ are trying to use privacy settings as a way of distinguishing themselves to end users. In recent years we have seen these sites move away from placing all the privacy settings on a secondary page, and start selectively putting them near the data element they control.

1.1 Access-control proximity displays

In order for users to better understand the implications of their access-control policy as well as how it is used we need to provide greater transparency to end users. In this thesis I am proposing the use of proximity information displays to make users more aware of how their resources have been used in the past and how they could be used in the future. Proximity information displays (Figure 1.1) are interface components that show users information about their access-control settings and who has been accessing their resources in a way that is easy to understand and enables them to create policies that better match the users' preferences. They are referred to as *proximity* information displays because information is always placed in close proximity to places where the user interacts with or thinks about their resources.

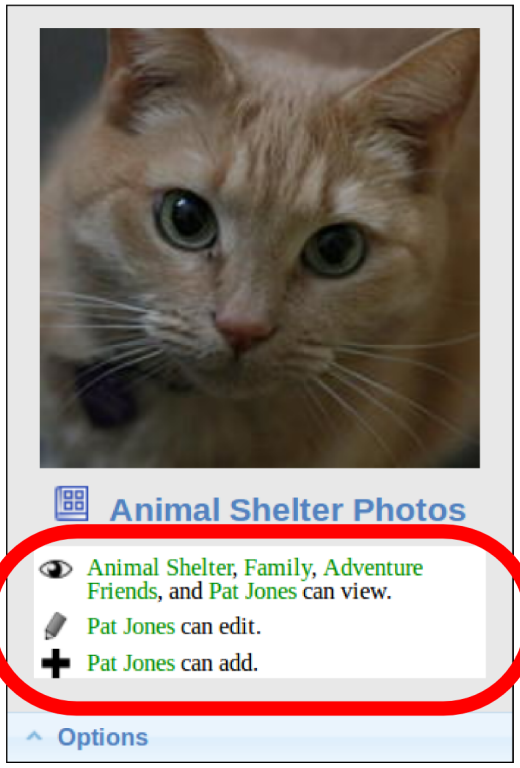

Figure 1.1: Proximity display showing access-control settings under an album thumbnail.

Proximity displays are designed to enable users to notice permissions. They should enable users to 1) identify miss-matches between what they want and the current setting, and 2) become aware of the settings.

Egelman divides the space of security indicators into *passive* and *active* [36]. Active indicators force the user to make a decision before progressing. Passive indicators present information to users but do not force the user to notice or engage with the indicator. Proximity displays are intended to be passive indicators. While many end users' access-control policies do not necessarily match their access-control preferences, programmatically identi-

fying access-control policy miss-matches is error prone with potentially high false positive rates. Proximity displays are intentionally designed to be passive. The end user should be able to easily notice permission errors, while not experiencing a negative impact to their normal workflow.

In this thesis work I explore the design of proximity information displays and the effect the displays have on users' ability to identify issues with, and be aware of their access-control policies.

1.2 Thesis statement

The objective of this thesis is to test the hypothesis:

> Users of a system that includes proximity information displays of access control-information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface.

1.3 Research questions

In this thesis I take an in-depth look at how people notice access-control errors and the impact proximity access-control displays have on that behavior. My work addresses a range of questions intended to support my thesis topic. These questions, enumerated below, express the specific issues I will be looking at in this thesis.

1. How do security professionals manage access-control in organizational settings? Chapter 3
2. How do end users manage access-control? Chapter 3
3. How do people react to access-control setting information being presented on the same screen as their photos? Chapter 4
4. How do people react to information about who has previously interacted with their photos being presented on the same screen as their photos? Chapter 4
5. What is an effective lab enviornment design that enables participants to both understand the goal and still treat security as a secondary task? Chapter 6
6. Do proximity displays improve people's ability to identify access-control permission errors over having the information on a secondary screen? Chapter 8
7. Do proximity displays improve people's ability to remember their access-control permission settings? Chapter 8
8. Do proximity displays negatively impact people's performance on their primary task compared to having the information on a secondary screen? Chapter 8
9. Do proximity displays which show who has previously interacted with the photos improve peoples' ability to identify access-control permission errors over having setting information on a secondary screen? Chapter 8

10. Does the position of the proximity display impact people's ability to notice errors? Chapter 8

1.4 Overview of studies

I present the results of six studies we conducted in order to examine how people react to access-control information placed in close spatial proximity to the item it controls.

1.4.1 Interviewing users

The concept of proximity displays came from an interview study I conducted that explored how both security administrators and end users manage access control in physical and digital security environments. One of the outcomes of this work was the observations that 1) participants were not always aware of the actual access-control policy implemented on the system, 2) participants used a set of tactics to effectively change allow/deny decisions into a finer grained set of effects, and 3) participants manage their access-control policy when not near a computer by using their memory of the current policy to create the effect they want. The interviews highlighted the need for users to be able to determine their current policy quickly, at a glance, and to remember the policy for latter off-line interactions. In Chapter 3 I describe these interviews and the lessons we learned that motivated the development of proximity displays.

1.4.2 Reactions to Access-control proximity display content

To better understand how people would react to different types of information and different presentation methods I conducted a focus group study. The interviews suggested that people had need of a detailed display which gave them concrete information on which to base their mental models of their security policy. The interviews also suggested that presenting detailed information about who had previously accessed their photos would assist users in their continued reevaluation of their policies and social networks. However, participants considered detailed information about who had viewed their photos to be highly invasive because it "forced [them] to stalk [their] friends." Participants were generally positive about showing setting information, provided that it did not take up too much screen real estate. Several users commented about the positive effect of finding and changing permissions easily. In Chapter 4 I discuss the high level take aways from the focus groups.

1.4.3 Proximity information display quantitative and qualitative evaluation

The positive view of focus group participants suggests that proximity information displays which show permission setting information are perceived as useful. However, I wanted to know if these displays are actually useful for participants in terms of assisting them to 1) identify errors in their policies, and 2) improve their awareness of the content of their

policies. To do this I conducted several role play lab studies where I asked participants to come into the lab and work through several tasks while playing the role of a fictitious person who managed an online photo sharing site. This person was responsible for fixing permission and non-permission errors, such as spelling, orrientation, and tags. Participants were told what the access control policy should be for different types of photos in the albums. They were then given a series of emails that requested various changes to the photo albums. In the course of fulfilling these requests they had the opportunity to detect and correct permission errors. I conducted four lab studies using this format to quantitatively and qualitatively understand the effect proximity displays have on participants' permission error identification and policy awareness behavior. In Chapter 7 I detail the methodologies used to test the proximity displays, and in Chapter 8 I detail the combined results of these studies.

Study 1, pre-study: Based on the responses to proximity displays in the prior studies we decided to focus on presenting permission information on the proximity displays and leave the information about who has seen the photographs for a latter study. Based on the focus group feedback I was concerned about the amount of screen real estate required by the proximity display. I was also concerned about the effect of putting the display in too hidden of a location. To evaluate these concerns I tested the display both on the sidebar and under every photo and album thumbnail. The outcome of this study was inconclusive and the study was stopped early due to several methodological issues (Chapter 6), but the behaviors of the participants strongly indicated that showing permission information in close spatial proximity enabled participants to notice errors in their permission settings.

Study 2, eye tracker study: In the pre-study I observed a participant behavior I term "checklisting." Participants who checklisted would appear to finish with a task, pause, go through a check list of all the types of actions we had trained them on, and then explicitly check the permissions. The methodology from the pre-study was redesigned to reduce this behavior by reducing the number of error types, both permission and non-permission, present in each task. I also added some qualitative data collection mechanisms, including an eye tracker, to better capture how participants were interacting with the proximity displays. The result of this study was that placing proximity displays under every album thumbnail and photo enabled participants to identify statistically significantly more errors than placing it on the sidebar or placing access-control setting information on a secondary page. I also observed that participants who see proximity displays under the photos tend to see the displays mid-way through the task, but change the permissions at the end of the task.

Study 3, lab study: In the prior study I saw a statistically significant difference in the number of permission errors identified, but I did not see a difference in participants' ability to remember the permissions. Results from the interview study indicated that participants reason about their security settings off line, and make decisions that depend on their memory of their settings being correct. In addition to enabling participants to find permission errors, I also wanted to make them more aware of their current settings. I hypothesized that the lack of difference in memory in the prior study was caused by 1) forcing control participants to repeatedly access the permission modification interface and, 2) providing participants with a permission modification interface which showed the policies

for all albums, not just the album the participant was currently working with. In this study I decided to test the style of the permission modification interface used, in addition to the proximity display. I also increased the amount of qualitative data collection by adding a post-study interview where I used a cognitive interview approach to ask participants about the choices they made during the study. I found that the permission modification interface used impacts participants' ability to notice errors, but had no impact on memory. I also learned that participants were able to glance at the displays and some participants had a natural tendency to correct all permission errors in one single pass.

Study 4, online study: The prior studies indicate that proximity displays help people identify permission errors. However, these studies were done with a small number of participants. The results of the prior study also highlighted the high level of participant variability; some participants are more inclined to check permissions than other participants. To address this, I conducted a within-subjects study — every participant saw the control condition and one of the proximity-display conditions. I also increased the number of proximity-display conditions including proximity-display designs that mimic the Facebook proximity-display design and a proximity display that contains information about who has seen the photo album. I found that conditions that used proximity displays that showed permission setting information under photos/albums or under album thumbnails and on the sidebar were statistically significantly better than control at enabling participants to identify errors. However, similar to the prior studies, I saw no difference in memory.

1.5 Outline of the thesis

In this thesis I start with a discussion of the related work (Chapter 2), discussing both what is currently known about the way people manage access-control, and systems similar to proximity displays. Then I discus the interview study that motivated the need for proximity displays (Chapter 3). This is followed by a focus group study to explore peoples' reactions to variations in proximity display content and design. The details of the proximity display design and implementation in the Gallery photo sharing system are described in Chapter 5. Designing the methodology for the four studies that tested the effectiveness of the proximity display was an informative experience with several lessons learned (Chapter 6). I detail the methodologies of the last three studies (Chapter 7), then describe the results (Chapter 8). Finally, I conclude with a discussion of the effect, future work, and design recommendations for proximity displays (Chapter 9).

Chapter 2

Related Work

Making the process of managing access control usable is a difficult and important problem. In 2003 the Computing Research Association released a list of four Grand Research Challenges including: "Give end-users security controls they can understand and privacy they can control for the dynamic, pervasive computing environments of the future" [50]. The National Academy of Engineering agrees, listing cybersecurity as one of their Grand Challenges and specifically noting that understanding the psychology of computer users is a vital component of improving the state of cybersecurity in general [75].

> "Security features in IT systems are, in a sense, like brakes on automobiles. Although brakes are used to slow or stop vehicles, their real purpose is to enable drivers to go faster by enabling them to avoid accidents caused by external threats (such as mechanical failure in other vehicles, rude or reckless drivers, road hazards, stop signals and heavy traffic). Better security is an enabler for greater freedom and confidence in the cyber world" [50].

Systems that allow end users to configure privacy settings may be thought of as access-control or security systems, as they involve policies that govern access to a user or to a user's personal information. In this thesis we will be using the terms *privacy settings* and *access-control settings* interchangeably to refer to the set of settings a user can manipulate to control who can see what part of their information. The term *implemented* is used to refer to the current state of the access-control settings on a system. The term *ideal* refers to the access-control policy the user would like to implement on the system.

In this chapter we will discuss what the research community currently knows about how end users manage implemented and ideal access-control policies. End users currently have difficulty managing their implemented policies, which results in negative consequences for users (Section 2.2). Researchers are designing systems that better support users in their access-control management tasks (Section 2.3). Finally, there exist models that explain end user behavior towards warnings in general and warning specific to computer security (Section 2.4).

2.1 Placing security information in spatial proximity

Several studies have looked at the effect of showing privacy and security related information in proximity to items the user is working with [36, 65, 102, 103, 107]. These studies generally show that displaying security and privacy information in proximity to related items can positively influence end user behavior provided the user understands the import of the information.

Tsai et al. showed that placing a graphical representation of each websites' privacy poilcy next to search results increased the amount of money people were willing to pay when purchasing privacy sensitive items [102, 103].

Lieberman et al. was concerned about the imact of accidentally emailing too many people thorugh lists or the "reply to all" feature [65]. He designed an email interface to enable people to easily spot errors, the interface showed the photos of all the people being emailed near the box for email addresses. He saw a significant improvement in people's ability to quickly identify who was being emailed.

Wang, in his thesis work, displayed a large privacy related proximity display on the side of a fictional book selling website [107]. He observed that participants self reported interacting with the display and liking the options.

Egelman divides the space of security indicators into *passive* and *active* [36]. Active indicators force the user to make a decision before progressing. Passive indicators present information to users but do not force the user to notice or engage with the indicator. Sunshine et al. showed that passive indicators were less effective than active indicators in alerting users to the dangers of self signed certificates [100]. Sotirakopoulos et al. repeated the sunshine study and found that users were ignore both active and passive indicators [95]. The authors hypothesize that users were now more familure with the active indicators and were now habituated to ignoring them.

Kelly et al. attempted to use eye tracking to better understand how people, particularly experts versus non-experts, looked at Facebook's proximity display icons while normally interacting with their own Facebook profiles [58]. Their study was plagued by issues related to the eye tracker technology used. However, their preliminary results show that some users do look at the access-control information.

2.2 Users' policy management

End users have trouble managing permissions in their online social environments. They are not aware of what their current permission settings are, and they incorrectly assume permissions to be correct when errors exist. They do not regularly check the permissions for errors or adjust their policies when their context changes. This can lead to a loss of privacy for individuals, resulting in potential harm and embarrassment. We take a look at how people interact with security at home and at work.

2.2.1 User awareness

Users lack awareness of their current policies whiling inaccurately assuming that the policies are fine. An empirical study of Facebook users compared participant's sharing intentions to the implemented privacy policy, and found that every participant they interviewed had at least one mismatch [68]. A 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 [59].

Whalen et al. conducted an online survey on end-user experiences with sharing and access control. They found that users have dynamic access-control needs that vary with task and are often frustrated by current access-control mechanisms that are difficult to use and not well-suited to users' workflow [110].

Research in the area of photo access-control management shows that end users care about the privacy of their photographs. Similar to other data sharing domains, end users claim to care about privacy but have difficulty managing it. Ahern et al. [6] found that sharing decisions are often related to the people in photos and the environment in which they are taken. Besmer and Lipford [19] also report that concern over "impression management" is a major factor driving concerns about photograph sharing.

2.2.2 Policy re-evaluation

As part of interacting with others, people continuously construct, interpret and reevaluate their social context based on actions others take [76, 80]. In file security the actions of others are often hidden by the system, and even the settings are placed on secondary pages where they are not easily visible. Without this visibility it can be challenging for users to take the access-control permission settings into account as part of their natural social reevaluation [17]. Users may not realize that their access-control policy no longer accurately represents what they want until something happens to bring it to their attention [34]. Prior work in domains such as location tracking and photo management tells us that end-users' policies can be very dynamic and are often based on the current relations between the user and the requester [53, 60, 72, 77, 96]. However, when users set static policies to dictate who can see them where, unanticipated or out of character access patterns cause requests to be denied [27]. Studies of home and Internet centralized file storage environments also show that end users start out by creating one access-control policy and, based on observations of how it is used, they may realize that an alternative policy is more appropriate [67, 71, 88].

2.2.3 Managing permissions in the home

Home users tend to view their immediate threat model as non-malicious [71]. When working in small groups, people establish social rules that allow them to function without tight security. These rules work as long as the group is small but break down in larger settings [9]. Home users trust the other members of their home and expect them not to pry beyond clearly marked boundaries. Instead of using technology to protect their files, users hide the files or put them on clearly marked personal devices or in personal spaces [71, 87].

People tend to think about physical and virtual security holistically, not separating the two concepts [34, 71].

The home is not a structured environment, where each person has her own account and all her data and accesses are tightly tied to the account. Home users tend to share a single account on the "family computer" preventing a clear tie between account and person [25, 87]. Account sharing is primarily driven by convenience issues, being able to quickly access the computer outweighs the privacy and security concerns having multiple accounts solves [37]. Users also have dynamic access-control policies that can change quickly [13, 66]. Using focus groups on ubiquitous shopping technology, Little, Silence and Briggs found that users are concerned that computer controlled devices cannot properly respond to the unpredictable day-to-day behavior changes of the home environment [66].

Users are not necessarily skilled at managing their computing resources on their own and tend to seek help from trusted people when they need it [34, 78]. Once a trusted person has been consulted, the user tends to blindly believe that the device is now secure even if the trusted person is no longer present [34]. Users also appear to learn about "correct" security behavior from stories told to them by other users. These stories allow users to learn from the negative experiences of others [80]. Users create heuristic rules about the types of data that are stored on different devices and therefor accessible to different sets of people. These rules are rarely if ever updated [88].

2.2.4 Managing permissions in an organizational setting

Organizations put a large amount of time and money into the creation and maintenance of internal file systems which track important documents, preserve confidentiality, and ensure security protections but, unfortunately, do not encourage end users to engage in secure sharing behavior. Many of these systems prevent internal organization employees from reliably sharing files with others and themselves [110]. The lack of reliable file sharing support forces users to circumvent the perceived pointless impediment of the file system and turn to alternative file sharing technologies such as email, instant messaging (IM), third party storage (Dropbox), and USB drives [32, 73]. These alternative sharing mechanisms enable users to quickly and easily share the document with whomever they want but they lack many of the security properties of the original file system.

Unlike the home environment, where the number of users is small and the assumption of non-malicious users may be reasonable, the office environment can be large and contain malicious users. Schneier writes [90]: Access control is difficult in an organizational setting. On one hand, every employee needs enough access to do his job. On the other hand, every time you give an employee more access, there's more risk: he could abuse that access, or lose information he has access to, or be socially engineered into giving that access to a malfeasant. So a smart, risk-conscious organization will give each employee the exact level of access he needs to do his job, and no more.

Malicious employees are a major concern for organizations. A study by CERT of 49 insider attacks found that 59% of the "insiders" were former employees and 43% still had authorized system access at the time of the attack [8]. A TELUS Security Labs study found that 33% of security breaches reported in 2009 were due to insiders. Insider breaches were

reported by 17% of Canadian organizations.

Not all insider attacks involve data breaches. Some insider attacks are just people making use of resources the company provides in ways the company disapproves of. Dwayne F. Cross, a government worker, was convicted of computer crimes for looking at over 150 passport files. His reason: curiosity [45].

Another challenge is that members of the information technology (IT) field often perceive end users as insecure and the cause of many security incidents [56]. End users similarly view security professionals and even their own IT departments as being overly paranoid and generally getting in the way of the work end users need to get done [5, 46, 49], and end users aren't always wrong in assuming that dealing with extra security tasks isn't worth their time [49].

A few studies have surveyed needs for access-control systems from a holistic organizational perspective. Ferraiolo et al. studied the access-control needs of 28 commercial and government organizations and identified seven access-control approaches. One approach they discuss is *discretionary access control* (DAC), in which access is assigned to individuals and groups, who in turn may delegate that access to others. The authors note that DAC is well suited for organizations where end users have rapidly changing information access needs and must be able to specify access-control policy for resources they control. Although DAC is usually implemented through *Access Control Lists* (ACLs), the authors point out that when these ACLs are centrally administered they "can become clumsy and difficult to maintain." They also note that the DAC approach is not suitable for organizations concerned with maintaining tight controls on access rights [41]. The introduction of Role Based Access-Control (RBAC) [40] was partially intended to address this issue by making the setting of access-control permissions better fit how organizations actually manage their settings.

2.2.5 The social statements access-control settings make

Smetters and Good examined the acquired access-control rights of employees in a large office environment. They found that access rights tended to be collected over time at the company and treated as a status symbol [94]. Sinclaire et al. observed a large financial institution during an entitlement review of its employees' current permissions to resources within the company including file permissions. As part of the review auditors asked employees to review their own access to files and applications and remove permissions to resources they didn't actually need. Employees voluntarily removed 15% of their own access permissions because they "'just didn't want to worry about' having access to applications they didn't need" [92].

In addition to access-control being viewed as a status symbol, too much focus on keeping things secure can be viewed as paranoia. Gaw et al. studied a non-profit organization where maintaining security was an important part of employee's job descriptions [42]. The found that employees only used secure communications for important documents and not for other types of communications. This was partially because doing things like encrypting all email was perceived as paranoid.

2.2.6 Automatically create policies

Theoretically, the best way to assist end users in their permission modification is to automate the problem away. If computers could automatically determine the correct policy and just enact it with a high degree of accuracy, our problems would be over. Unfortunately, we do not yet have a system capable of reliably predicting the correct access-control permissions and enacting them on our behalf. Researchers have endeavored to study and predict our access-control preferences [31, 60, 86].

Fischbein et al. found that users' preference for sharing or hiding their location information varied across time even when the location, time of day, and requester remained the same [18]. They hypothesized that the changes were due to contextual factors not visible to the system. Sadeh et al. found that end users' were able to specify policies that matched their *ex-post* preferences only 79% of the time [86]. Cranshaw et al. used machine learning with user feedback and was only able to accurately match end users' *ex-post* preferences at best 87% of the time [31].

There have also been attempts to use an Attribute Based Access-Control model [54] to automatically create rules based on pre-existing attributes. Klemperer et al. explored the use of photo tags, combined with user specified rules, to manage access-control policy for photos [60]. They found that using organizational type tags resulted in 27% of photos being erroneously marked as allowed or denied for at least one "friend." Though only 7.8% of friend, photo combinations were erroneously allowed or denied access.

Researchers generally view full automation of access-control policy creation as unlikely to happen soon. Partially due to the high false positive rates described above, but also because of the issue of exceptions and emergencies. As Rissianen et al. says, there can be many different situations in which an access request could be made and only some of those situations are possible to anticipate [85]. As Edwards et al. points out, automating security enforcement may not always be beneficial [35]. Other researchers have similarly observed that users plan ahead for exceptional or unanticipated situations and need an access-control system capable of supporting this type of forward thinking and planning [12, 14, 24, 32, 79, 97, 99].

If we accept that full automation is unlikely to happen in the near future, then we must rely on end users to actively be involved in the creation and maintenance of their own access-control settings. Users need computer systems that enable them to manage security as part of their workflow. Researchers have proposed several systems which enable users to manage security as part of their normal system interaction.

2.3 Enabling access control management

In access-control literature we tend to think about access-control policy specification as a user task and policy enforcement as a computer task. Stevens and Wulf coin the term *Computer Supported Access Control* or CSAC to emphasize that the technological mechanisms behind access control are only one part of how access control as a whole is practiced. They argue that access control should be designed as a supporting system, where humans

work with computers to actively manage how files are accessed, instead of an automation system, where the computer automatically enforces policy without additional input from the user [98].

In his work Lampson described the task of setting access-control permissions in terms of proactive permission setting and automated enforcement. In other words he assumed that permissions would be set ahead of any access attempts and that the system would be solely responsible for judging the appropriateness of the request based on previously expressed access-control lists and enforcing it [62]. These assumptions that access-control is set before the access and that enforcement should be automated are common in the security community [40, 43, 52, 70, 89].

Stevens and Wulf [98] and other researchers [12, 79, 85] have postulated that access-control management tasks are actually conducted in one of three ways. *Ex-ante control* is when the resource owner sets the policy before any anticipated accesses occur and the computer enforces it at the time of the access. *In-medias-res control* is where the access permissions are defined by the resource owner at the time of the access attempt. Finally, *ex-post control* is where the computer automatically grants access and the legitimacy of an access request is evaluated by the resource owner after the access has already taken place.

Researchers have looked at many different ways to assist users with their permission modification tasks. While there are many different ways to assist users, the approach taken by researches depends largely on how they assume users will interact with their technology. In this section we discuss different solutions proposed by researches to address users who manage their access-control policy *ex-ante*, *in-medias-res*, and *ex-post*.

Ex-ante control

Users engaging in *ex-ante* control create their access-control policies pro-actively in advance of any access attempt based on how they anticipate the resource will be used in the future. The policies are then automatically enforced by the computer system that is responsible for interpreting the policy expressed by the user based on the current context.

Traditionally, end users interested in engaging in *ex-ante* control pro-actively seek out an access-control management interface and use it to specify the policy. In Windows XP, for example, a user must pro-actively right click on a file and select the "Sharing and Security" option before they can see or modify the file's policy. Many different researchers have looked at how to support this type of access-control policy management [23, 82, 83, 105].

Johnson et al. [55] built a system where end users could, through their email client, upload a document to a document sharing system, automatically grant access to all email recipients, and include a link to the document in the email instead of the document itself. Though they were never able to get the system to a fully deployed state, the researchers were able to observe participant's positive reactions to it. As a result of this work they put forward the idea of Laissez-faire access-control [55]. Similar to De Paula et al. [33], the Laissez-faire work proposes that access-control needs to be less restrictive, fit naturally into the end user's workflow, and better match current behavior where access control is more continuous and less about definitive *allow* and *deny*.

Brodie, Karat, and Karat built and tested the SPARCLE natural language policy man-

agement interface [23]. SPARCLE assists knowledge workers in writing machine-readable natural language privacy policy rules in a guided environment [57]. Vaniea et al. explored the use of syntax highlighting in the SPARCLE interface. They found that when writing rules in natural language, end users need the interface to expressly support the planning/translating and revising tasks normally associated with natural language writing [105]. Using the SPARCLE system, Reeder et al. identified five general usability challenges that policy authoring systems must address to be considered usable. Amongst these challenges are: 1) making default rules clear, 2) communicating and enforcing rule structure, and 3) preventing rule conflicts [82].

Maxion and Reeder observed that, when interacting with access-control permissions, end users rarely care about the individual rules and instead want to see the effective permissions [69]. Effective permissions are the result of considering all relevant access control rules together to determine whether access will be granted or denied. Maxion and Reeder designed the Salmon system, which showed users the effective permissions and how those permissions were computed. They found that users who used the Salmon system were better able to perform basic policy management tasks, such as giving a person access to a file [69].

Reeder et al. introduced an interface paradigm for access-control policy management they call the Expandable Grid, which gives users both a high-level view of all the effective permissions in a system and the ability to drill down and examine any particular permission [83]. They found that end users using the Expandable Grid to perform basic policy management tasks, such as give Bob access to fileA.txt, were faster and more accurate than users who used the default Windows interface for policy management [83]. Further exploration by Reeder et al. found that the conflict resolution strategy used by the system to compute effective permissions had a significant effect on end users' ability to accurately make policy changes [15].

The Grey system [11, 12, 13], similar to Beaufour and Bonnet's proposed personal servers with digital keys system [16], is an implemented distributed discretionary access-control system that was constructed and studied in a live environment. The Grey system enables end users to manage access-control in a discretionary way, while maintaining strong audit. Every access to a resource requires a certificate based proof that access should be allowed, thereby ensuring that the logs contain both the access attempt itself and details about why the access was allowed. The system is distributed in that the certificates and proof statements are all developed and stored on smart phones, so no central server is required. The mobility of the devices enables end users to make, and change their policy from anywhere, with little effort. A within-subjects study of Grey users found that they created more restrictive access-control policies using Grey than with the physical key system they had used previously. The study also found that Grey users were more able to easily change their policy and this resulted in them giving out less access "Just in case." However, one of the issues with such a system is that more than one person can change the implemented access-control policy, without necessarily informing other people who have access. This observation was one of the motivations for this thesis.

Proximity information displays are partially intended to provide additional support for *ex-ante* control that is not provided by existing technologies. Unlike existing policy

management solutions, proximity information displays will provide end users with a passive way to review their existing access-control policy without having to proactively locate a policy management interface. By providing end users with information about who could access their resources I will provide them with an easy way to engage in *ex-ante* control. To my knowledge, there is no existing work which examines placing policy information on the interface to encourage users to engage in *ex-ante* control.

In-medias-res control

In-medias-res control, otherwise known as uno-tempore control by Stevens and Wulf [98] and reactive control by Bauer et al. [12, 13], is somewhat less studied. Stevens and Wulf [98] define uno-tempore control as "The permission is defined at the moment of the access attempt." Bauer et al. describe reactive policy creation as any policy decision made in reaction to a access attempt or request [13].

In-medias-res control is performed on a case-by-case basis for a specific access in a specific context. Unlike users engaging in *ex-ante* control a resource owner participating in control has understanding of the context under which the access is taking place and potentially even knows the reputed purpose of the access [12, 53].

In *in-medias-res* control there is little to no automation on the part of the system. An access request is not approved by the system, instead it is manually or automatically forwarded to one or more users who decide the outcome which the system enforces. Alternatively, a request could be created out-of-band which the resource owner responds to by creating permanent or temporary permissions. Bauer et al. created a physical access-control system called Grey that allows end users to directly request access to offices from office owners who can choose to either allow or deny the request [13]. They found that end users made use of this functionality to manage offices that are accessed only occasionally by people other than the occupant. Mazurek et al. also explored having end users approve or deny access to their files in real time [72]. They found that users responded differently when asked to describe their policy *ex-ante* than when they were asked at the time of the access request (*in-medias-res*).

Another type of *in-medias-res* control is the creation of temporary permissions that can only be used only once or for a short time period. Whalen et al. observed a need for granting temporary access to files [110]. Bauer et al. also observed people using *in-medias-res* control to give others one-time or temporary access to an office [12].

Proximity information displays are not intended to support *in-medias-res* control. *In-medias-res* control requires that the resource owner be notified in a timely manner. Because proximity information displays are spatially located on the interface near the resources they refer to, there is no guarantee a user will be looking at them at the time when *in-medias-res* control needs to be performed. This makes proximity information displays an inappropriate medium to encourage *in-medias-res* control. However, proximity information displays can help end users engaged in *in-medias-res* control by making it easier to locate the policy modification interface to make changes. Existing research shows that when trying to solve a problem users tend to start at the problem source, the resource, and iteratively search outward for a way to solve it [74, 114]. I am unaware of existing research which looks at

the effect of placing links to security controls in close proximity to resources.

Ex-post control

Ex-post control is defined by Stevens and Wulf as "Permissions are checked after access was granted" [98]. In *ex-post* control the resource owner sets little to no access-control restrictions *ex-ante* and instead relies on accountability to ensure that resources are used in a responsible way. The system logs the details of each access. The resource owner then reviews the accesses after they have happened.

Ex-post control has several major advantages. If the resource owner is in an environment where the majority of users are trusted, managing individual permissions may take more effort than it is worth. As Zhao and Johnson observe "rigid access control delays an organization's response to the changing markets, resulting in missed opportunities or degraded service quality" [115]. Engaging in *ex-post* control allows the resource owner to let other users use their good judgment and quickly gain access when access is needed. *Ex-post* control also allows the resource owner to evaluate the appropriateness of an access once all the facts are known. As Blakely suggests "make users ask forgiveness, not permission" [21].

Similar to *ex-ante* control, in *ex-post* control the system is responsible for automatically granting access based on a previously expressed set of preferences. The difference is that in *ex-post* control the resource owner takes an optimistic view and gives access to all people who might ever need access. The system is responsible for automatically enforcing this policy and the resource owner is responsible for manually reviewing the appropriateness of each previously allowed access.

Ex-post control is based on the observations that end users do not always know who should or should not have access to which resources in the future and that end users have limited time to manually approve and deny every request. In their work Jaeger, Edwards and Zhang look at the permission assignment state of individual users in terms of actions that are expressly allowed and actions that are expressly denied. They found that often a significant portion of the access-control space has neither an express allow nor an express deny defined [52]. Rissianen, Sadighi and Sergot took this observation one step further and applied the idea of access-control spaces to policy creation and enforcement. In their work they argue that unanticipated and unenforceable policy should be enforced with an "Allowed - with override" policy that is enforced via *ex-post* control [85]. Stiemerling and Wulf expanded on this idea by building negotiation functionality into their group-ware document sharing tool. The tool notified users when specific documents were changed and gave the users a technological medium for negotiation and resolution if the change was unacceptable to someone [99].

Stiemerling and Wulf observe that in multi-user collaborative environments, users have need for more complex policy controls than a simple allow or deny. They observe that in collaborative work environments people have to access each other's files while at the same time respecting the other person's privacy. In their work they look at how people use awareness, trusted third parties, and negotiation to handle situations that are unforeseeable or simply outside the abilities of the access-control system to specify. They then add a

tool to their group-ware software that allows users to create complex conflict-negotiation rules. The negotiation system allows another user to override the existing access-control permissions provided several requirements, such as an email to the owner or n number of people agreeing to the override, are met [99].

Polvey introduces the concept of "optimistic security" in which the resource owner places few, if any, policy restrictions on the resource and instead relies on accountability and the ability to roll back the system to ensure integrity of the data. In optimistic security the potential accessors are considered somewhat trustworthy, and it is assumed that the majority of accesses will be "good." If one of the users performs unacceptable accesses the resource owner has options for recourse via system roll-back and change accreditation. So, while another user can freely read and make changes to resources the resource owner can easily attribute each change to the person who made it and they can easily return the system to a prior state [79].

Gutierrez et al. proposes a system where end users can negotiate the amount of tracking information visible to the content owners of pages the user visits [47]. Content owners would explicitly set the level of audit, information required to view each piece of content, in this case: detailed information, anonymous information, or no information collected. Users also explicitly state the level of collected audit information they find acceptable to be visible to content owners. Each user will only be shown content that matches both their and the content owner's preferences. Users who are more willing to give up privacy can see more content, and owners who are more willing to display content without tracking will display that content to a wider audience. Though built, this system was never tested with end users.

Proximity information displays are intended to support *ex-post* control by providing information about who has been using which resources. The information allows the resource owner to casually perform an *ex-post* review of the accesses and determine if anything unacceptable is going on without having to proactively open a dedicated interface. There has been limited research on how to construct interfaces which support *ex-post* control and the majority of that research has looked at multi policy author environments.

2.4 Behavioral Models

Research from cognitive psychology, behavioral economics and the warning sciences provide useful information on how humans react to and think about different situations. When designing interfaces that are intended to alter end-user behavior and may only be viewed for a few seconds, it is valuable to consider how users will process and think about the interface and the information it presents. In this section I will talk about the Communication-Human Information Processing Model (C-HIP) [112] which describes how humans process warning information. I will then discuss an expansion of the C-HIP model called Human In The Loop Framework (HITL) [30] which adapts many of the central principles of C-HIP to the computer security domain.

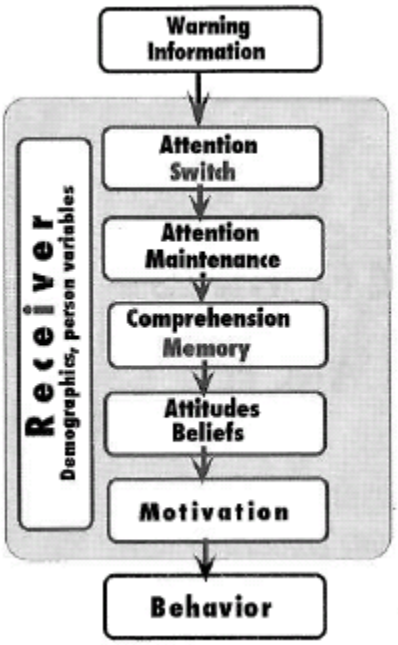

Figure 2.1: Communication-Human Information Processing Model (C-HIP).

C-HIP Model

Wolgalter proposed the Communication-Human Information Processing Model (C-HIP), pictured in Figure 2.1, for structuring and discussing research about warnings. C-HIP model is useful to understanding how people process presented information in terms of noticing it, understanding it, deciding if it is important, and finally doing something about it. According to the C-HIP model, interaction begins with the display of a warning through a channel to the end user. Once the warning has been delivered via the channel there are several stages the user can go through. Each of these stages is described below.

Attention Switch In the initial stage the warning needs to get the user's attention by getting them to look at the warning. To do this switch the warning must be designed to be noticeable. It also needs to be positioned such that it can be noticed by the user.

Attention Maintenance Once the user has switched attention to the warning their attention needs to be held long enough that they acquire the information presented by the warning. Legibility and form factor can have a strong influence on attention maintenance. If the warning looks difficult to read or unclear the user may not dwell on it long enough to attain the needed information.

Comprehension Memory Even if the user looks at the warning for a long enough time they may still not be able to internalize the information from it if they are unable to comprehend it or if it fails to activate relevant information from memory. For example a "Warning! May contain musca domesstica" sign is useless to someone who doesn't know that musca domesstica is the scientific name for the common house fly.

Attitudes and Beliefs A fully comprehended warning may still fail in its purpose if

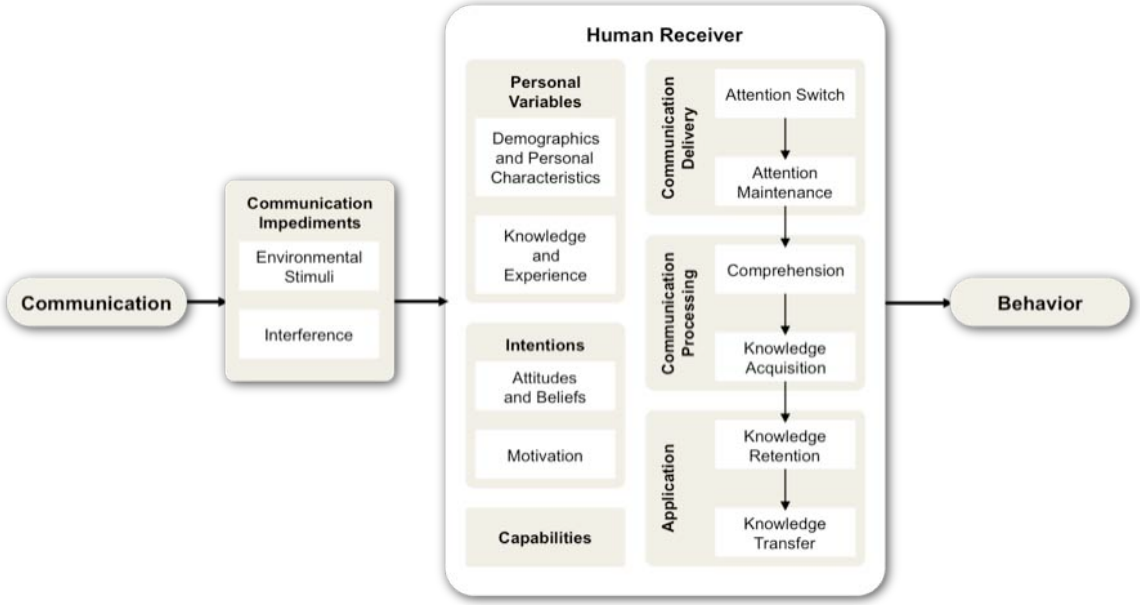

Figure 2.2: Human In The Loop Framework

it fails to adequately influence the user's hazard-related attitudes and beliefs. Beliefs and attitudes form the user's current mental frame-of-reference based on the user's experiences. For example a "your files are visible to all people on this computer" warning may be ignored by someone who believes that no one would ever go looking for their files.

Motivation In the final stage the user is either energized to engage in behavior appropriate to the warning or they are not. A motivated user will move on from this stage to the *Behavior* stage where they engage in a behavior appropriate to the warning.

HITL framework

The Human In The Loop (HTL) Framework proposed by Cranor [30] expands and adapts the C-HIP model to the domain of computer security. The work also postulates that through the HITL framework lessons from C-HIP are applicable to additional communication mediums including notices, status indicators, training communications, and policy communications.

Cranor also introduces the concepts of *knowledge retention* and *knowledge transfer*. If a user, through a communication, learns about a particular hazard are they likely to remember that the hazard exists in the future when they encounter it again? If the same user encounters a similar hazard are they able to apply the lesson they learned from the warning in a new domain. For example if a person uses proximity information displays and learns that all of ProjectA is world readable, will they remember that fact latter when they try and save secretFile.txt to ProjectA?

Chapter 3

Security administrators interview study

The research for this thesis began with a an attempt to better understand how access-control is managed inside of organizations [1]. This research was conducted as part of a larger set of studies intended to understand how people approach access-control and if a system like the Grey smartphone based physical access-control system we were designing would actually help solve existing problems. The purpose of this study was to better understand how security administrators, managing both digital and physical systems, managed security in their organizations. The results of this study lead to a better understanding of how access-control is managed and the every day issues people have. The proximity displays proposed by this thesis are intended to address some, but not all, of these issues.

Effectively controlling access to resources within an organization is a challenging problem for access-control policy professionals. Making sure the correct person has access to the correct resource at the correct time often requires communication within and between departments. Access-control policy changes are needed when employees are hired, terminated, or change job roles. Temporary changes may be needed when employees are given temporary assignments. Changes to specific access-control policies may be needed when company-wide policies change. The introduction of new computer systems or the retirement of old systems, as well as changes in physical office space are other reasons for access-control policy changes.

A study of 23 insider attacks found that "in 78% of the incidents, the insiders were authorized users with active computer accounts at the time of the incident. In 43% of the cases, the insider used his or her own username and password to carry out the incident" [81]. A related study of 49 insider attacks found that 59% of the insiders were former employees and 43% still had authorized system access at the time of the attack [8]. These findings indicate that current systems may be inadequate at supporting policy professionals' needs, even for routine tasks such as revoking access when employees leave an organization.

In this study we sought to understand the challenges policy professionals face in their

[1]This chapter is based on a previously published paper: L. Bauer, L. F. Cranor, R. W. Reeder, M. K. Reiter, and K. Vaniea, *Real life challenges in access-control management*, Proceedings of CHI, 2009

Pseudonym	File or Physical	Organization	System Managed
Ann & Kristen	Physical	University A	Department-wide swipe-card, physical-key and key-pad systems
Henry	Physical	University B	University-wide swipe-card system
Tony	Physical	University B	Department-wide swipe-card system and physical-key management
Kevin	Physical	University B	Department-wide swipe-card, physical-key and key-pad systems
Fred	File	University B	Department-wide Windows and Unix-like file systems
Jerry	Physical and File	University B	Physical-key and electronic systems for a lab
Sue	Physical	University B	Department-wide physical-key system
Seth	File	Organization A	Organization-wide file system
Ralf	File	Organization A	Organization-wide file system
David	Physical	Organization B	Organizational-wide swipe-card and physical-key system
Beth & Sara	Physical	Organization C	Department-wide swipe-card and physical-key systems

Table 3.1: List of the interviewees, the type of system they worked with and the role they played in managing the access-control policy for that system. All interviewees are referred to by pseudonym.

daily tasks. We focused on understanding to what extent policy management technology was successful or unsuccessful in helping policy professionals meet these challenges. To do this we conducted 11 interviews with 13 policy professionals in 5 organizations.

The data from these interviews lead to three key findings. First, we find that policy professionals take different roles in creating policy—some are high-level policy architects, others are implementers of policy designed by others. Current policy-management technology does not acknowledge this distinction, and hence fails to provide tools specifically suited to each role. Second, we find that policy is often jointly managed by several people rather than a single individual. Although technology sometimes aids these individuals in coordinating their activities, such tools are typically poorly integrated with the mechanisms for creating and manipulating policy. Third, we find that some commonly desired policies cannot be fully enforced with the access-control mechanisms that are used to implement them, leading to cumbersome workarounds.

In the remainder of the chapter we discuss the methodology of the study, the results of our interviews, and how they support each of these key findings. For each key finding we suggest ways in which technology that supports policy professionals could be improved to better match the needs of its users.

3.1 Methodology and data analysis

The study was designed to elicit an understanding of the challenges access-control policy professionals face and how current technology helps them meet these challenges. We began the study with no hypothesis and through the interview and data analysis processes we incrementally constructed theories concerning access-control management.

3.1.1 Interviewees

Interviewees were recruited using existing contacts. We interviewed thirteen policy professionals from five organizations. In two cases two professionals who shared the same job function were interviewed together. Eight of the interviews were with policy professionals

who manage a physical-access-control system and three were with policy professionals who manage access control for a file system.

We purposely chose to consider both physical and file access control in our work because both use increasingly similar computer-based management interfaces. Every administrator who was interviewed worked with at least one access-control system that had digital components and was administered using a computer interface. Additionally, researchers are starting to create computer technologies to solve problems with physical security systems [11, 16], further eroding the line between physical and file access-control systems.

Prior work indicates that, in some organizations, responsibility for administering access-control policy tends to be delegated, with a central authority delegating responsibility to department administrators, who in turn pass on the responsibility to other people in the department [11]. To better understand the needs of professionals at different levels of an organization's hierarchy, we specifically selected participants from multiple levels of the organization.

3.1.2 Organizations

We use pseudonyms to identify the universities, organizations, and policy professionals discussed in this chapter.

University A is a public university that has approximately 37,000 faculty, staff and students at its main campus. We interviewed two administrative assistants, Ann and Kristen, who manage physical access control in their department using a swipe-card system, physical keys and key-code pads. Their department contains roughly 150 faculty, staff and graduate students. They also interact with undergraduate students, who have a high turnover rate.

We conducted separate interviews with six policy professionals at University B, a private university with a campus population of approximately 12,000. Henry manages the campus-wide swipe-card system that controls access to most university buildings. Kevin and Tony manage physical access control for their respective departments using the swipe-card system, keys, and key-code pads. Both Kevin and Tony support departments of approximately 1,500 people. Fred manages a file system for the same department as Tony. Sue manages another building at University B that houses 170 people from several departments. Finally, Jerry is the lab manager for a lab located in Sue's space. Jerry's research group includes 70 people, but researchers from other groups occasionally need access to Jerry's lab.

Organization A is a large non-profit membership organization that has research departments. Organization A has approximately 1,200 full- and part-time employees and about 1000 volunteers. The organization is divided into five divisions, which are physically located at different sites around the city where it is based. Each division has its own departments, systems, policies and cultures, which are loosely linked by the main organization. Seth is the Security Director for all the divisions in Organization A and Ralf is the central network administrator.

Organization B is a non-profit organization that spans multiple states. The organization takes security very seriously; for them, a single breach could be detrimental to their business

model. David is their central administrator and he controls access using swipe cards and physical keys.

Organization C is a smaller non-profit of around 200 people that researches and evaluates data provided by other organizations. They also take security very seriously because their business model depends on other organizations trusting their security measures. Beth and Sara are the two administrators tasked with overseeing the physical-key and swipe-card systems for the organization.

The individuals we interviewed represent a broad range of policy professionals from a variety of different types of organizations with differing organizational structures and access-control needs. However, this study did not include interviews with policy professionals in for-profit companies or very large organizations. While we expect that most of our findings are likely applicable to for-profit companies and very large organizations, interviews with policy professionals in these organizations would likely reveal additional issues not discussed in this chapter.

3.1.3 Semi-structured interview

We used semi-structured interviews as our method of inquiry because they allowed us to focus on several primary questions but still have the flexibility to explore comments made by the interviewees. We designed our questions to focus on typical policy-management tasks but we also asked if the person had ever had to deal with specific incidents such as quickly revoking access rights from a terminated employee. The majority of people we interviewed performed policy-management tasks as only a portion of their job and in some cases fairly infrequently. The questions were designed to not only explore topics of interest, but to specifically bring up common incidents as a way of encouraging interviewees to remember specific events.

Our questions focused on several topics:

- Overview of interviewee's role in the organization.
- Technologies used by the organization and interviewee to control access within the organization.
- Policy changes caused by employee movement in the organization, including new employees, terminated employees, temporary employees, and employees who have moved internally in the organization.
- Written and unwritten procedures for making changes to the implemented access-control policy for a resource.
- Security incidents that have happened or could happen in the organization.
- Procedure for reviewing the implemented access-control policy for errors and checking the access logs for irregularities.

3.1.4 Data Analysis

The interviews were conducted in concurrence with the data analysis to better facilitate theory building. After conducting each semi-structured interview we used the audio recordings or detailed notes collected during the interview to analyze the interview content by building workflow, artifact, sequence and cultural models [20]. During these analyses we identified interesting topics which were recorded and added to the list of questions used in successive interviews. When the interviews were completed we used affinity diagrams to organize the topics we identified in our interviews. Topic groups were then used to construct theories. This approach is similar to other studies presented at CHI and SOUPS [22, 51, 113].

We constructed affinity diagrams [20] using comments, issues, breakdowns and successful solutions identified while constructing the work models. We wrote each piece of information on a sticky note and organized them into groups of similar topics. When reviewing topic clusters on the affinity diagram, we noticed that some topic groups described both problems and solutions but others only described solutions that indicated the presence of unmentioned problems. Using information from the work diagram and the affinity diagram, we identified the problem and solution (if any) that each topic represented.

Using the complete list of problems and solutions discussed in our interviews, we identified common problem themes. We grouped the problems based on similarity and causation in order to better understand the larger issues.

3.2 Roles of policy professionals

Our interviews revealed two types of policy and two roles for policy professionals. *Policy makers* formulate *intended* policy—policy that they believe should be enacted. Intended policy represents a single person's or a group's intentions; multiple intended policies that refer to the same resource could potentially be inconsistent with each other. For example, an employee at Company A may want to give access to her files to her friend at a competing company, but this is inconsistent with the general policy of Company A. *Policy implementers* translate the intended policy into the *implemented policy*—policy that is enforced by the access-control technology deployed by the organization. In doing this, they may need to adapt abstractly defined intended policy to fit the capabilities of the access-control mechanism and recognize or resolve inconsistencies or oversights in the intended policy. The distinction between policy makers and policy implementers is key to understanding how access-control policy is managed in an organization.

A person filling a policy-maker role is both empowered to make decisions concerning portions of the organization's policy and has (some of) the knowledge required to know what changes should be made. Policy decisions include assigning users to groups and giving individual users access to specific resources. Policy makers do not necessarily know how to change or view the implemented policy.

A person filling a policy-implementer role has the ability to make changes to and view the implemented policy. Unlike a policy maker, a policy implementer does not necessarily have insight into what changes need to be made or why, or what policy needs to be put

into place. Implementers depend on policy makers to decide what changes should be made and what the policy should look like.

It is possible for a single person to simultaneously fill the roles of policy maker and policy implementer. For example, an end user of a file system may both know what the policy for his files should be and have the ability to change the access-control permissions for those files. In a centralized system, an end user may be forced to request certain changes to the file permissions, which the central system administrator then implements. Our interviews were focused on policy professionals in centralized environments and all our interviewees were either policy makers or implementers.

We found that the largest issue faced by implementers is knowing what changes need to be made to the policy and when to make them. Conversely, policy makers know what the policy should look like but have limited to no ability to view or manipulate it. This issue arises because makers and implementers are typically different individuals, and because coordination can be difficult.

3.3 Policies are managed by multiple people

Many policy professionals expressed concerns about managing a policy where multiple people are capable of changing the policy with little or no notification. Issues mentioned by policy professionals ranged from concerns about synchronizing policy edits across multiple professionals to the difficulty of managing exceptions to the policy. A common theme was a need to have a way to know at all times what the policy says and whether it is still accurate.

3.3.1 Maintaining an understanding of the implemented policy

For many policy professionals the biggest challenge is maintaining a good understanding of the current implemented policy. Policy implementers need a working understanding of the policies they maintain because they are asked to make decisions based on the policy and it is not always convenient to access the policy itself to answer the questions. While we were interviewing him, Ralf received a phone call on his mobile phone concerning an employee in department A who was filling in for an employee in department B. The employee couldn't log into the computer of the employee she was replacing. Because he has an excellent working knowledge of his organization's network access-control policy, Ralf was able to identify the problem, determine whether a temporary exception was needed and instruct his assistant to fix the problem, all without accessing his computer. Ralf explained that being able to solve problems over the phone was very valuable because he was rarely at his desk and didn't always have access to a computer where he could look things up.

When only one person manages the policy it is easy to maintain a working understanding of the policy. However, 11 of our 13 interviewees worked with at least one *policy implementation peer*, a person with similar responsibilities and abilities as themselves. With multiple policy implementation peers making changes to the implemented policy, it

is difficult for any one policy professional to maintain an understanding of the state of the current policy.

The policy implementers we interviewed solved the issue of not knowing what other policy professionals were doing by using a standard set of heuristics for dealing with policy maker requests and by notifying others about changes. David is the primary policy implementer for a physical-access-control system. Whenever an incident occurs in any of the buildings he manages, he is the one who gets called and asked to explain why the incident happened. David likes to know what changes are being made to his system because he may be asked about the implemented policy at any time. When discussing how he coordinated policy changes with his team, he told us that he trusts his fellow policy implementers to know what a normal request looks like and to address the request appropriately. However, he still wants to be notified after any such change.

Ralf has a more complex coordination problem amongst his policy implementation peers. Each of Ralf's coworkers is responsible for a different part of the organization's policy. For example, one of Ralf's coworkers manages the firewall policy. Another coworker manages the file-system policy for one of the departments. Only Ralf has an understanding of how all the different systems and policies interact. When any of his coworkers has a question about another part of the system, the coworker goes to him. Ralf told us that he makes sure that he is aware of all changes occurring on his system. He instructs all his fellow coworkers to report any changes they make to him so that he always knows the state of the system. Having a holistic knowledge of the system lets him make decisions without having to consult anyone else or dig through system files. Ralf commented how his memory was the most complete set of documentation at the organization. His manager wanted him to start documenting the information he was collecting because it isn't written down anywhere, and if Ralf ever had an accident then no one would know what was going on in the system.

Only one policy implementer, Fred, wasn't concerned about maintaining a working knowledge of the implemented policy he worked with. Fred's department has about ten policy-implementation peers and the department is known for having a lot of employee turnover. Each policy maker who uses the file system has the ability to make changes to the implemented policy for their files. With so many individuals making changes, maintaining a working knowledge of the policy is infeasible. Fred doesn't bother trying to understand the current state of the system before making changes and instead simply verifies that the requested changes don't conflict with the high-level intended policy of the university, which is fairly loose.

Implementers also discussed giving two independent groups of policy professionals responsibility for making alterations to a resource's implemented policy. All four implementers who mentioned it felt that it was a bad idea. Kevin and Tony both felt that either they should manage the implemented policy for a room themselves or the policy maker should manage it directly, but they didn't want to be placed in a situation where they might be blamed for a change they did not make. Henry, who manages a physical-access-control swipe-card system for all of University B, has a similar opinion. He makes sure that every door in his system has exactly one group that can change its implemented policy. Tony talked about how he had once requested that Henry let him share manage-

ment responsibilities for a door with another department. Henry had refused the request and told Tony that either Tony's department could manage the door's policy or the other department could, but not both.

3.3.2 Exceptions are hard to manage

Exceptions to normal policy were a problem even for groups who had established an effective method for communicating policy changes. An exception is any change to the implemented policy that violates the "normal" intended policy of the organization. For example, giving an office key to someone who doesn't work in that office when the organization has a "one office, one key" policy is an exception. Normal policy changes such as adding a new user have a well-defined set of tasks associated with them. Adding an exception to the policy means the implementer must manage the exception separately. Implementers who tried to maintain knowledge about the current policy state found exceptions especially irksome.

None of the implementers like exceptions and four of them attempt to ban exceptions from their systems. Ralf dislikes allowing exceptions because they are hard to manage, and worse, it is hard to remember that the exception exists. On his file server, Ralf has a policy that each user gets her own directory that only she can access and each project group gets a common directory that can be accessed only by members of that group. Ralph explains:

> They have that common [disk] drive, and occasionally they get into this situation where they're like, "I don't want anyone else to see that," you know, because anyone in their department can see that.... And you're like, "OK, so, like now I have to make another folder just for you two?" It actually starts to become an administrative nightmare.... I try not to make too many changes and I try and explain that to them upfront and say, "Look if you want I'll do this once but I don't want to be doing this five times."

David, Beth and Sara were concerned about the possible negative effects of allowing exceptions to their policies. Because their organizations take security very seriously, it is important that employees such as security guards be able to spot abnormal access behavior. One way this is done is by using chemically-treated temporary badges that change color over time, allowing anyone to identify temporary visitors who have stayed too long. Similarly, they want employees to be able to identify odd access behavior. Allowing exceptions makes the policy non-standard and makes it harder for employees to determine whether someone has legitimate rights to a space or not. In general, the policy professionals from both organizations attempt to limit or prevent exceptions. In the rare case where an exception is necessary, Beth and Sara grant the exception, but their resistance to exceptions and the small size of their department means that there are only ever a few exceptions in place at any given time. David manages several departments so he completely refuses to add exceptions to the system's implemented policy. Instead, the security guards maintain a list of people whom they can let into certain rooms. A new person, room pair can be added to the list by filling out a form at the guard desk. This workaround allows people to be given access without adding exceptions to the implemented policy. The solution also allows the security guards to identify odd behaviour.

3.3.3 Getting policy-change notifications

Many policies depend on information from multiple sources. A common type of policy, for example, gives all employees access to a resource. The policy maker who formed this policy does not, however, know who all the employees are; this information is managed by the human resources department, which in this way also plays the role of a policy maker (e.g., granting access to newly hired employees and revoking access to departed employees). The implemented policy is a result of appropriately combining input from the two policy makers. Accounting departments, which typically allocate internal charges for network access and other services based on each employee's home department, are also potential sources of information about employee internal movement and employee termination.

Four of the interviewed implementers mentioned the benefits of setting up their access-control system to use records maintained by another department. Three other implementers mentioned how they were currently trying to establish better relations with the human resources or accounting departments in an effort to more quickly get information about changes in employee status.

Henry manages a swipe-card system that controls access to physical and virtual resources at University B. The turnover of people involved with the university is so high that he doesn't want to individually add and remove each person from the system. Instead, his department works closely with the Registrar, which monitors the status of all faculty, students and staff at the university. The swipe-card system is linked in with the Registrar's system so that when new people join the university they are automatically given access to communal university resources. When people leave the university their access rights are automatically removed.

Henry's arrangement with the Registrar also helps Tony who manages access control for one of the departments at Henry's university. Since Henry's department automatically adds and removes swipe-card accounts, Tony doesn't need to worry about routine university events and can focus on department-specific access-control concerns.

3.3.4 Documentation is old or wrong

In several cases policy implementers discussed making decisions based on information stored inside the system that was out of date or wrong. In these cases, policy implementers had to recognize that the documentation was not valid and find alternative ways to get the data they needed.

One such example came from Fred, who receives requests from people who want access to files and folders on a file server. Fred uses the access-control list in the file system to determine who owns the folder or file and treats that person as the only person who is allowed to make decisions about it. Occasionally, he will tell a requester that they have to get the folder's owner to send him the change request only to be informed that the owner is no longer at the university. This is problematic for Fred since he must then locate the new person in charge of the folder's policy and update file ownership in the system.

Another example of information being entered into the system and becoming stale comes from Kevin, who manages physical access for his department at University B. Every

time a person is given a key to a space, this fact is noted on an index card titled with the person's name. If the lock is re-keyed, however, this information is not added to the card. In order to determine if someone has access to a specific room, her card must be pulled up and the number of the key she was given must be compared with the current key number for the door, which must also be looked up in a separate record. Information stored on the card about which door the key opens cannot be trusted since it may be old.

Documentation can also be completely missing. Kevin told us that his department also uses a swipe-card system to control access to some resources. However, since not all students, staff or faculty need access to these resources, swipe cards are issued only on an as-needed basis. An administrative assistant gives out the card and notes this fact on a piece of paper, so that later an implementer can activate the card and add it to the system when he has time. Consequently, the database of swipe cards is often incomplete since the implementer doesn't always have all the information available when he enters the card into the database. Without complete information, knowing who has what card is difficult.

3.3.5 Discussion

Working with multiple policy professionals can cause problems with keeping relevant people apprised of the current policy state and keeping the policy synchronized. Policy implementers feel they need to be notified about changes in the policies they manage. This suggests a need for technologies that provide notifications when policies change and provide methods of documenting why a change was made. They also need a way to incorporate parts of the implemented policy that are maintained by other departments.

Make documenting implemented policy changes part of the natural workflow. The majority of the problems described by policy professionals trying to coordinate edits to the implemented policy centered on their need to know what the current implemented policy looks like. One solution to this problem is to encourage policy implementers to document their changes to the policy. Good documentation would allow other policy professionals to learn about the policy without having to memorize it.

It would be better, however, if documenting the reasons for a change in the implemented policy was an integral part of making a change to the policy. This could be done in two ways. First, policy management systems could require users to document a policy change (and aid them in doing so) before the change was accepted by the system. Second, the implemented policies could be specified in a self-documenting policy-specification language [4, 7, 64]; i.e., the implementation of the policy could preserve many of the abstractions of the intended policy. For example, the implemented policy could explicitly encode the sub-policies, "John is a student," and, "students can access the lab," instead of encoding just, "John can access the lab," as is more common. Some policy-management systems provide such functionality and could be further improved to support implemented policies that are even closer to intended policies.

Provide a way to keep policy implementers apprised of changes to implemented policy. For many of our policy implementers, having good documentation that could be consulted in case of need wasn't enough. They needed to have an excellent understanding (without referring to documentation) of the implemented policy at all times to properly do their job.

For these policy implementers, we recommend using a publish-subscribe technology where the system automatically sends out updates when the policy changes and implementers can indicate that they want to receive updates about certain parts of the policy.

Automatically update compound policies. Implemented policies may depend on information that is maintained on separate information systems, e.g., databases spread among several departments may house different pieces of information relevant to the policy. Integrating these different systems so that the implemented access-control policy is automatically kept up to date has many potential benefits. Several groups we interviewed used systems that had this functionality.

3.4 Policy makers are distinct from policy implementers

Another major issue expressed by both implementers and policy makers is the challenge of knowing when a change needs to be made and determining what that change should be. Policy makers expressed concerns that the implemented policy does not match their intended policy and it is difficult to review changes made to the implemented policy. Implementers discussed problems with knowing when a change needs to be made, verifying that the person requesting a change has the appropriate authority, and maintaining records demonstrating the request.

3.4.1 Viewing implemented policy

Unlike policy implementers, policy makers typically do not have the ability to view or manipulate implemented policy directly. Instead, they have to find and query an implementer to get an understanding of what the implemented policy looks like. Everyone we interviewed mentioned at least one incident where they had to ask for or were asked for a report about the implemented policy.

Since policy makers do not know what the implemented policy looks like, they have no way of knowing if it is correct or not. Those policy makers who are concerned about the wrong people accessing resources for which they are responsible request portions of the implemented policy from an implementer and review it for errors. According to Sara, both she and Beth review the access-control lists (ACLs) for each door in their department once or twice a year—however, they would like to do so more often. Since they don't have direct access to the ACLs they send a request to the implementer who sends them the ACL for each door. They then go through these lists looking for anyone with inappropriate access. Sara tells us that occasionally they do find people who shouldn't have access. She attributes this to a "slip of the finger" on the implementer's part. After they review all the ACLs they send a list of corrections to the implementer who makes the requested changes.

Kevin and Henry both mentioned that they occasionally get requests from policy makers wanting to know the implemented policy for their resources. Henry says that he gets general requests asking for a list of everyone who can access a specific area. Kevin's system doesn't support the ability to create a list of everyone who can access a space, and so he doesn't get many of those requests. Instead, he gets asked about specific people. Kevin says that

a few times a year he will get an email from a policy maker asking if a specific person has access to a specific room because someone has just spotted the person there and is not pleased about it.

3.4.2 Getting notifications about policy changes

A major problem faced by implementers is knowing when the policy needs to be changed. Since an implementer doesn't always know the intended policy, it is difficult for them to detect inconsistencies without the help of a policy maker. Implementers either ignore these problems, trusting that a policy maker will notify them when a change needs to be made, or they proactively attempt to get change information from policy makers.

Ralf, Seth, David, Kevin, Sue and Fred all discussed instances where they were not notified about pertinent personnel changes which should have resulted in changes to implemented policy. Ralf, in particular, was annoyed about not being told when employees leave the organization:

> I try to disable an account as soon as I know that account [holder] is gone.... As soon as you do it then all of a sudden they are complaining because they will try and bring somebody else in and say well they were using that account and I'm like, "No, that doesn't work, they need a new account...." I just don't like them using [an account] under somebody else's name.... Who knows if someone else knew what their password was or that person got back into their account again and is using it along with this person.

David found a more proactive solution. Instead of waiting for a policy maker to complain or request an access-control list for review, he proactively sends out lists to all policy makers on a monthly basis. This method allows David to find potential errors in the access-control policy for his organization before they become issues. It is unclear how effective this method actually is since Beth and Sara, whom we also interviewed, regularly receive these periodic lists but only review them once or twice a year.

David explained how he also tries to encourage policy makers to send him information about policy changes—such as employee termination and internal movement—in advance so that he can schedule the changes and ensure that the employee loses and gains access at the appropriate times. Temporary employees such as students are entered into the system with a start and end date so they are automatically removed once they leave. David says the system works well but every so often a policy maker will forget to tell him about a change in plans and someone will be denied access.

3.4.3 Verifying requests and keeping records

Since an implementer does not have perfect knowledge of the intended policy, she has to trust policy makers to make the correct decisions about what policy should be applied to the resource. However, when a problem occurs, implementers are concerned they will be blamed since the state of the implemented policy is their responsibility. To address this issue, implementers perform sanity checks on requests, verifying that the request matches

the organization's policy and that the person making the request is authorized. Implementers also keep records of the change requests they receive so they can reference the records if there is ever a problem.

One of the first issues an implementer encounters when presented with a policy change request is validating that the requester has the authorization to request the change. Of our interviewees, six know who owns each of the resources they support. The rest of the implementers either consult documentation to determine ownership or find another trusted person to ask. For example, when Fred gets a request to give someone access, he consults the file or folder in question and determines if it is owned by the requester. If the folder's system-indicated owner is no longer at the organization (a reasonably frequent occurrence), Fred sends an email to a trusted administrative assistant and asks who has taken the previous owner's place in the project group.

Implementers are also concerned about accountability. Most implementers we talked to keep records of who requested each change along with some sort of proof. Typically, these records are the emails requesting the change. Our interviewees expressly pointed out to us that they keep these emails specifically for accountability. Other types of records are also kept by implementers. Kevin's department requires that the requester sign a form before new access is given to someone. Fred's department also uses a form that must be filled out for a new user and includes who the requester is. For other types of requests, Fred uses a help request tracking system that allows him to tag requests involving policy changes.

3.4.4 Discussion

There appears to be a natural divide between policy makers and policy implementers. Both policy makers and implementers perform their own specific sets of tasks, but they need to communicate with each other to accomplish their tasks. Access-control systems should seek to reduce this divide or better facilitate communication across it.

Allow policy makers to directly edit the implemented policy. The principle of least privilege—that a person should be given the minimal rights needed to do her job—is a well established axiom in security [89]. We observed that, in practice, policy implementers often do not have sufficient understanding of the intended policy to accurately enforce the principle of least privilege. Providing policy makers with the ability to make changes to the implemented policy themselves would let them leverage their greater knowledge of the intended policy to create a more accurate implemented policy. To enable this sort of policy creation, access-control systems would have to support interfaces tailored to policy makers, exposing and allowing the policy makers to control only the portion of the policy for which they are responsible and only in ways that match their authority. There has been some success in designing experimental systems with such features: Bauer et al. found that when they gave policy makers a more flexible access-control system the participants created less permissive policies that better fit their needs [13].

Provide feedback to policy implementers. Giving policy makers direct access to their portion of the implemented policy makes some implementers uncomfortable, as they worry that policy makers will introduce errors or leave the policy in an inconsistent state. Beyond building in safeguards that ensure that policy makers cannot implement policies that

they are not authorized to make, as described above, systems could improve the feedback implementers receive as a result of changes being made to the implemented policy. In the limit, systems could allow implementers to preview and approve the policy changes introduced by potentially technically unskilled policy makers, thus providing the flexibility for policy makers to implement policy and still allowing other policy implementers to ensure the changes are reasonable.

3.5 System can't enforce desired policy

Policy professionals also have to consider how policies will work in combination with the technology that enforces them and what will happen when people do not follow secure practices. The topic of policy enforcement is very broad and mostly outside the scope of this study. However, we touch here on enforcement issues that arise as a result of the decisions policy professionals make about managing their resources.

3.5.1 Choosing an access-control technology

When implementers discussed the access-control technologies they used, they almost always began by discussing the enforcement abilities of the system. Implementers were very interested in features such as reliability, the ability to fail gracefully, and simplicity. They had less, if anything, to say about the types of policies the system supported or the management interface.

Nearly all the implementers who managed physical access-control had participated in the selection of the system they worked with. Kevin explained to us how he used different combinations of technologies on every door to get the perfect mix of reliability, security and usability for each lab. Technologies such as keys and key pads were used by implementers because of their reliability, stand-alone qualities and the ease with which access could be shared.

Implementers were very interested in the capabilities of the systems but David was the only implementer who seemed interested in the management software associated with the system. He purposely selected his system because it had a management interface that was easy to use and integrated information from his other security technologies. The other implementers only minorly considered the management software in their selection process. University C's requirement document for their new building primarily specified physical requirements of the system and only occasionally mentioned a requirement for the software (such as the ability to schedule exceptions at least a year in advance).

Implementers didn't mention the usability of the management system when selecting technologies, but they were annoyed by poor management systems. David talked about the old system his organization used which could require that a user be added or removed from as many as ten databases when making a change. Tony tried to show us the management system for the swipe-card system he works with and quickly gave up. He told us that his co-worker had received the training and performed all interactions with the system.

3.5.2 Knowing who has an access token

In the previous sections we have assumed that the person exercising an access right is the person who is intended to be doing so. Unfortunately, this is not always the case, as several commonly used access-control technologies do not link the access request to a person (e.g., physical keys).

Using key codes may be more convenient for policy makers, since they can give access without interacting with management technology, but it means that no one knows exactly who can access a resource. Ann and Kristen talked about the problems with using key codes for lab doors. Each lab door has a single key code which is given out to all the occupants. Theoretically, the key code is only known to the room occupants, but nothing stops the occupants from sharing the code with others. As a way of ensuring that only the room occupants have access, the key codes are changed once a term and the new code is emailed out to all the room occupants.

Kevin's department also uses key codes for some doors. He has similar problems as Ann and Kristen, except that his department is much larger and he is not always certain who should be given the new key codes. To solve the problem, he instead emails out the new code to the administrative assistants who work with people in the space, and asks them to distribute the new code. Ann talked about a specific incident where a door code had to be changed:

> We just had a professor the other day who sent an email saying, "Some people I don't want in the lab sought access." So [we changed the code] ... then she just gave it out to two people that she said was ok.

Physical keys can also cause problems. Users frequently give keys to others to facilitate achieving a goal, even if this is not consistent with an organization's intended access-control policy [13]. Keys, also, can be easily copied even if stamped "Do not copy." As a result, it is hard to know who has a key to a room even if accurate records are kept. Tony, Kevin, Ann and Kristen all talked about the need to periodically re-key doors just to be certain that only the correct people had keys.

3.5.3 Unexpected events

Dealing with unexpected events is another important part of enforcement that needs to be considered in the implemented policy. Owners do not think of all possible events *a priori* and unexpected events do occur.

Jerry spent the most time talking about unexpected events since he is the policy maker for a lab filled with expensive equipment. His intended policy for whom he will allow to have access to the lab is fairly restrictive, with only a few staff members given access. However, if an incident occurs (e.g., a fire) he trusts a much larger number of people to enter the lab (e.g., in order to shut down expensive equipment). He would like a system that allows him to give out a type of access that could only be used in emergencies and would immediately warn him when it was used.

Kevin has also had issues with unexpected events. This happens often enough that Kevin started adding key-pad locks onto all lab doors. He puts an administrator code onto

each door so that if he gets a call and can't come personally, then he can simply relay the code to the caller and change it the next day.

> What was happening is we were having different people coming in at night, emergencies and what not, not everyone has a key, not everyone has a card. So if I get a call at home. A guy calls and says "Hey I'm down here, I can't get into [a lab]," so I have a code in there that I can give them. I've given it to a lot of maintenance people and security people which gets them in there.... doot doot doot you're in.

3.5.4 Discussion

Managing the actual implemented access-control policy in the wild is a challenging task. Policy implementers are limited by the types of technology available to them. Even when they can choose the technology that best suits their needs, they still have trouble configuring it to their specific situation. Current technologies don't necessarily support policy implementers' need to change intended policy into implemented policy.

Prioritize management interfaces and ability to implement desired policies when choosing systems. The policy implementers who worked with physical access-control systems viewed reliability as a major requirement. Physical keys, key pads and some of the swipe-card systems were selected because they could function autonomously if necessary and they had a low failure rate. However, an insufficiently flexible management interface or the inability to enforce desired policies can be as great a detriment to security and convenience as unexpected failure of the system, e.g., a power outage. Hence, we suggest that these features be given greater consideration when systems are being chosen.

Take advantage of new technologies. New access-control technologies make it possible for access-control systems to achieve a previously unprecedented degree of flexibility and security. Smart-cards, RFID badges, or even software on commercially available mobile phones can all be used to enable access-control systems that make it cheap and convenient to extend access to new users, delegate access on demand and in an ad hoc manner, and yet provide a high degree of auditability and assurance that unauthorized access will not be allowed. These technologies can make it unnecessary, for example, to share keys or key codes, and we believe that adopting them would benefit many organizations.

3.6 Related Work

To the authors' knowledge there has been no other study of physical or file access-control policy professionals. Other researchers have studied computer system security professionals in general but have not focused on the specific role of access-control policy management. Barrette et al. studied security professionals who worked in a system administrator role. They found that administrators are very collaborative and work together combining their specialized knowledge to solve problems [10]. Much of the information administrators use is both specific to their organization and exists in many places, requiring administrators to combine the information using custom tools [10, 22].

Few studies examine the challenges of managing a physical access-control system. Bauer et al. interviewed members of a university department prior to the creation of new a physical access-control system. They determined that authority to grant access to resources was passed down the departmental structure with the department head delegating to the building administrator, who delegated to various staff members [11].

Designers have considered the problem of creating tools to assist policy professionals. One example is the SPARCLE Policy Workbench, which allows policy professionals to write privacy policy in natural language and parses the policy and converts it into an implemented policy [23]. The Expandable Grid is another example of a tool which allows users to manipulate the implemented policy for File System rules [83]. These tools are promising but neither are based on actual experience of policy professional issues and tasks.

Other studies have focused on users and how they use security enhancing technologies. Gaw et al. studied the use of PGP for securing email communication. They looked at an environment where security was very important and the employees were motivated to secure communication. They found that employees still did not regularly encrypt their email for various social and convenience issues [42].

Dourish et. al. explored end user's use of security technology. They found that end users tend to delegate security concerns to trusted individuals or groups and trust that their resources are secure. In organizations this sometimes caused a mismatch between the security settings and the current needs of a group [34]. Singh, Cabraal and Hermansson conducted interviews of banking customers in Australia. They found that banking customers often used insecure methods to manage their finances because the secure methods failed to match the social or cultural situation or were very inconvenient to follow [93]. Similar to Gaw's work, this shows that users do not use secure but inconvenient technologies unless they must.

3.7 Results in terms of other studies

In this study we found that existing access control technologies do not always completely support security administrators in their task of creating and maintaining access-control policies. Administrators make use of a varienty of technologies to create the security effect they want. In other work we further explored this behavior in administrators, office workers, and end users [12, 13, 71].

Access-control technology designs frequently assume that users want to divide the world into two groups of people, those that should be able to perform a specific action on an object, and those that should not be able to perform the action on the object. However, most users manage their personal access-control using more fine grain distinctions.

We found that when managing access-control people use a wide range of tactics and social pressure to enact security policies that would be technically infeasible using only system settings alone. The tactics used fall into five main categories: planning for the unexpected, in-the-moment, witnesses, obfuscation, and audit.

Planning for the unexpected – People would give physical keys to another person

with the explicit instruction that they key was not to be used except in an emergency. While the other person was trusted, the goal of the permission granting was not to give them daily access. A combination of trust and social pressure was used to make sure the access was not abused. An "emergency" was defined as any unexpected event where the access granter either was unavailable or had remotely authorized the access.

In-the-moment – Privacy and security are often highly contextual, giving access is not only about *what* and *who* but also *why*. People who normally did not want anyone entering their offices would mention several highly context dependent specific situations where they wanted to allow someone into the office just once for a specific purpose. Those who needed to give in-the-moment access would typically call someone who had an emergency key to open the door or verbally state a key code.

Witnesses – Offices, homes and even folders are spaces that can contain many types of content. Giving a marginally trusted person access to one, even for a specific purpose, was perceived as risky. When giving in-the-moment type access to an untrusted person the permission granter would require that a trusted person be present. This trusted witness would provide access credentials on behalf of the untrusted person and be physically present to witness all actions that were taken in the space.

Obfuscation – Physically or digitally hiding an object that needed to be protected was a simple low-tech tactic. Hiding required limited understanding of how the security system worked and participants had confidence that no one would target them sufficiently to find the hidden item. Hiding allowed someone to give access in-the-moment without using a third party by verbally telling the accessor where the object or credential was hidden.

Audit – Another tactic was to place trust in a group of people, give them access, trust them to behave correctly, but have a way to audit their actions later. This was enforced either with logs or by placing the object being accessed in an open space visible by many people. This tactic uses minimal technological mechanisms to force correct behavior and instead uses social pressure and the threat of punishment to encourage correct behavior. It also allowed the person whose item it was to make judgments with an understanding of the actual consequences of the actions.

The act of controlling access is not just about allowing someone into the office or not. Issues such as context, purpose, levels of trust, and the ability to reserve judgment until effects were known, were all major factors. An effective system should support users in these behaviors.

3.8 Conclusion

We interviewed thirteen policy professionals from five organizations in an effort to understand the challenges involved in policy management. We found that policy management had three sets of real-world requirements that were either ignored or not adequately addressed by technology: 1) policies are made/implemented by multiple people, 2) policy makers are distinct from policy implementers, 3) current access-control systems can't always implement the desired policy. Based on our observations, we suggest a number of improvements that could be made to access-control system.

Access-control systems should support easy communication between policy professionals. By encouraging policy implementers to document the policy changes they make, it may be possible to provide vital information for those who will manage the implemented policy in the future.

System designers also need to be aware of the existence of two policy professional roles: policy implementer and policy maker. Each of these roles is associated with a different set of skills, abilities, and tasks. Policy implementers have the ability to make direct changes to the implemented policy. Policy makers have the ability and knowledge to know what changes should be made. Designers of policy-management systems should understand the tasks and limitations of both roles and design to support the differences.

The capabilities of the enforcement technology itself are important. New technologies make it possible to enforce security policies that older technologies, like keys and key-code locks, cannot. Access-control systems also differ in their policy-management interfaces, some of which are far more flexible and expressive than others. In addition to more typical concerns like the ability of a system to withstand a power outage, these capabilities need to be given careful consideration when selecting an access-control system.

Finally, we discuss the results of this study in relation to other studies we have conducted. We find that users in general use a variety of tactics when managing their own security. Some of these tactics are based on technology but many are a combination of technology and human trust relationships.

Chapter 4

Focus Group: User reactions to proximity security information

In our studies of how administrators and end users manage access control, two major issues stand out. First, the access-control policies currently enforced by the access-control systems frequently do not match the policies people want to have enforced. Second, people don't actually know what access-control policy is currently enacted by the system. Effectively this means that what people want to have happen does not match what is actually happening, and people are not aware of this fact. Awareness is one of the corner stones of security management, if you don't know something is wrong you may make inaccurate and potentially dangerous decisions based on a false sense of security.

We theorized that the best way to address this issue was to improve end users' awareness of their current access-control policies. When we talked to administrators in the discussed in Chapter 3, one of our observations that there is a the separation between policy implementers and policy makers. While these groups were distinct in our organizational research, we believe that digital systems are moving towards a model where end users are expected to embody both roles. Therefore we believe that the best way to support digital access-control in the future is to make end users aware of what their current security settings actually are.

4.1 Theoretical approach

When we initially conceptualized proximity displays for photo sharing systems we had the following goals in mind:

Proximity Place security information in spatial proximity to photos instead of on a secondary page where it cannot be seen. This was important because people don't tend to think something is wrong until they observe a problem. If the data necessary to observe the problem is hidden people may never identify the issue.

Audit Display information about who has seen the photos. In Section 3.7 we observe several tactics people use to manage security, several of which involve giving other people access and expecting them not to use it. The access recipients are generally

trusted, but the person who owns the resource has no way to identify if the access is used when it should not be. Presenting information about who has accessed their resource would give people a chance to reassess their policy decisions.

Policy settings Display information about the current permission settings. Studies, in addition to our own, have indicated that people frequently have mismatches between the settings they say they want and the actual settings enforced by the system [13, 68]. Additionally, some of the tactics observed in Section 3.7 rely on the user accurately knowing what their current policy is. Displaying the current policy to end users would increase awareness and help them to identify errors.

Layered Enable layered data exploration by moving some details to a secondary interface. Proximity displays should help people realize that an access-control problem exists, but they may not be the best mechanism for supporting the user in identifying the scope of the problem or correcting the issue. We expect that some information needs to be visually apparent to the user, while other information should be *layered* – shown to the user only when they interact with the display. Layered data can be anything from tool tips that only appear when the user's mouse rests on a component to a secondary page where permissions can be explored in depth.

Specific names of people/photos Display who, not just what group, saw or could see information. When talking to users and security experts we noticed people bringing up examples where a problem had occured because the membership of a group or folder was not quite what was expected. This was especially true with user groups where multiple people could alter the membership of the user group. We wanted to support users' ability to identify individuals who shouldn't have access, not just groups.

Comprehensive Show detailed data, not just an overview icon. Icons and other small passive indicators take up a small amount of space, and if you know what they mean, can present detailed information. Unfortunately end users do not always know what the icons mean and may never bother to find out. We wanted our displays to have enough information visible that a user could, with high accuracy, determine if there was a problem with the policy or not.

We decided to use focus groups as a way to rapidly test both the reasonableness of these goals and how people will react to permission displays. Using focus groups allowed us to get information on what aspects of the display designs people felt would be most useful to them. Though the results of focus groups provide only a general sense of how people will actually react to a finalized product, they are an excellent way to present and iterate through many ideas quickly.

4.2 Methodology

The methodology used was fairly typical of focus groups. Participants were asked several questions to get them thinking about privacy and security, then they were asked to comment on each of several website designs. Based on feedback we altered the designs between

focus groups so each group saw a slightly different set of interfaces. We transcribed the notes and the audio and grouped the comments by concept.

4.2.1 Participants

Participants were recruited from an existing pool of people in the Pittsburgh area who had previously indicated interest in behavioral research studies. We recruited a total of 28 for five focus groups. Each group had between four and six participants. The majority of participants were students, and all had previously shared photos online.

4.2.2 Protocol

The focus group was conducted in a conference room at Carnegie Mellon. The sessions were audio recorded and lasted for an hour. To get participants thinking about security and privacy we started by asking participants:

1. About the websites they used to share photos.
2. The last thing they shared online.
3. A time they discovered that someone they didn't want to share with could see their content. (Not all participants were expected to answer.)
4. A time they tried to share content and had not been able to due to technical issues. (Not all participants were expected to answer.)

Participants were then handed two pages with cartoons on them (Figure 4.1) which illustrated use cases for providing a user with privacy and security information in an easy-to-notice way. Participants were then asked: "Can you imagine an instance where you or a friend might experience a situation like those Alice and Joe encountered?" Participants were encouraged to briefly discuss the answer to this question amongst themselves.

Participants were then given a packet of website designs. Participants were told that they would be going through the packet as a group and were asked to not look ahead. For each webpage in the packet the researcher gave the participants a brief presentation of the site, its features, and any interactive components. If participants had any questions about how the site worked they were allowed to ask them at this point. The researcher then asked the participants to fill out the questions included in the packet in silence. When everybody was finished the researcher started the conversation by asking each participant what the best and worst parts of the website were for them and encouraged them to discuss the answers. Participants were informed that they could write in the packet at any point, so if they were unable to voice an opinion they were welcome to write it.

At the end of the study session the researcher asked each participant what their favorite and least favorite website design was and why.

4.2.3 Interface designs

We tested a variety of user interface display designs during the focus group and we iteratively modified them between focus groups to get a better understanding of what was

Alice is looking through her online photographs for ones she can use for her screensaver.

Alice is delighted to discover that her good friend Sue found time to look at the pictures of their trip to Chicago together.

Alice hasn't had a chance to talk to Sue much since they got back so she decides this is the perfect time to reconnect. Alice writes Sue an email saying how much fun she had and she can't wait to see Sue again at Christmas.

(a) Usage scenario where Alice notices that a friend has viewed one of her photos which results in a positive experience.

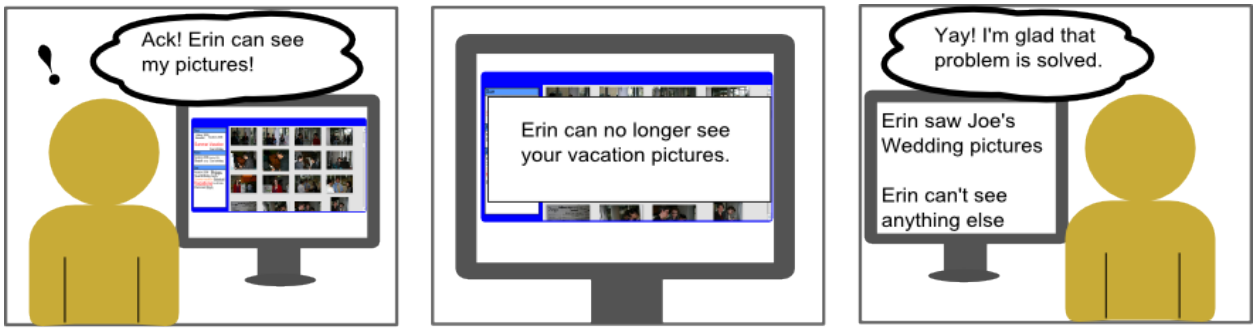

Joe is posting pictures from his vacation on his favorite photo sharing website. He is looking through the pictures to make sure they are all ok when he notices that his ex-girlfriend, Erin, can still see his vacation pictures.

Joe is very upset and quickly removes Erin from the list of people who can see his vacation pictures.

Joe is very concerned that Erin may have seen some pictures he doesn't want her to see. So he looks through the remainder of his photo album to make sure she can't see anything else and to check that she hasn't seen anything she shouldn't.

(b) Usage scenario where Joe notices that an ex-girlfriend can see his vacation photographs. Joe views this negatively which causes him to update his access-control policy.

Figure 4.1: Usage scenarios illustrating how a end user might use proximity displays both to cause a positive social experience, and to notice an issue.

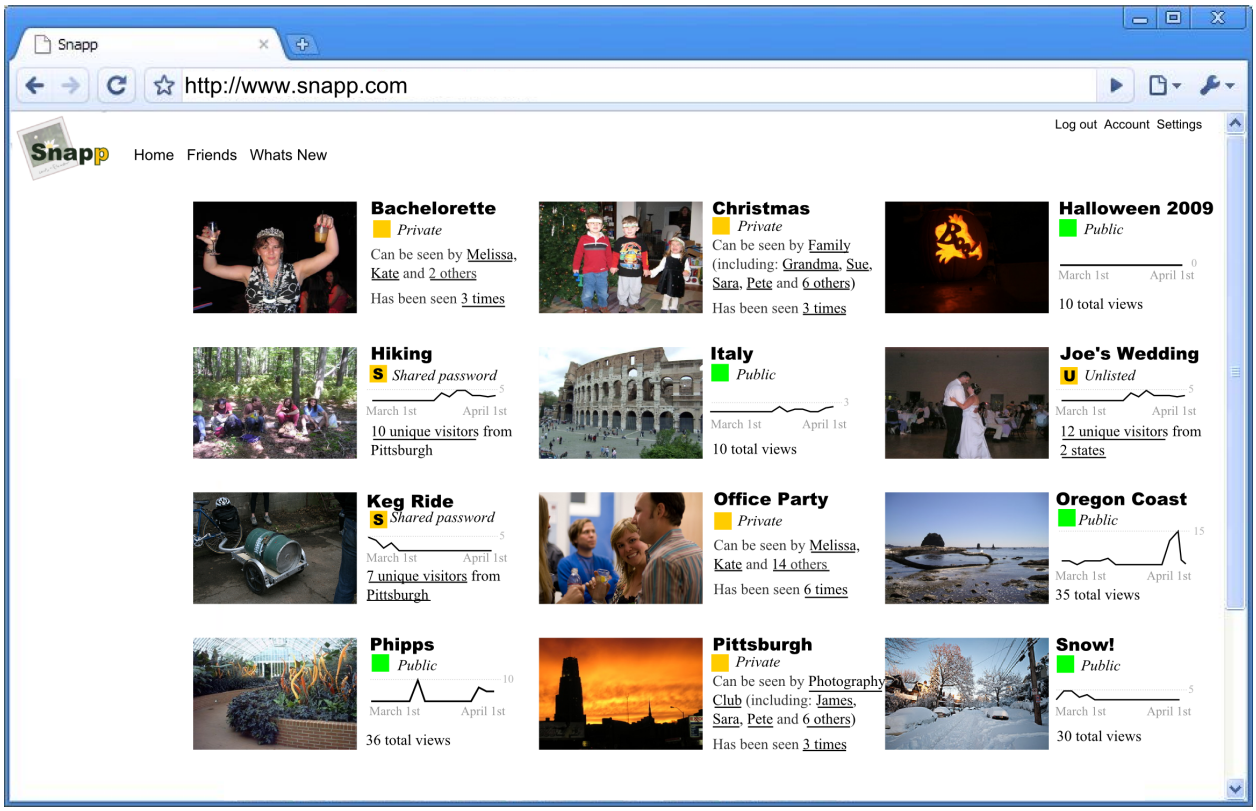

Figure 4.2: Example interface typical of the ones shown to focus group participants.

working and why.

The interfaces we decided to use were designed to loosely resemble real interfaces but left out many of the details an actual interface might have, such as the tools for photo editing or the Facebook and Twitter share buttons. The permission information shown was also mildly exaggerated, taking more screen real estate than would be normally feasible and showing more details than should be needed. We wanted participants to focus on the permission information we were showing them and have minimal distractions from other user interface components. We also wanted them to comment on multiple aspects of an interface with the end goal of decreasing the information shown to a more manageable set based on users' reactions and perceived usefulness. Figure 4.2 shows an example of the type of interface we showed to participants.

The interfaces chosen were intended to represent the following features to a greater or lesser extent:

Spatial proximity We wanted information to be spatially co-located with the photographs, but how close did it need to be? We tested a combination of placing information next to the album thumbnails and placing information on the sidebar.

Policy information We wanted to test different methods of communicating the current access-control settings and the impact of varying levels of detail concerning those settings. Some interfaces were more detailed, specifically naming people who could or could not access the albums. Other interfaces were very general, listing only the number of people with access and omitting individual and group names. We also tested interfaces that showed a high-level icon indicating public/private followed by a list of groups, which is less specific than names and more specific than a single number.

Audit information We wanted to see how participants reacted to seeing information about who had previously viewed albums and the impact varying levels of detail had on user perceptions of the people doing the accessing. In some interfaces we presented large amounts of detail, listing information such as how long a person had looked at a photo, and the date/time they looked. Some interfaces were more general showing only a total number of views or unique viewers.

Layered data We recognize that it is impossible to put all the information a person might be interested in onto the main interface. To accommodate for this we *layered* the data, making some information visible on the main screen, some information visible only on moue over, and some information visible only if a participant clicks on a link. We wanted to test several styles of layering data, types of data, and level of detail to show. Interactive layered features were demonstrated by the researcher during the presentation of the interface.

4.3 Results

Participants' comments were collected and transcribed from the packets and the audio recordings. The transcribed comments were sorted by topic using an affinity diagraming methodology [20]. The high level take aways from the studies are detailed below.

4.3.1 Why is privacy important to me?

Similar to what others have seen from the college-age demographic, we observed an interesting mix of participants who consider anything posted on the Internet to be potentially public and participants who want fine-grain control over who can see what [2].

The difference between these two view points appeared to be based in two mental models of how information uploaded to websites was going to be treated by companies. Participants who considered everything on the Internet to be public had the view that the companies they gave their data to were going to loose it, change their privacy policy without warning, or in some way accidentally expose their data. These participants did seem to use individual privacy settings, they just felt that the settings expressed what they wanted to have happen, not what was actually going to happen. For them it was important to only place data online that would have limited negative impact if it became public, and spending significant time setting up a detailed policy was a waste of time. This view point was more predominate in the younger participants and appeared to be drawn from experiences with companies like Facebook which has been known to retroactively change the visibility of pre-existing data, a practice which contributed to a ruling against them by the Federal Trade Commission [3]. Existing research also shows that people loose faith in a company after what they consider to be a betraying change to the company's privacy policy [26]. Participants with the other metal model were more optimistic about how online companies would protect their data. These participants generally agreed that putting embarrassing photos online wasn't smart, but they also felt that actively managing their privacy settings would make a difference and would be honored by the websites.

Participants who consider everything online to be public tended to dislike the idea of proximity information displays because they cluttered the screen with "useless" and "creepy" data. These participants also tended to refer to people who want to control access to each album or picture individually as "control freaks," or "micromanagers." This is very similar to the culture Gaw et al. observed where people who take extra security precautions tend to be considered "paranoid" [42]. These participants tended to respond to privacy and security information by suggesting that we remove it or change it to a small icon similar to what is now used on Facebook and Google+. One such participant wrote in his packet "Want a lock sign."

Participants who felt there was a need to control access to pictures tended to like the idea of locating information about who could and who has seen the pictures next to the pictures. They felt that proximity information displays made the information easier to find and easier to understand. In the words of one such participant: "I like that you can immediately see who viewed your pictures without necessarily accessing the album in question." They expected that being able to see the controls would better enable them to identify issues and be more aware of their settings.

4.3.2 Who has viewed my photos?

One element we wanted to experiment with was the effect of showing people not only who could see their photos but who had seen their photos. We refer to this information as *audit*

information. We felt that this type of information would assist people in re-evaluating their policies and finding issues by giving them a sense of how their permission settings were actually being used. However, our focus group participants strongly disliked being shown audit information.

A participant in the fifth focus group explained the concern: "Too much specific information about who has been seeing what, it makes me uncomfortable as a poster and as a viewer. I don't want stalking information available." Several participants across the focus groups brought up the term "stalking" in reference to detailed audit information. Participants felt exposed and uncomfortable by the data both from the perspective as a viewer, but also as a photo owner. The issue stemmed not only from the data itself but also its location. Participants felt that this data might be acceptable if the photo owner had to go out of their way to find it, but by placing it on the main interface we were encouraging stalking and taking away their right to choose to see the information or not. One participant explained it as "You are forcing me to stalk my friends."

The more detailed the audit information the more concerned participants became. One display showed the date, time, and duration of every view. One participant circled this display and wrote: "Creepy!!!" The concern wasn't just with the data exposure, it was also with how others might misinterpret the data. A specific example concerned the duration of a view, one participant commented that it was creepy that someone would view the photos for an hour, then a different participant pointed out that they might have walked away from the computer. This sparked a conversation about how data might be misinterpreted and thereby cause someone to think ill of another when no wrong had actually been committed. Participants were very concerned about how others might misinterpret actions taken online.

Participants were also concerned about trying to extract permission setting information from the audit information. They didn't want to accidentally confuse situations where a person used to have access, and therefore showed up in the audit information, with the person currently having access. Early focus group participants primarily saw interfaces that showed detailed audit data or detailed permission setting information. Due to concerns about screen real estate we showed one or the other but not both. In latter studies we placed specific current permission information with specific audit information and saw less concern about confusing if someone could currently see the album or not.

While the majority of participants did not like audit information, there were a few participants who saw the information as potentially valuable if used in the right places. These participants commented on usefulness of understanding who had seen their photos and the interface components we had added that explained why that person could view the photos. In the words of a participant: "I like the additional features that help you see who viewed your photos and why (what groups they are a part of ect.)"

Participants hated specifics about who had viewed their photos but they loved high-level statistics, particularly information that helped them answer the question "Is my photo popular?" On interfaces with no audit data shown we would get comments asking for the number of views or some very high-level sense of what was popular and what was not. Because of the push back about audit data being specific we tried showing participants a display which graphically sorted items into most to least viewed, and another that showed sparklines. However, participants made it clear that they really wanted the specific number

of times their photo or album had been viewed.

4.3.3 Who could see my photos?

Participants were more positive about seeing information concerning who could see their photos than who had seen their photos. They considered this information useful and enabling. Similar to other topics, what participants liked or found useful about this information depended on their mental model of how effectively companies would honor the settings associated with the participants' data.

Participants who were convinced that all their photos were essentially public anyway generally considered the majority of the information shown to them to be "irrelevant" or "unhelpful." They especially disliked the amount of screen real estate the various displays required. These participants preferred the idea of using an icon or something very small and high level to express the policy. If they needed to know more they felt that they could easily click through to some other screen and see it.

Participants who considered permissions to be worth setting liked how easy it was to see their policies. They felt that the displays made it easy to change their policies, indicating that the idea of using displays as a segue to modifying permissions made sense. In the words of one participant: "I like the control over who can see [the pictures] and how simply that control is apparent." Participants liked the comprehensive approach to policy display. Interfaces that used words to explain the information shown tended to be liked by participants and considered to be easy to understand. More detailed interfaces or ones that made heavy use of icons, explained via layered tool tips, were less well liked. Participants were concerned that they, or less computer savvy people, would not understand the meaning.

While the idea of showing who could view a photo or album was generally liked, participants were concerned with the detail and space required to show the information. They disliked showing individual names of people who could view the photos, because they were concerned that the visualization would not scale well. They also talked about how they thought about the people they shared with as groups not individuals: "I think about my friends in clusters, bicycling, activist, college." Participants were also concerned with their ability to think about that number of people at once, and the consequences if they forgot someone: "When controlling who can see what on a per person level you have to be aware of every person. If someone is not able to view something you can have hurt feelings even if it is easy to change settings." Participants were of the opinion that by using groups the interface would be easier to glance at and the policies easier to manage.

4.3.4 Proximity displays in personal and work environments

In the previous sections we have discussed participants' reactions to different aspects of proximity displays in an online photo-sharing environment focused on sharing personal photos. In the last two focus groups we added two website designs based on document sharing to the end of the website lineup. We were curious if participants would react differently to proximity displays in a more work-oriented domain. Participants' comments

on the document websites were nearly a complete reversal of their previous comments and concerns.

Audit information, which was heavily disliked in the photo-sharing context, was in high demand in the document sharing context. It was considered creepy and stalkerish to look at who had viewed photos, but it was considered very useful to see who had interacted with a document and the exact type of interaction conducted. If changes had been made to the document then participants wanted to see the exact changes. Unlike photos, the more specific the audit information the better.

Detailed information about who could view documents was also in high demand. We showed a version of the Expandable Grid [83] to the participants in focus groups 1, 2 and 3 on the photo-sharing website and focus groups 4 and 5 saw the grid on the document-sharing website. Participants disliked the grid on the photo-sharing website because it was "too much information," but loved it on the document-sharing website: "Loved the grid concept, good to know if everyone in group has been looking." Participants who considered the grid "too much for the main [document] interface" suggested it be moved to the documents main page, not completely removed from view like the photo-sharing participants had suggested.

4.4 Conclusion

In this chapter we have explored how people react to different types of access-control related information when it is presented on the main screen of a photo-sharing website. We found that the mental model participants have concerning how their data will be treated and protected by websites impacts the types of information they perceive as useful. Participants who feel that their data will likely be exposed have limited interest in showing permission information on the main screen, while those who feel that their settings will be honored are more interested in seeing this information. Regardless of their belief in permission setting effectiveness, participants found information about who had previously seen photos to be creepy and similar to stalking. Conversely information about who could, in the future, view their photos was considered to be useful and enabling.

Chapter 5

Proximity access-control information 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.

In this chapter we introduce the proximity information display designs that we tested as part of our research. We discuss our motivation and design approach and then detail specific features of our designs.

5.1 Theoretical Approach

Proximity information displays for access-control make use of policy relevant information to inform the user about how their files have and could be used. The displays also make use of policy relevant data that may be interesting to the user for curiosity reasons to encourage the user to periodically review the policy information and identify errors.

Based on our previous work and review of the literature, we set out to design proximity displays with the following goals:

- Users need to be able to identify permission errors.
- Users need to be aware of their current permission settings.
- Users need to be able to easily access and change the permission settings.
- Users need to be able to see when other people have accessed protected data.

5.1.1 Scenarios

To help the reader better understand how we expect access-control proximity displays will fit into the normal work flow of users we present three scenarios, two of which were used in our focus group sessions, where a user interacts with a proximity display.

Alice wants to select some photographs for her screen saver. She goes online to her photo albums and starts looking through some of her albums including some photos from a trip to Chicago she took with her good friend Sue. While looking at the album Alice notices that Sue has recently viewed the album. Alice hasn't talked to Sue since they returned from the trip. So she sends Sue an email to catch up on events since the trip.

Joe goes to his online photo sharing website to share some pictures from his latest vacation. He uploads all the photographs into a new album and starts going through them to make sure they are all correctly oriented, have good titles, and generally look good. While going through the photos he notices, via the proximity display, that his ex-girlfriend can see his new photos. Joe is very upset by this and immediately wants to make changes to his privacy policy. He uses the link on the proximity display to open the permission-modification interface for this album. He changes his policy so his ex-girlfriend cannot see his new vacation photos, and then returns to his new album. He uses the proximity display there to double check that the ex can no longer see these photographs. He then returns to checking the status of his new photos and making sure they are presentable.

Sam likes looking through all the comments people make about her photographs. Sam takes great photographs and enjoys having other people comment on them. As she is going through her most recently posted album she wonders who she shared this album with. Sam generally doesn't change the access-control settings on any of the websites she uses. She trusts that the websites are popular and likely have good defaults. Since nothing bad has happened she is disinclined to waste time looking through multiple pages of settings. However, she realizes that the settings are visible on the same web page as her album photographs. She glances at the proximity display, primarily out of curiosity, and realizes that her poetry group can see these photographs. Sam didn't expect that the poetry group could see her photos, but she decides that she is fine with them viewing her work, so she returns to reading comments. Later she writes a poem about a quirky apple she photographed, and points her group to the photo knowing that they can view it.

5.1.2 Design properties

In order to be effective proximity information displays have the following properties:

Spatial proximity - Information about who has and who could view an item is located physically close to the item to which the information pertains, where it can be easily noticed.

Glanceable - Easy for a person to comprehend the proximity display quickly find important information.

Useful and interesting information - The information presented should be interesting to the user so that they want to look at it.

Layered data exploration - Only the most important or most motivating data is shown on the proximity information display, other relevant and important data is shown only if the user clicks on or hovers over something that they want more information about.

Easy to segue to policy modification - Once a user notices an error it should be ob-

vious how to proceed with fixing it. The design of policy-modification interfaces is not the subject of this thesis but the proximity-information displays should provide a clear path showing the user how to begin the policy-modification task.

By placing relevant and interesting policy information in close spatial proximity to the items with which a user is currently working I will enable the user to easily perform the "Attention Switch" and "Attention Maintenance" steps in the C-HIP warning model [112]. By making the information easy to glance through I will support the user's ability to quickly decide if any of the presented information is important or interesting, thereby keeping or loosing their attention as appropriate.

5.2 Proximity display design

The proximity-display designs used in this thesis fall into three categories. The first was the initial grid based design, showing only permission setting information, that we showed to users as part of the first and second evaluation studies. The second was a list based design that also showed only permission setting information. We moved from grid to list designs to reduce user confusion over the symbols and to make the display more compact. The third type of design also used a list but instead of permission settings, it showed information about who had previously interacted with the photos.

5.2.1 Grid based design

The initial proximity display design (Figure 5.1(a)) shows access-control policy in grid form, with each row of the 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 5.1(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 5.1(a), for example, mousing over the icon next to "Everybody" would display "The group Everybody cannot view Animal Shelter Shared Albums."

The idea of using close spatial proximity to link concepts is well known and part of Gestalt principles [39]. These principles describe how humans visually group and associate objects, visual objects which are in close spatial proximity are considered to be related. We use this principle 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

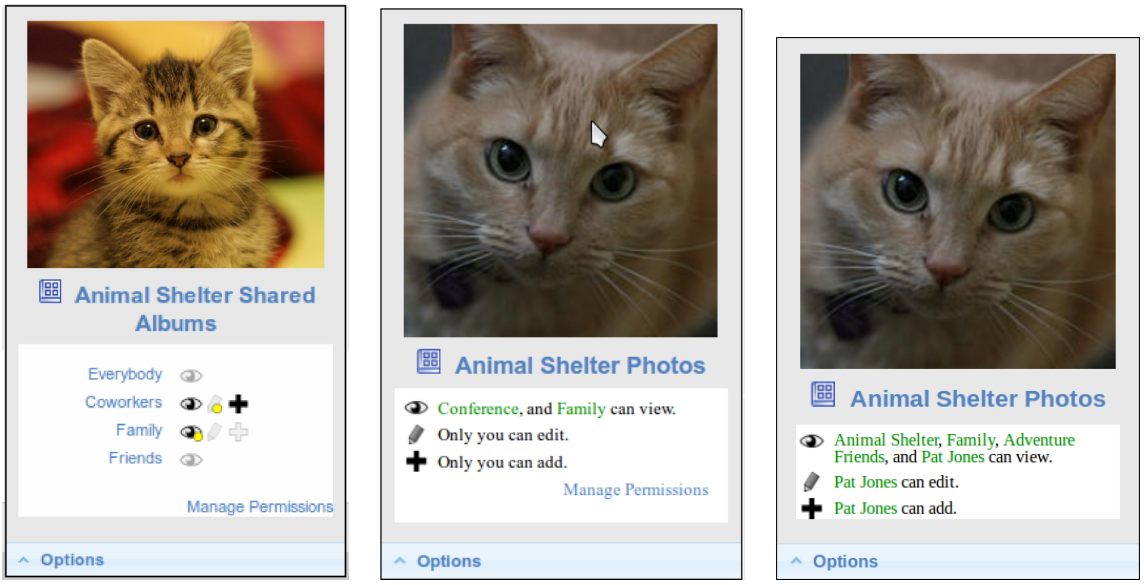

(a) (b) (c)

Figure 5.1: The proximity displays shown to users in the four evaluation studies. The display used in the first and second studies (a), is based on a grid style design. The displays used in the third (b) and fourth (c) use a list-based design. Displays (a) and (b) include a "Manage Permissions" link, participants were rarely observed to use the link, so it was removed in design (c).

file permission settings [83, 84]. 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 privacy notification layouts, intended to be shown during application installation, against participant comprehension speed [101]. They found that layouts 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. Tam's work used a different set of actions and the layouts were not for proximity displays, however, we have a similar goal in that we want people to comprehend our displays quickly.

5.2.2 List-based design

The list-based design came out of our experiences with the grid design. We found that while the symbols readily made sense to most people, other people became confused and therefor set the policy incorrectly. The displays which explained each icon when it was moused over proved to be of limited assistance to participants with incorrect mental models. The researcher also observed that glancing at the grid was challenging. The grid design aesthetic meant that a user had to first read the group name on the left then mentally connect it to the icons next to it. While this was easy to understand, we felt the design was overly challenging for users to take in "at a glance." Too much focus was required

to parse the information. Finally, the grid based display was fairly large and had limited scalability.

The list based design (Figure 5.1(b)) incorporated the same icons as the grid (view, edit, and add photo). However, these icons were displayed statically and were only intended to graphically indicate the action the list refers to. These choices were consistent with layouts Tam et al. showed to be effective [101]. Even if no groups had access, the icons where shown. If child albums had different permissions from the parent this was indicated on a completely separate line that just read: *some subalbums have different permissions.*

The list itself is designed to be both understandable and scalable. Where the grid design was organized with each line corresponding to a group, this design made each line correspond with an action (view, edit, add). The icon at the beginning of each line visually indicates the action and makes it clear where the beginning of each line is. The group names are embedded in a sentence that clearly states, in words, what action these groups can engage in. The group names are shown in a different color to make them easy to visually separate from the static parts of the sentence. This was intended to assist participants with pattern matching the word shapes. The design is also scalable. If there are too many groups to list then the display shows "and 3 more groups" as a link so participants can easily see all groups without switching pages.

In the last evaluation study we decided to make the display appear on the screen at all times. In the earlier designs (Figures 5.1(a) and 5.1(b)), when the display was located under an album or photo thumbnail, it was only visible if the user placed their mouse over the thumbnail. When it appeared elsewhere on the screen it was always visible. We made the display under album/photo thumbnails always visible to make the displays more comparable with other designs such as icons (further discussed in Section 7.3.1). Researcher observation of how people used the designs in Figures 5.1(a) and 5.1(b) suggested that they were ignoring the "Manage Permissions" link and instead using the link on the Options menu. Log data supported this conclusion, so to save space we removed the "Manage Permissions" link from Figure 5.1(b), resulting in Figure 5.1(c).

5.2.3 Audit-based design

One of the goals of the thesis is to help people not only be aware of their current access-control policy but also assist them in re-evaluating it based on past performance. We received some negative critiquing from the focus group participants, because they felt that the idea was creepy and a bit stalkerish. However, we received more positive feedback when we discussed using it in a work type environment. So we feel that this is an interesting direction to explore.

The resulting audit-based proximity display design, shown in Figure 5.2(b), was designed based on the reactions from focus group participants, pre-study lab participants, and the way people interacted with the grid and list based proximity interfaces.

The focus group participants were very clear that displaying detailed information about who had viewed their photos made them feel uncomfortable. However, they were interested in understanding how popular their albums were. They also felt that audit type information was more acceptable if it was either general, or on a separate page where someone would

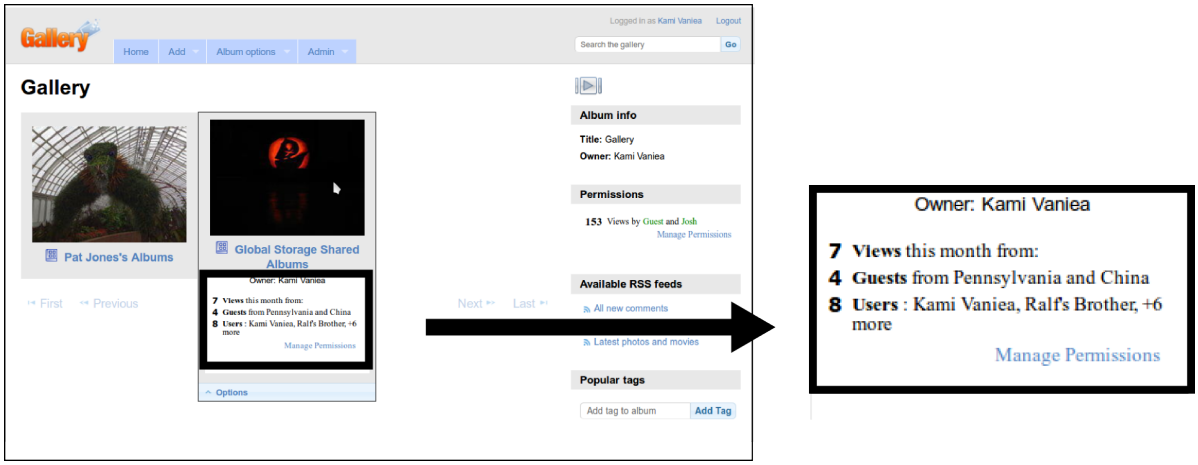

(a) Audit proximity display shown in the 2nd evaluation study.

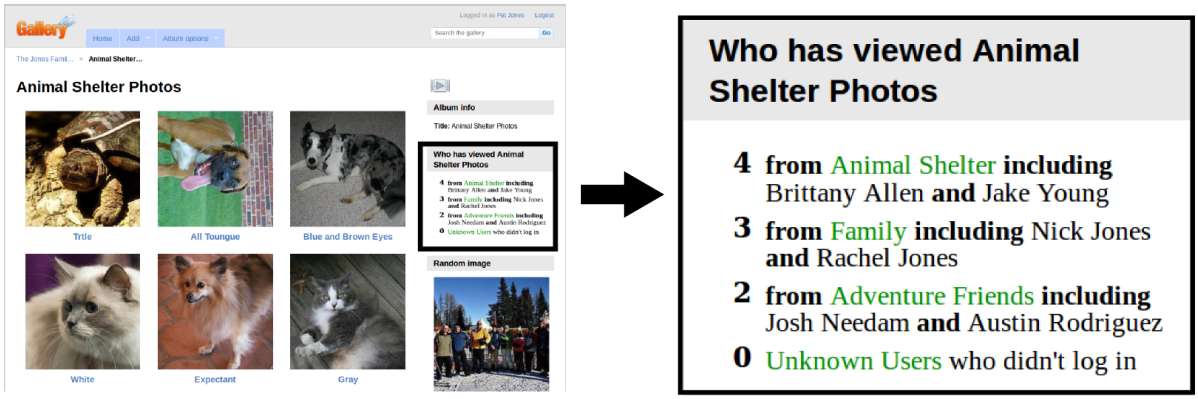

(b) Audit proximity display shown in the 4th evaluation study.

Figure 5.2: The proximity displays showing who had accessed the album (audit). Unlike the displays shown in Figure 5.1 which show who could access the album in the future, these displays show who has accessed the album in the past. Figure (a) was pilot tested during evaluation studies 2 and 3, resulting in the Figure (b) design, which was evaluated in the final evaluation study (study 4).

have to actively attempt to find it.

Our initial audit-based design, shown in Figure 5.2(a), was intended to make it clear how many times the album had been viewed (popularity), and who had seen it. When we tested this design with users they rarely interacted with it. When asked, they stated that the information shown did not help them find errors, so they had ignored it. Our focus group participants had similarly expressed concern about the helpfulness of displaying individual people's names, stating that they thought about their friends as groups.

The display used in the forth evaluation study (Figure 5.2(b)) is organized by groups with each group displayed on a single line. Only groups who's members have accessed the album are shown, with the exception of "Unknown Users," which is always visible. When we talked to end users we found that they cared about random people on the internet viewing their photos. So, to add some consistency between the displays, and to reassure users, we added "Unknown Users" to every display. The group name is shown in a contrasting color so that a user can easily see which groups have access. The number at the front of the line indicates the number of people in that group that have viewed the album. Focus group participants expressed a disinterest in details, including people's names, but we wanted to encourage scenarios similar to the Alice use case in Section 5.1.1. To balance these conflicting goals, we decided to put the names of the group members who had accessed the album in a less intense font so as to deemphasize them. Additionally, we only list the people who have seen the album on the album page but not on any of the photo pages.

We observed users working with the audit-based design during the second and third evaluation lab studies, and tested it on the online evaluation study (study 4). We initially wanted to test the audit display during the second evaluation study but when we put it in front of pre-test participants they ignored it as irreverent information. This type of interface is challenging to test in a role play style lab environment because participants are working with a fictitious ideal policy and only in the lab for 1-1.5 hours. In that time frame there isn't really time to observe participants re-evaluating their permission choices. In the online evaluation study we wanted to know if the audit design would be useful for identifying permission errors.

5.3 Access-control policy modification interface

We were concerned that the default permission-modification interface for Gallery might be confusing for end users. We also wanted to make sure our proximity information display design sufficiently matched the permission-modification interface so as to not confuse our users. To this end, we created two permission-modification interfaces: a full-page interface which shows the access-control policy for all albums on a single page, and a dialog interface which shows only the access-control policy for a single album in a pop-up dialog.

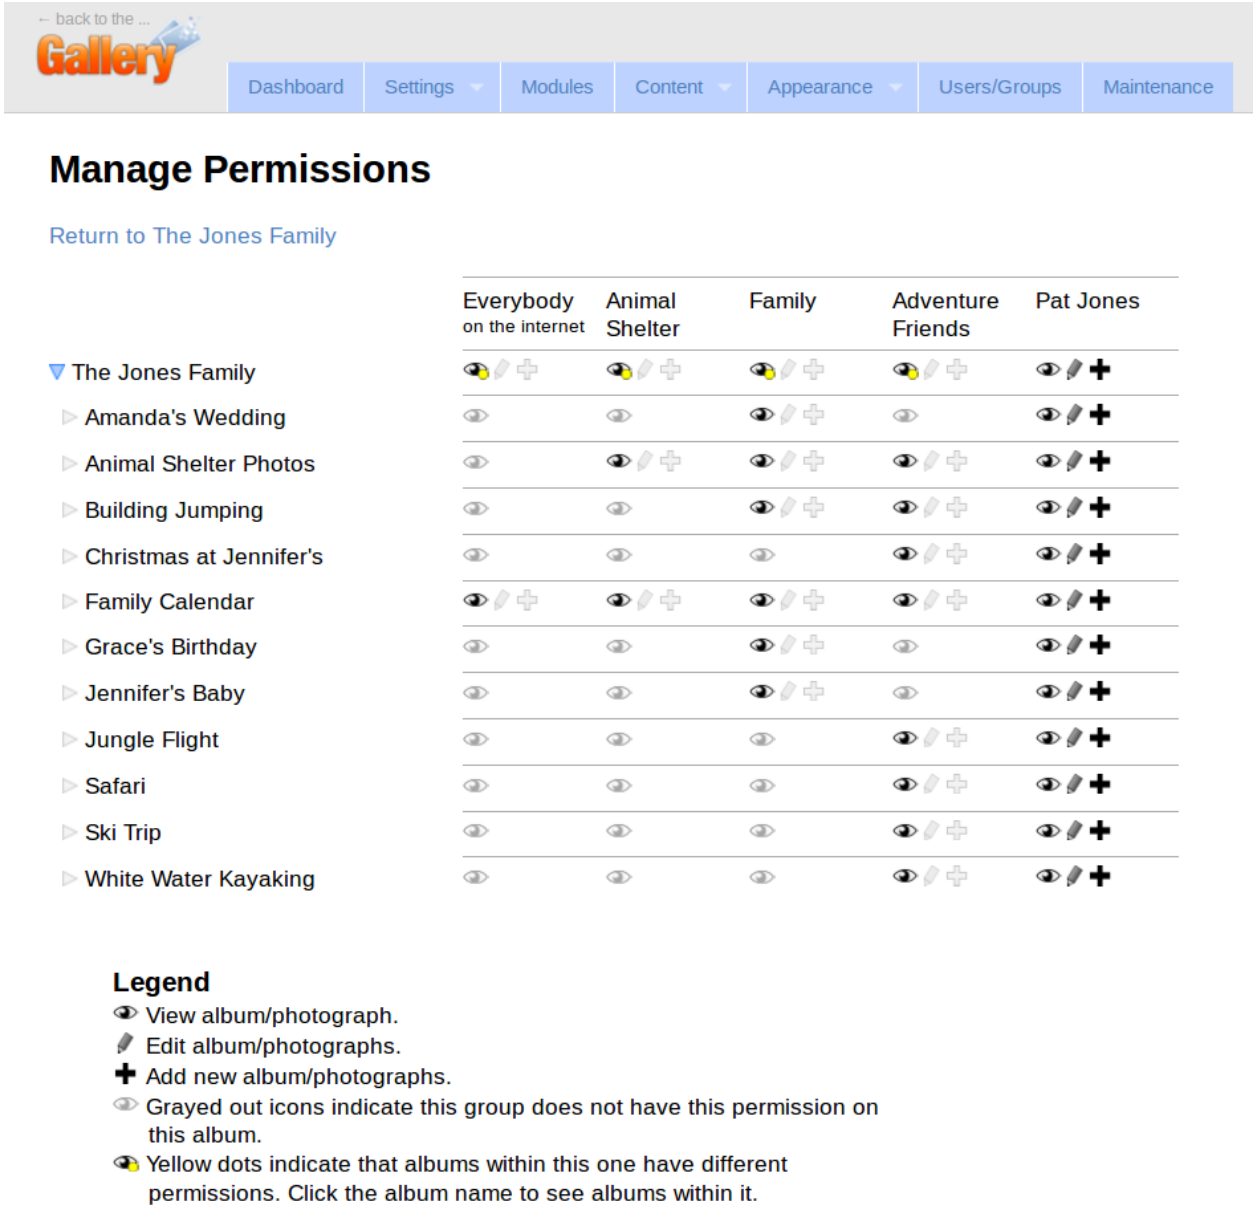

Figure 5.3: Full-page policy-modification interface used by participants to make changes to the access-control policy. All the albums are listed along the left; user groups are listed along the top of the grid; and view, edit, and add permissions are shown as icons in the central grid. This interface also contains a legend at the bottom left.

5.3.1 Full-page interface

The full-page permission-modification interface (Figure 5.3) displays the access-control policy for all the albums in Gallery on a single page. This interface was designed based on Reeder's Expandable Grid [83], which was shown to be effective in assisting users in understanding and accurately managing their access-control permissions.

Each row of the grid is associated with an album, and each column is associated with a user group. Each cell contains one or three icons indicating the actions this group can current perform on this album (black icons), as well as the actions the administrator can currently grant (grey icons). The add and edit actions are not possible when the view action is denied, so the icons for these actions are not shown at all. For example: The *Animal Shelter* group cannot view *Amanda's Wedding* photos; consequently they also cannot add or edit this album. Giving the *Animal Shelter* group the ability to add without the ability to view would not actually give them any rights, since viewing is necessary to add. Hence, the icons for add and edit are removed.

A user can grant/deny any action by clicking on the icon, which will change it from black to grey or vice versa. To assist users in understanding the icons, the interface includes a legend in the bottom left.

5.3.2 Dialog interface

The dialog permission-modification interface (Figure 5.4) shows the permissions associated with a single album. The display opens as a JavaScript dialog box, so the user can easily view permissions and make changes without having to switch pages.

The design of the dialog is intentionally very similar to the grid-based proximity display. We use the same icons, organization, and mouse-over effects. Similar to the full-sized permission-modification interface, the user can change the access-control policy by clicking on any of the icons.

5.3.3 Conflict resolution and effective permissions

By default, Gallery shows users the access-control policy rules rather than effective permissions, but this makes it very difficult for a user to accurately understand why a particular group has access, or how to change the permission. When we designed the proximity displays and permission-modification interfaces, we decided to show the user effective permissions (the result of evaluating all relevant policy rules) rather than the sets of policy rules that induce them. Prior work by Bauer et al. [15] and Maxion and Reeder [69] has shown that people better comprehend access-control policies when they are shown *effective permissions*. In our design we show users effective permissions, and allow them to change permissions by indicating the effective permission they wish to change.

Albums in Gallery can contain subalbums, and permissions on the parent album affect its children, but can be overridden by the permissions set on the children. Similarly, Gallery has two built-in groups: Everybody, and Registered Users. The group Everybody includes all users with accounts on the website and all guest users. The group Registered Users

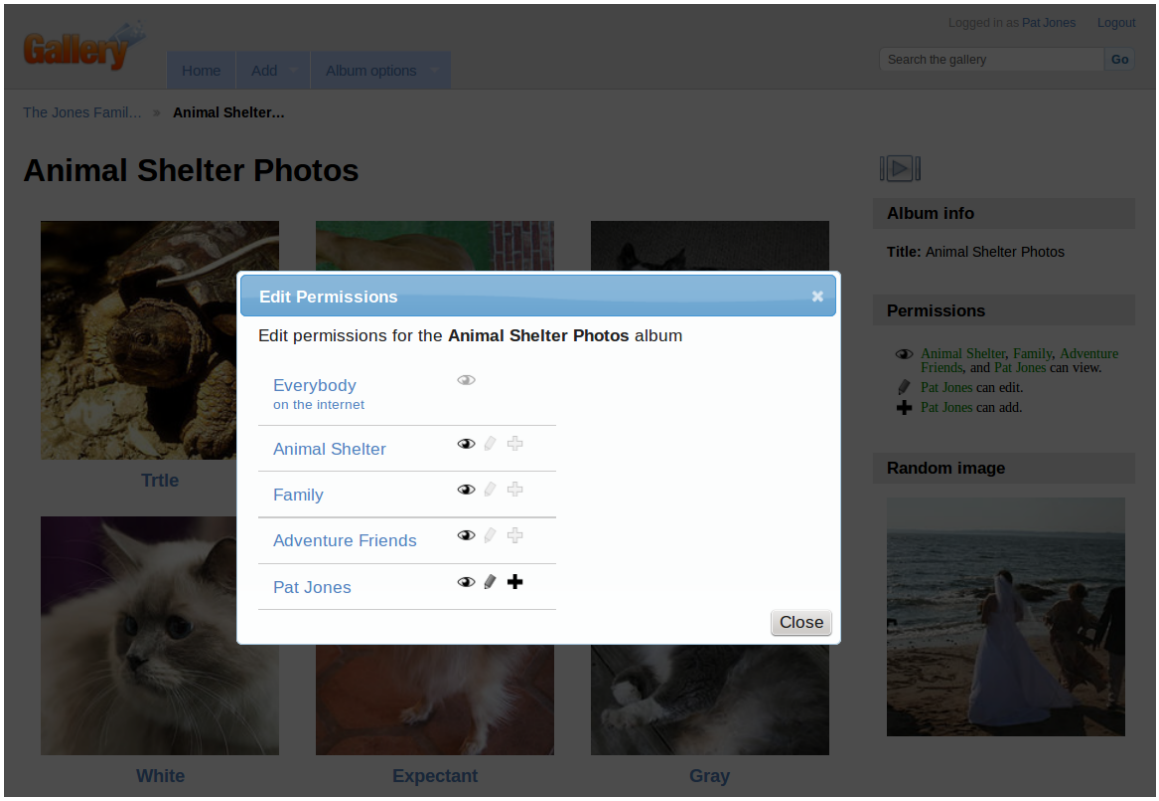

Figure 5.4: Permission-modification dialog. Sentence at the top of the dialog reminds the user what album the permissions refer to. The group names are listed along the left side, followed by the different actions (view, edit, and add) that are allowed or denied for that group. A black icon indicates that the permission is allowed; a light grey icon indicates that the permission is denied. Placing the mouse over any icon produces a tool tip indicating the meaning of the current icon. For example: "Animal Shelter can view this album." Clicking on an icon toggles it between allow and deny.

includes all users who have an account on the website. We decided that this group caused needless confusion and removed it for the evaluation studies. The existence of these groups effectively makes groups hierarchical, because permissions set on these built-in groups effect the permissions on the other groups.

To address potential permission rule conflicts, Gallery uses the conflict-resolution strategy shown in Table 5.1. A conflict can occur any time two or more rules apply to the same user group, album, and action, but have different outcomes (allow, deny). When this happens, the system must decide which outcome to use. In Gallery, if a person cannot view a parent album, then they cannot view any child albums regardless of the permissions on the child; however, if they can view the parent, then they may or may not be able to view the child based on the permissions on the child. Conflicts in the user group dimension always result in a decision of allow. If a user is a member of any group which can view the album, then the person can view the album regardless of the rules on other user groups to which she may belong.

		Relationship between ALLOW rule's user group and DENY rule's user group			
		Contains	Peer	Same	Contained by
Relationship between ALLOW rule's album and DENY rule's album	Contains	DENY	DENY	DENY	DENY
	Same	ALLOW	ALLOW	ALLOW	ALLOW
	Contained by	DENY	DENY	DENY	DENY

Table 5.1: Conflict-resolution strategy used by Gallery version 3.1.

Displaying effective permissions on the proximity displays and permission-modification interface worked well: participants understood the current permission state. Enabling intuitive permission modification was more challenging because participants were attempting to manipulate effective permissions, rather than specifying rules in which access-control is implemented in Gallery. We addressed this issue by a translating users' effective permission change requests into sets of rule changes.

When a participant indicated that they would like to change an effective permission, toggling it from deny to allow or vice versa, our algorithm computed the set of rule changes necessary to produce the least number of effective permission changes. For example, assume the album Animals, which has subalbums Dogs and Cats, was not visible to the group Family. The user indicates that they would like Cats to be visible to the group Family. Our algorithm would add an allow rule for (Cats, Family), and (Animals, Family) in order to give Family the ability to view Cats (the parent album must also be visible to Family). The algorithm would also create a deny rule for (Dogs, Family) to ensure that the effective permissions on the subalbum Dogs do not change.

Participants seemed to find these side effects intuitive. Some participants were briefly surprised when they clicked an icon on the full-page interface and more than one icon changed. However, most participants quickly realized why the change had occurred and

did not seemed bothered by it. Participants who saw the dialog interface only noticed that multiple permissions were changing when they tried to remove permissions from a group and the permissions were also removed from the group Everybody. However, similar to the full interface, they rapidly determined the reason for the change.

Chapter 6

Designing an access-control study were security is a secondary task

In four proximity information display evaluation studies, we investigated an interface intended to help users stay aware 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. [1].

Our primary hypothesis was that if the current permission settings are shown in close spatial proximity to the resources they control, instead of on a secondary page, users are more likely to notice and fix permission errors. To test our hypothesis we need our participants to interact with the display as as secondary task, where they have a non-security primary task and interacting with permissions is secondary.

Other researchers have studied security as a secondary task using various approaches [48, 100, 107]. One approach, used by Haake et al. [48], is to conduct a long-term study where the participant is made aware that security is a part of the study but the study is run for long enough that the user stops focusing on security. Another approach, used by Sunshine et al. [100], is to not make the participants aware of the security nature of the study, but the study design forces participants to engage in a security behavior while trying to complete their primary task. A final approach, used by Wang [107], is to keep the participant unaware that the study is about security and give the participant the option of whether or not to interact with the security functionality.

To test our hypothesis we decided to use the last approach. We conducted a lab study where participants performed various photo management tasks. Depending on condition, participants were shown permission information under the photos, elsewhere on the page, or on a secondary page (control).

When designing the initial study methodology, we wanted to meet the following goals: make security a secondary task (Section 6.3), give the participant ownership/responsibility

[1]This chapter is based on a published paper: K. Vaniea, L. Bauer, L. F. Cranor, and M. K. Reiter, *Studying Access Control Usability in the Lab: Lessons Learned From Four Studies*, Proceedings of Workshop on Learning from Authoritative Security Experiment Results, 2012 (to appear)

for the albums (Section 6.4), make sure the participants understood the policy they needed to enact (Section 6.5), and develop clear metrics for measuring the outcomes (Section 6.6). Despite careful planning we encountered methodological issues on every one of these goals.

In this chapter, we discuss this study and three subsequent ones, each of which took into account the methodological issues that arose in the proceeding study. We focus our discussion on aspects of the methodology that tried to accomplish the four goals described above. We describe the difficulties encountered during each study, and changes to the methodology designed to address those difficulties. Through this process, we shed light on the challenges intrinsic to many studies that examine security as a secondary task, and convey a series of lessons that we hope will help other researchers avoid some of the difficulties that we encountered.

6.1 Study Goals

The purpose of all four studies was to test the hypothesis:

H: Users who see information about access-control permission settings on the main interface notice permission errors more often than users who have to proactively open a second interface to view permissions.

When designing study 1 to test H we wanted to create a study environment that met the following four goals:

Secondary permission task Participants should be in an environment where there is little encouragement to engage in security tasks and the benefits, if any, are not immediate. Users treat security as a secondary task because the benefits of security are often hard to envision but the cognitive and time costs of engaging in it are immediate [109].

Other researchers who study security technologies have successfully simulated the secondary task mindset in the lab. Whitten and Tygar's work on email encryption had participants focus on sending and receiving emails while they measured the usability of PGP [111]. Similarly Sunshine et al. asked participants to find information on websites while studying their reactions to SSL errors [100].

Participant responsibility Participants should feel they are sufficiently responsible for the experimental content to be comfortable making changes they deem necessary. Because changing permissions is secondary, the framing of the study should make it clear to participants that they should make changes outside the bounds of their primary task.

When replicating the SSL study described above, Sotirakopoulos et al. experienced issues with participants claiming that the lab was a "safe" environment so they behaved differently [95]. Witten and Tygar overcame this issue in their work [111], but doing so requires careful study design.

Ideal-policy comprehension Participants should be aware of and comprehend the *ideal policy* – the correct set of permissions for the content. The participant needs to have a clear ideal policy associated with the content they are working with. Participants need to be able to consistently decide when a permission setting is "correct" or "wrong."

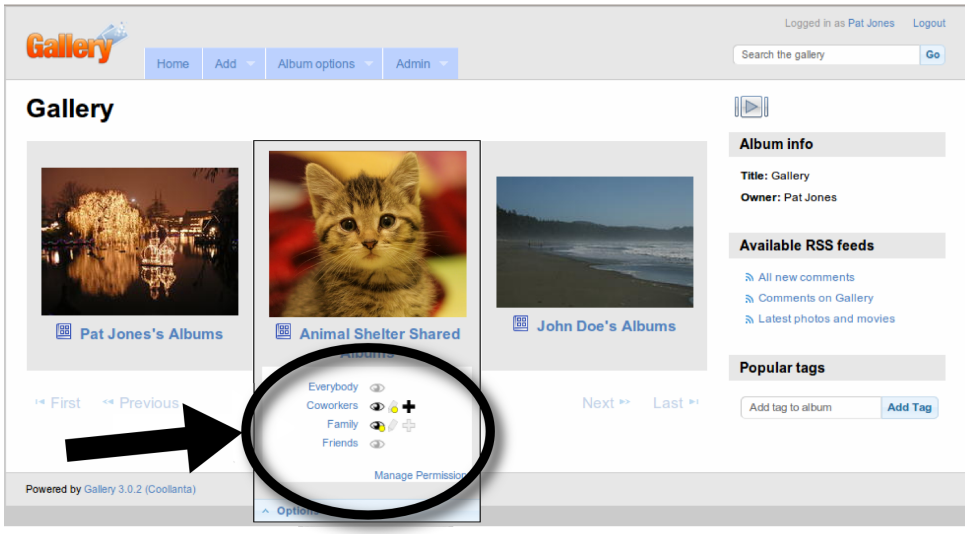

Figure 6.1: Example of proximity display used in studies 1 and 2. The interface for studies 3 and 4 had a slightly different permission display interface design.

Effective outcome measurement We need to be able to accurately measure if participants are noticing and fixing errors. In real world environments the presence or absence of an error can be very subjective and dependent on context [13, 27, 72]. To accurately test "noticing" errors, we need to be able to differentiate between environments with no errors, environments where participants are not noticing errors, and environments where errors have been noticed.

6.1.1 Overall System Design

We decided to use a photo-management website as the domain because it is a common type of website where end users might set access-control policy. We chose to use an open source photo-management website software, Gallery3 (version 3.1) [1], because it was easy to modify and unknown to general users, thereby ensuring minimal bias from prior experience or training.

We built a Gallery module which displays permission information in a small display that appears under the photos/albums (Figure 6.1), or in other parts of the interface. We also built a new permission modification interface that shows the permissions for every album on a single page (Figure 5.3). The permission modification interface was designed to be easy to use and comprehend based on prior work [83, 84] and was not the focus of this research. Access-control permissions in Gallery are expressed as four-tuples of (user group, album, action, decision). Permissions cannot be expressed for individual users or photographs.

6.2 General Study Design

Our initial study design was intended to test the following hypotheses in addition to our main hypothesis H.

H1 Users who see permission information under photos/albums notice errors more often than users who see permission information in other spatial locations.

H2 When a permission is changed to an error state by a 3rd party, users who see permission information under the photos/albums or on the sidebar notice errors more often than users who see permission information only if they click to a second page.

H3 The type of error, too many permissions or too few, has an effect on the number of errors noticed.

H4 Participants who see permission information under the photos/albums or on the sidebar can recall those permissions better than participants who see permission information only if they click to a second page.

H5 Participants in each of the conditions take the same amount of time to complete each task.

In this work we discuss the methodologies of four similar studies briefly. More detailed methodology descritipns are given in Chapter 7. In this section we present the core methodology used in all four studies. In the following sections we detail the unique methodological choices made in each study to meet the goals described in Section 6.1. We discuss the outcome of the choices and how they informed the methodological choices in the next study.

The first three studies were between-subjects lab studies and the last was a within subject online study. All studies used a round-robin assignment to experimental conditions. Participants in all conditions performed the same tasks. Each study had a slightly different set of conditions, but two conditions were present in every study: the control condition was the default interface, which included a link to the interface for changing permissions; the under-photo condition additionally included a proximity display under photos/albums (Figure 6.1).

Participants were asked to role play [38, 91, 111] the part of Pat Jones, who manages online photo albums using Gallery. Role playing is a commonly used method of encouraging user engagement. Whitten et al. successfully use it to encourage participants to view security as a secondary task. Tasks were communicated to the participant in email format. In the first three studies the emails were delivered to the participant on paper by the researcher administering the study, in the last study they were shown in an html frame above the website.

Participants started with a training that showed them how to perform several actions on the website including: changing titles, rotating photos, and changing permissions. Participants were asked to perform all actions described in the material to ensure that they understood how to manipulate the interface. In studies 1-3 this training was done on a separate instance of Gallery with fewer albums than the rest of the study. In study 4 the training and the tasks were done on a single Gallery instance.

After the tutorial, participants in study 1 and 2 were given several short warm-up tasks. These tasks were to ensure that the participant had understood the training. It also gave them an opportunity to acclimate to using the interface. Participants in studies 3 and 4 were given 1-2 full task sized warm-up tasks to acclimate to the interface.

The bulk of the studies were composed of a set of tasks presented to the user in sequence. Each task was composed of a set of *subtasks* – issues with the album that the participant is expected to correct to successfully complete a task. A primary subtask was directly expressed in the email and several additional subtasks were implied by observable errors such as rotated photos, misspellings, and incorrect permissions. All tasks contained at least one explicit and one implied title, rotate, delete, or organize subtask intended to distract the participant.

Some tasks were *prompted* in that if the participant failed to correct any subtask, permission related or otherwise, they would be presented with an email pointing out the mistake and asking that it be corrected. *Unprompted* tasks refer either to tasks with no associated prompting or to participant interactions with a task prior to receiving prompting. Participants were unaware of which tasks were prompted until they received a prompt.

Some albums were *changed* mid-way through the study. The participant first interacted with an album and was made aware of the current state, including permission settings. When the participant was distracted by an unrelated task the researcher made changes to the album. The participant was then instructed to interact with the now changed album again.

Finally, participants filled out a survey that asked them to recall permissions for a selection of albums they worked with, as well as non-task albums with correct and incorrect permissions. For each combination of album, group, and permission the participant could answer *True*, *False*, or *Not Sure*. The survey also asked demographic and prior experience questions.

Study 1 was an hour long between-subjects lab study. Participants were given printed training materials that they worked with for about six minutes. This was followed by five short warm-up tasks which took an average of eight minutes in total. Participants were then given 8 tasks which took an average of two and a half minutes each. Tasks appeared in the same fixed order for all participants. Finally, they filled out the survey. There were five prompted tasks and two changed albums. This study was run on 26 participants and three conditions. It was stopped early because of issues with the methodology.

Study 2 was an 1.5 hours long between subjects lab study. Participants were given printed training materials that they worked with for about five and a half minutes. This was followed by five short warm-up tasks, which took approximately 8 minutes to complete in total. They were then given 12 tasks to perform, which took an average of 3.5 minutes apiece. Tasks appeared in the same fixed order for all participants. Finally, they were asked to fill out the survey. There were five prompted tasks and three changed albums. This study was run with 3 conditions and 34 participants, one participant was excluded, resulting in 11 participants per condition. Further details of this study can be found in Vaniea et al. [106].

Study 3 was a 1.5 hours long between subjects lab study. Participants were given printed training materials that they worked with for about five and a half minutes. This

was followed by two large warm-up tasks taking approximately 13 minutes to complete. They were then given 15 tasks in a random order which took an average of 3.5 minutes apiece. Finally, the survey was verbally administered by the researcher, followed by an unstructured debriefing interview. There were three prompted tasks and no changed albums. This study had two independent variables: proximity display and permission modification interface. The proximity display was shown either under the photo (under photo) or not at all (control). The permission modification interface was either a separate page with all permission settings shown or a dialog with only one album's permission settings shown. There were 9 pre-study participants and 33 actual participants in this study.

Study 4 was a hour long within subjects online study conducted on Mechanical Turk. All participants performed training, warm-up, and tasks for both the proximity display condition and the control condition. The order in which participants saw the conditions was assigned round robin. Participants completed a set of training tasks which took an average of four minutes. Then they completed a warm-up task that took an average of three minutes. They were then given 7 tasks, with a maximum of two minutes to complete each of these tasks. Tasks appeared in the same fixed order for all participants. When finished with both conditions they were given a survey to fill out that asked questions about both conditions that the participant worked with. There was one prompted task and one changed album per condition. There were 300 pre-study participants and just over 600 actual participants in this study.

6.3 Secondary Permission Task

Participants should be in an environment where there is minimal encouragement to engage in security tasks, and the benefits, if any, are not immediate.

6.3.1 Study 1

We decided to give participants a primary task that would take the majority of their attention while still being sufficiently open ended that they would consider engaging in other subtasks. We communicated the tasks through printed emails because the structure allowed us to give context, such as the ideal policy, to the task without drawing too much attention to it. To prevent users from perceiving permission content as explicit direction, we stated all permission information in passive voice and all primary subtasks in active voice. For example, the email in Figure 6.2 explicitly asks that the titles be changed, but also implies, that the Friends group needs to be able to view the photos. The ideal policy components, that could not be expressed passively, were embedded in information pages about Pat's friends, family, and co-workers.

We were concerned about giving participants too much permission *priming* – the amount participants are encouraged to engage in permission behaviors. Every time a participant reads or interacts with permission information they are being primed to think about permissions. We compromised by creating three blocks of tasks separated by information pages. Two of the tasks had permission errors and in the third task permissions

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: New photos

Yo Pat,

Here are the better photos from the Building Jumping trip last weekend. Could you put them up on your site? Just set it up like any of your other albums. Also could you title the photos with the people in them? I had the red parachute, George had the green one and of course your's was blue.

When you are finished send me back a link so I can forward it to the rest of our friends.

Thanks,
Josh

Figure 6.2: Email from Pat's friend stating in passive voice that everybody in the Friend's group needs to be able to view the photographs.

were never mentioned. This third task was to give the participant time without permission priming.

To test behavior in the absence of prompting, the first two tasks were unprompted. If the participant did not correct permissions on these albums, the researcher did not make them aware of the issue. Participants were first prompted about permissions after the third task. We prompted here to be sure participants knew what the album permissions were before they were changed by the researcher.

Outcome Participants rapidly deduced that this was an error-finding study and tried to find and correct all the errors. However, none of the participants noticed that the study was solely about permissions. While participants may have been biased to look for errors, only 67% of participants noticed any permission errors without prompting and no participant noticed all the errors. For comparison 86% of the title errors were corrected.

Over-priming participants to identify and fix errors in general may have caused a control condition behavior we termed "checklisting." Participants who checklisted would reach the end of a task, pause and appear to go through a mental check list. One participant did this out-loud, listing all the types of errors she had seen in the training material, making sure she had checked all of them before moving on.

Additionally, many participants never obviously consulted the permission display to determine if there was an error before opening the permission modification interface. We hypothesized that since all emails mentioning permissions were associated with albums containing permission errors, participants always needed to open the modification interface and had no need to consult the display.

6.3.2 Study 2

In study 1 all tasks that expressed permission information in the email had permission errors. Effectively there was no "cost" to checking permissions because participants could determine from the email that there was a permission error. To address this concern we added a new hypothesis:

H6 Participants who see permission information on the main screen are, in the absence of an error, less likely to open the permission modification screen than users who have to proactively open a second interface to view permissions.

New Read-permission tasks We added three new tasks where the email expressed the ideal policy but the current settings matched the ideal policy, so there was no permission error. After this change, 50% of tasks expressed the ideal policy and had permission errors, 25% of tasks expressed the ideal policy but had no error, and 25% of tasks did not express an ideal policy. Two of the new tasks were prompted. If the participant did not obviously check the permissions, the researcher prompted them with an emailed question about the permissions. The new tasks were also intended to test if participants used the displays to determine the lack of an error (H6).

Outcome The addition of the new tasks appears to have reduced permission priming. We observed no participant engage in checklisting type behavior. Additionally, 53% of participants corrected permissions on 3 or less of the 12 tasks before being prompted and no participant corrected all permission errors. In comparison, over 90% of spelling errors were corrected. This suggests that participants were not overly primed to look for permission errors.

The reduction in priming allowed us to observe more subtle issues with our methodology. Participants' permission-checking frequency was impacted by the different tone and wording of the ideal policy in the task emails. Emails with stronger wording resulted in permissions being checked more frequently by participants in all conditions and emails with weaker wording were checked less. This meant that while we had a valid study-wide result, we couldn't compare the permission identification behavior between tasks. The wording strength added a confounding factor.

6.3.3 Study 3

Reducing the number of tasks with permission errors to 50% and providing ideal policy information in the absence of errors appeared to cause less checklisting behavior. However, the wording of tasks caused participants to check permissions on some tasks more than others, suggesting that participants did not have consistent priming. In study 3 we wanted the tasks to provide a consistent level of permission priming independent of the presence of a permission error. We also wanted to maintain the "cost" of checking permissions at a 50% chance of there being no error.

One ideal policy We used a single ideal policy that applied to all albums because it 1) better mimicked normal usage where a single user has a consistent set of requirements, 2) was clearer for the participant to understand than getting a new policy with every email,

and 3) eliminated wording variability since the participant would only see one policy. To counter differences in participant memory, participants were allowed to look back through any piece of paper the researcher gave them, including the page with the policy.

The ideal policy we ultimately selected had five rules, three of which involved access-control permissions. We were concerned that having a single policy that clearly mentions permissions would overly bias participants to look for permission errors, so we tried the protocol with seven test participants. We found that despite the priming, participants infrequently checked for permission errors but frequently checked for the other types of errors mentioned in the rules.

Consistent task structure Previously the emails were two paragraphs and important information appeared wherever it was most natural based on the email content. For this study the first paragraph was contextual only, indicating how it related to Pat but contained no vital data. The second paragraph clearly explained the primary subtask the participant was to engage in.

Unlike studies 1 and 2, the warm-up tasks in study 3 used the same structure and wording style as the other tasks. Based on observations in the prior studies, the tutorial was sufficient for understanding the system and the warm-up tasks were only necessary for the participant to acclimatise to the system and how tasks were presented.

Randomized tasks We decided, with the exception of the warm-up tasks, to randomize both the order that tasks were presented in and which tasks had permission errors. The goal here was to remove any ordering effects and by removing any effect task wording might have on a participant's inclination to check permissions.

Outcome The use of a single ideal policy allowed us to reduce the number of times we presented the participant with permission information. Only 11 of the 31 participants checked permissions on more than 50% of the tasks suggesting that for the majority of participants permissions remained a secondary task.

Our primary concern was that having explicit permission rules expressed in the beginning of the study would overly prime participants to check permissions regularly. Behavior of practice participants suggested that this would not be the case. However, the results of the full study showed that over priming did impact participants.

Our changes to study 2 appeared to eliminate the checklisting behavior observed in study 1 participants, but the design of study 3 brought it back. A graph of number of tasks where control participants checked permissions shows a non-normal distribution with peaks at 0 and 100. The other conditions showed similar distributions. This suggests that the permission priming effected some participants more than others.

6.3.4 Study 4

In study 3 we saw no difference between conditions because participants corrected all or none of the permissions with few participants in the middle. Using a single ideal policy worked well in study 3 as did the mix of 50% of tasks having permission errors. Because study 4 was within subjects, we decided to use a fixed permission order for easier comparison.

Time limitation We hypothesized that in study 3 that providing participants with clearer instructions made it easier for them to know what to do, but the only cost to participants for checking permissions was the time required to perform the check. In real life that time would be an opportunity trade off since the user could be doing something else with that time. In study 4 we decided to limit participants to a maximum of 2 minutes per task, forcing them to value their time and make trade offs. The primary researcher, as an expert user who knew where all the errors were, required a minimum 1.5 minutes complete each task, so we tried 2 and 3 minute limits on practice participants. We determined that a limit of 2 minutes created the largest differentiation amongst users.

Compensation variation For our practice participants we were concerned that Mechanical Turk users would not take the tasks seriously and do the minimum to advance through the study. So we offered a bonus based on performance. However, study feedback suggested that participants were deeply concerned that failure to get everything correct meant they would not be paid. They also felt a level of personal responsibility to get all the subtasks correct. So we adjusted compensation to a single rate and explicitly stated that all participants who got more than 25% of the task components correct would be compensated.

Outcome The combination of time limitations and reduction of emphasis on accuracy worked well. Permissions were changed unprompted by 66% of participants. In the under-photo condition only 4 of the 62 participants corrected all permissions. We also saw a reduction in feedback about the number of tasks participants had correctly completed.

6.4 Participant Responsibility

The framing of the study should make it clear to participants that they could and should make changes outside the bounds of the subtasks expressed in the emails.

6.4.1 Study 1

By having participants role play we were able to inform them that they had a responsibility for some albums by telling them it was part of their job or that their mother regularly relied on them for assistance. We wanted participants to be aware of what types of errors (rotations, spelling, ect.) were within the bounds of the study without overly priming them towards permissions. The tutorial that covered several functionalities of Gallery, included permissions and followed by five prompted warm-up tasks, two of which involved permissions.

Outcome The open-ended nature of the tasks combined with the imparted responsibility made participants uncertain about how to react to tasks and prompts. For example, after a prompt from Pat's mother, in which the mother is panicking about seeing a photo of Pat sky diving, one participant simply responded "Sorry Mom." Another participant asked how old Pat was, then slapped the paper down on the table and declared loudly "I am NOT answering this!"

Some participants didn't feel it was their place to change permissions. A couple of participants noticed an error and verbally decided not to correct it because the album belonged to someone else and they expected that the album owner knew what they were doing, even if the permission was odd. Participants were not instructed to talk aloud during the study so we had no way of knowing how many participants noticed an error and chose not to correct it.

6.4.2 Study 2

Based on observations of participants we theorized that the general uncertainty was caused by a lack of clarity in the task descriptions.

Clearer instructions When observing participants complete the study 1 methodology we noticed numerous small confusion points that together made participants uncertain about what to do in the study. For example, a warm-up task tells participants that a photo of a poster has an incorrect title but doesn't say the correct title. Participants needed to read the title from the photo, but participants became confused. In study 2 we clarified that the titles can be read from the posters in the photos. Another example is from study 1's task 13 where Pat's sister apologizes for messing up Mom's photos and asks Pat to put the photos "back the way you had them." The participant is supposed to undo changes made by the sister so that the album looks like it did at the end of task 11. Some participants tried to change the album back to what it looked like when they first saw it at the beginning of task 11. We clarified the explanation. When running these tasks on practice participants we specifically asked them if these points were clear.

Outcome Participants appear to have taken responsibility for the albums and considered permissions to be in the bounds of the study. We did not observe any participant choosing to not change permissions due to concern about who owned an album. The clarification in wording resulted in less participant uncertainty over how to handle situations.

6.4.3 Study 3

Directly telling participants that they were responsible for the albums, combined with clear wording, appeared to have caused study 2 participants to sufficiently take responsibility for the albums. In study 3 we tried to keep these themes.

Prompts We initially decided to make only warm-up tasks 1 and 2 prompted tasks to make sure that participants were capable of performing all the actions necessary for the study. As part of the prompting emails, the participant is directly told that it is their responsibility to find and fix these types of errors.

After running the protocol on several practice participants we discovered that around the 5th task, participants would start to become lazy and stop taking responsibility for correcting all the errors. We solved the problem by making task 5 a prompted task. Similar to warm-up tasks 1 and 2, the participant was told in the email that fixing errors is their responsibility.

Outcome Participants took responsibility for the albums and considered permissions to be in the bounds of the study. When asked after the study if they felt they could change

permissions, all participants asserted that they felt they were allowed to do so.

Making task 5 a prompted task was very effective in reinforcing participant responsibility. Throughout the study participants would get lazy or careless around this task, receive a strongly worded email from their boss, and immediately start paying more attention. In the debriefing interview we asked participants about their reaction to this email. Participants said that they realized that the boss would be checking their work so they needed to do a good job.

6.4.4 Study 4

The methodology for study 3 worked well so we made only minor alterations for study 4. We reduced the strength of wording in the prompted warm-up task so that it simply pointed out the error. Because participants only had eight tasks per condition and were limited to 2 minutes we decided to not prompt mid way through.

Outcome Because study 4 is an online study we have limited feedback on participant's feeling of responsibility. Participants who gave study feedback expressed a strong desire to get all the tasks correct. The number of permissions and non-permission subtasks corrected also indicated that participants took responsibility for the albums.

6.5 Ideal Policy Comprehension

Participants should know the ideal policy associated with the content they are working with.

6.5.1 Study 1

We considered conducting the experiment using participants' own albums and policies but ultimately decided against it. Prior work has shown that participants' ideal policies change over time [72], in reaction to new technology [13], and based on context [27]. Mazurek et al. asked participants to provide ideal policies twice: all at once in a single sitting and by answering the same questions in small batches over the course of a week [72]. They found that the same participants responded with different ideal policies depending upon when asked. We were concerned that participating in our experiment would impact participants answers concerning their ideal policy, negatively impacting our ability to get an accurate ground truth. Instead we decided to create a fictional ideal policy which would be consistent across all participants.

To make the ideal policy appear less like explicit instructions, we expressed it using passive voice in the emails. However, not all permission information, particularly who shouldn't see the albums, could be easily expressed in passive voice so some information was presented in instruction pages that described the people the participant was about to interact with. To make this information simple to internalize, we created characters. For example: Pat's mother was described as panicking easily, while Pat was described as

enjoying dangerous activities. The instruction sheet commented that Pat generally avoided telling his/her mother about the dangerous activities.

We decided to have two permission warm-up tasks to be certain that participants could accurately both read permissions as well as change the permissions. If they were unable to do so the researcher provided guidance. The first permission warm-up task simply asked the participant if a particular album was visible to everybody on the internet or not. The second permission warm-up task asked the participant to change the permissions on a specific album.

Outcome Participants seemed to understand the ideal policy without difficulty and participants who made changes tended to make the correct ones. However, we have no way to determine why participants who did not change permissions chose not to do so.

The warm-up task in which participants were asked to read a permission resulted in participants guessing instead of reading the permission. In the warm-up task, Pat's boss asks if people at other companies can see a particular album. Participants tended to correctly guess that the album was publicly visible and answered the question without even looking at the screen. We had prepared prompting emails in the event of an inaccurate guess, but had not anticipated that the majority of participants would guess accurately. For the non-control conditions there was no way to be certain they had guessed since we could not verify if they had looked at the display.

6.5.2 Study 2

Participants seemed to understand the ideal policy in study 1 so we made minimal changes to the way it was presented.

Changed permission-read warm-up task In study 1 participants were guessing that anyone on the internet could view the album in the permission reading warm-up task. In study 2 we changed the task so that the correct answer was that anyone on the internet could *not* view the album thereby making it the opposite of the common guess.

Think-aloud protocol For reasons discussed in following sections we made study 2 a think-aloud study. A side effect of this decision was that participants had to read all instruction materials and emails out loud, ensuring that all materials, particularly the ideal policy, were read. We were also able to determine when instructions were confusing.

Outcome In warm-up task 2 (read permission) we observed more participants consulting the display to determine what the permissions were instead of opening the permission modification interface. Participants were still inclined to guess that the album was public but the guesses were now wrong and the researcher was able to prompt them, so every participant understood how to read permissions.

Using a think-aloud protocol forced participants to read all text aloud, thereby ensuring that all materials, including information about the ideal policy, was not skimmed over. Based on the think-aloud statements, participants appear to have understood the ideal policy. However, the protocol had no explicit outcome variable with which to test ideal policy comprehension.

6.5.3 Study 3

In this study we decided to present one ideal policy to the participant at the beginning instead of presenting the policy in pieces. This was done to provide consistent permission priming (Section 6.3.3). It was also done to promote participant understanding of the ideal policy and make it easier to test that understanding.

Testing ideal policy comprehension Participants in studies 1 and 2 appear to have understood the ideal policy, but we did not measure their comprehension. Study 3 had a single ideal policy so we were able to perform a pre and post test of participants' ideal policy comprehension. The pre-test was administered after the warm-up tasks, participants were asked by a co-worker if a provided photograph was appropriate for the website and if they should do anything when posting it. The post test is part of the final survey, participants were asked what the permissions for several albums should have been.

Outcome Ideal policy comprehension was provably high in this study. Participants had no problem remembering the ideal policy and were able to apply it to different situations and albums with high accuracy.

In the pre-test 78% of participants correctly mentioned permissions for both comprehension questions and only one participant never mentioned permissions. Participants behaved similarly on non-permission comprehension questions. This means that participants were able to 1) recognize that permissions might need to be set for these photos, and 2) correctly apply the ideal policy. Across conditions participants answered an average of 91% and a minimum of 67% of post-study permission comprehension questions correctly. This shows that the methodology design enabled participants to correctly understand, remember, and apply the ideal permission policy.

6.5.4 Study 4

As mentioned in Section 6.3.4 we were concerned that the explicit listing of ideal policy rules in a bulleted list was over priming participants to look for permission errors. With practice participants in study 4 we experimented with several information page designs. We conveyed the ideal policy in paragraph form with varying levels of wording intensity and compared that with providing the policy in bullet point form. We found that presenting the policy in bullet point form lead to the lowest level of variance and the largest difference in permission correction between conditions.

Outcome In study 3 participants could answer "I do not know" to any comprehension question, but it was rare that they did so. In study 4, 50% of participants answered "I do not know" to at least one comprehension question, but only 4% answered all comprehension questions that way. Of the answered questions 90% were answered correctly. Interestingly the design of the information page which conveyed the ideal policy had minimal effect on ideal policy awareness. Participants who saw the ideal policy in paragraph form correctly answered approximately 87% of comprehension questions, with minimal variance between designs.

6.6 Effective Outcome Measurement

We needed to differentiate between environments with no errors, environments where participants are not noticing errors, and environments where errors have been noticed.

6.6.1 Study 1

We chose a lab study design because it offered us the most amount of control over potential variables. We could control the task design, types of errors, and when errors would appear. By using a role-playing scenario we could also control participants mindsets when approaching problems.

In order to test our primary hypothesis H we needed to detect when a permission error was "noticed." We anticipated that a participant who noticed an error was very likely to correct it. So for this study we defined "noticed" as "corrected." The number of people correcting a permission error is a strict subset of the number of people noticing errors and we anticipated a large difference in the number of permissions corrected between the conditions. So we were willing to accept that we might not detect a participant that chose not to correct a noticed error.

When designing memory questions we were concerned about participant fatigue leading to questions being guessed at or answered with the fastest answer. To counter this we limited our questions to six albums and only asked about two of the actions. We also required that all memory questions be answered with True, False, or Not Sure. This was to make providing answers the same amount of work as guessing.

Outcome Unfortunately, we did not see a statistically significant difference in the number of permissions corrected between conditions. We also observed participants noticing errors and choosing to not correct them which was not captured by our definition of "noticed." We considered changing our definition but determining if a participant had checked the permissions was impossible for participants in the non-control conditions who might or might not have looked at a proximity display. So, while it may be the case that H is supported if we define "noticed" as "checked permissions," our lack of measurement fidelity prevented us from testing this.

6.6.2 Study 2

In designing the outcome variables for study 2 we focused on being able to notice when participants checked permissions as well as when they corrected permissions.

Think-aloud and eye tracker Our inability to accurately measure when permissions were noticed but not changed was a major issue with the study 1 methodology. To adjust, we made study 2 a think-aloud study. Study 1 was deliberately not a think-aloud study so we could determine if participants took an equal amount of time to complete tasks (H5). Think-aloud protocols are known for giving inaccurate timing information. In study 2 we felt that accurate timing information was less important than accurately measuring participants' interactions with the displays.

To assist in measuring if and when a participant focuses on a display we decided to use an eye tracker. This data was intended to augment, but not replace, the think-aloud data.

Outcome The think-aloud data enabled us to determine when participants *checked permissions* using the following definition. Control participants were judged to have *checked permissions* if they opened the permission management interface and the permission was visible on the screen. Participants in the other conditions were judged to have *checked permissions* if they (1) opened the permission management interface; or (2) read permission aloud; or (3) clearly 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. This definition allowed us to measure if a participant paid significant attention to a display.

Data from the eye tracker was less helpful than anticipated. To operate, the eye tracker needed participants' faces to remain in a small area. This is possible for short studies, but our study was 1.5 hours. Participants would shift in their chairs or lean on the desk moving them out of range. We considered prompting participants when they moved outside the required area but decided this would distract participants and alter their behavior. We tried having participants experiment with the eye tracker before the study so they knew where the optimal area was. This helped, but participants still became distracted by the study and started moving outside the optimal area. While incomplete, the eye tracker data did give us a sense of when participants looked at displays.

6.6.3 Study 3

In study 3 we wanted to get more detailed qualitative data about how and why participants checked permissions. Our definition of "permission checking" from study 2 appeared to be working well so we did not modify it.

Permission modification interface In studies 1 and 2 we observed no difference in memory between the conditions (H4). We hypothesized that this was due to the full sized permission modification interface. Participants who visited the interface frequently changed more than one permission indicating that, even in the control condition, they were looking at other permissions. To address this issue we added the permission modification interface as an independent variable. The permission modification interface was either a separate page with all permission settings shown or a dialog with only one album's permission settings shown. We added the following hypothesis:

H7 Participants who see a comprehensive policy modification interface remember permissions better than participants who see a policy modification interface that displays a single album.

Post-study memory In studies 1 and 2 we asked participants to answer 128 memory questions about 13 albums, 4 groups and 2 actions (view and add) and saw no statistically significant difference between conditions. In this study we wanted more qualitative data to better understand what people remembered. We decided to verbally administer the memory questions and elicit free responses. We felt free form answers would get us a better

sense of what the participant remembered. Once all the memory questions had been asked the researcher prompted the participant about anything they had not yet mentioned. For example some participants only answer the questions in terms of the view action so the researcher would ask if they recalled the add or edit action for any of the albums.

When we asked practice participants, who had not checked permissions during the study, memory questions, we found that they started feeling embarrassed that they didn't know the answer, and after a couple questions they started guessing. To discourage guessing we interleaved the memory and comprehension questions. This meant that every participant could, at worst, provide an answer for every other question without having to guess. We found that this discouraged guessing and participants seemed more comfortable admitting that they could not recall the permissions for albums they did not check the permissions on.

Post-study debriefing Once all the questions had been completed we conducted a debriefing interview with the participant. In the prior studies participants had occasionally behaved unexpectedly. Initially we thought this was caused by methodology issues but some behaviors persisted through different methodologies. In this study we wanted to get the participant's perspective on why they engaged in these behaviors. However, many of the behaviors were short (1-2 seconds long) and we were concerned that participants would not remember why they had made a comment an hour ago. So we used a contextual interview approach [44] where the participant opened the album they were working with and the researcher explained the context in which the behavior occurred and asked the participant questions concerning what they were thinking or why they had done something.

Outcome This study design allowed us to accurately measure and test all the outcome variables we were initially looking for. The only issue was an unknown confounding variable that caused some participants to check permissions frequently and other participants to check them rarely.

The use of a single ideal policy allowed us to observe natural participant behavior that was inhibited by the design of prior studies. In prior methodologies the participant was unable to choose when to check permissions because they did not know the ideal policy until they started a task. With one ideal policy we observed several participants deciding at a single point in the study to check permissions for every album at once. This behavior was facilitated by the full permission modification interface. We found that participants who saw the full interface performed better across several measurements and were more likely to correct permissions regardless of if they saw the proximity display or not.

The combined use of a single ideal policy, randomized task order, and randomized permission error order allowed us to notice issues with our definition of permission checking. In the control condition we reliably determine when the permissions were shown. In the non-control conditions, we only determine when permissions were checked based on participant behavior. In study 3 non-control participants were statistically more likely to check permissions when there was an error than when there was no error. There was no statistical difference for the control participants. This suggests that participants were able to glance at the display and determine if there was an error fast enough to not vocalize [104]. This is good news for our display but it implies that we can only detect when a participant *focuses on checking permissions* rather than being able to detect every time they check

permissions. The eye tracker allowed us to determine when they fixate on a display but similarly did not tell us when they actually checked the permissions.

The use of contextual immersion during the debriefing session was very effective at getting participants to remember their reasoning behind specific actions. In cases where the participant couldn't remember they were still often able to make an educated guess as to why they would have done an action given their behavior up to that point. While a guess is not as good as remembering, participant's guesses as to reasons behind their actions were more accurate than researchers educated guesses.

6.6.4 Study 4

The prior studies had a small number of participants, and they exhibited a large between-participant variance, making it difficult to detect differences between conditions. In this study we wanted to increase the number of participants and account for the variance.

Within subjects In study 3 we observed that some participants internalized the need to check permissions while others did not. In the debriefing interview the participants who internalized considered it "obvious" and those that did not check permissions appear to have read the ideal policy and then forgot about permissions. To control for the pre-disposition to pay attention to permissions we decided to make study 4 a within-subjects study where every participant performs the training and tasks on both the control condition and one of the non-control conditions.

Measuring "noticing" Our hypothesis H is that participants in some conditions can "notice" permission errors more frequently than participants in other conditions. In studies 2 and 3 we equated noticing permission errors with checking permissions. However, measuring permission checking requires observation of the participant not possible in an online study. Additionally, we showed in study 3 that our measurement of permission checking was, at best, a lower bound for the number of times permissions were actually checked by participants. In study 4 we returned to our definition of "notice" from study 1 where we equate correcting permissions with checking them. This definition provides only a lower bound but with the larger number of participants and improvements to the methodology we did not anticipate a problem.

Permission modification interface In study 3 we observed that participants who saw the permission modification interface in a dialog had a larger difference in performance between conditions than participants who used the full page permission modification interface. Since our main hypothesis H is concerned with the impact of proximity displays, not permission modification interfaces, we decided to use the dialog for study 4.

Outcome Using the stricter definition of "notice" as "corrected" was effective in that we were able to show statistically significant differences between some of the conditions and control (not all conditions were expected to have a difference). We attribute this to both a larger number of participants and clearer, more tested, study materials.

Similar to study 1 we had a limited ability to measure why participants did or did not make changes to permissions. However, we collected extensive logs which we were able to compare to behaviors observed in prior studies allowing us to imply what users were doing.

6.7 Discussion

We discussed the methodologies of four studies designed to test our hypothesis. When designing our initial study we tried to account for anticipated methodology issues. Our initial design succeeded in some aspects and was lacking in others. Subsequent studies were adjusted to account for observed issues.

Secondary permission task Users treat security as a secondary task because the benefits of security are hard to envision but the costs of engaging in it are immediate [109]. In our studies we did not want to incentivize the participant to check permissions so we tried to balance the amount of priming with the cost of checking. We successfully managed priming on study 2 and 4, but in studies 1 and 3 we over-primed, first by mentioning permissions too frequently and then by using strong wording to express the ideal policy without forcing participants to consider trade-offs. We increased the immediate cost of checking permissions in studies 2 and 3 by adding tasks where the permissions were already correct and checking them cost time and effort. We further increased the cost in study 4 by adding a time limitation which forced the participant to make trade-offs. We found that at least 50% of the tasks needed to have no permission error in order to give checking a high cost compared to the benefit.

Participant responsibility Role playing was very effective in making participants feel responsible for albums that belonged to Pat. Our main issue was when we asked participants to be responsible for albums that belonged to people such as Pat's mother. We countered this issue in the second study by making it clearer that others trusted Pat to make changes.

Ideal policy comprehension We tried two methods of expressing the ideal policy to participants. The first was to have a different policy for each album. The policy was expressed using passive voice in the emails (studies 1 and 2). The second way was to have a policy that applied to all the albums. The policy was expressed using direct wording at the beginning of the study (study 3 and 4). Both methods sufficiently communicated the policy to the participant. The per-album policy gave participants less priming towards fixing permissions but was difficult to make consistent. The study-wide policy over-primed some participants to look for permission errors, but provided consistent priming to all participants on all tasks.

Effective outcome measurement Our primary issue with measuring the study outcome was defining and testing participants' ability to "notice" permission errors. In the first study we defined "notice" as changing permissions, but this definition was insufficiently precise to measure the difference between conditions. In later studies we changed our definition of "notice" to checking the permissions for errors. This definition allowed us to observe if participants were looking for errors independently of whether they found the error or decided to fix it.

In conclusion we presented the methodologies of four studies and discussed the decisions and outcomes of each study. We were able to describe our methodological successes and difficulties in terms of our four goals: 1) secondary permission task, 2) participant responsibility, 3) ideal policy comprehension, and 4) effective outcome measurement. Through this process, we have shed light on the challenges intrinsic to many studies that examine

security as a secondary task.

Chapter 7

Detailed methodologies

This chapter expresses, in detail, the methodologies and data analyses used in the last three evaluation studies. The first evaluation study was inconclusive due to methodological issues and is referred to solely as motivation for later methodological choices. Details of the first study and the reasoning behind subsequent methodology changes can be found in Chapter 6.

The second lab study, which we refer to as the *eye tracker study*, was a between-subjects lab study in which an eye tracker was used to better understand when participants were looking at the proximity displays. Tasks were given to the participant in a fixed order. This study showed that participants who see proximity displays shown under the photo and album thumbnails correct statistically more permission errors than participants who do not [106]. During this study we noticed several unanticipated user behaviors which we then explicitly tested for in the following study. Excluding results from the eye tracker, all results from this study are replicated with greater precision and power in later studies. Consequently we focus on the analysis of the eye tracker data when discussing this study.

The third lab study, which we will refer to as the *lab study* as it is the only one we discuss in detail, was a between-subjects lab study. Tasks in this study were presented in a random order and were randomly paired with permission errors to ensure that permission checking behavior was being measured separately from the influences of task wording or ordering. This study focused on collecting detailed observations of participants and post-study interviews.

The fourth evaluation study, which we will refer to as the *online study*, was conducted on Amazon's Mechanical Turk. This study was a between-subjects study where each participant saw both a control condition and an experimental condition. This study used a fixed task order and participants had a set time to work on each task. This study focused on evaluating the effectiveness of proximity displays with a large number of participants.

In the following sections we detail the methodologies and data analysis for all three studies.

7.1 Eye tracker study

The eye tracker study was 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.

7.1.1 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 full page policy modification interface (Figure 7.2(a)). 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 7.1): 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 verbally 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 7.1. Each task was composed of a set of *subtasks*, individual permission, rotation, deletion, spelling, or re-naming 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,

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 7.1: Tasks and information given to eye tracker study participants.

friends, or family (shown in the second column of Table 7.1). 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 7.1). 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 7.1). *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 7.1). 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.

7.1.2 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.

7.1.3 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.

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. This information was then matched with the fixation data from the eye tracker to determine when participants saw proximity displays.

7.2 Lab study

The lab study was a 1.5 hour between-subjects study where 33 participants were divided into four study groups based on two treatment types: proximity display and permission modification design. Half of the participants saw permission information on a proximity display located under every photo and album thumbnail, the other half saw no proximity display. Similarly, half of the participants modified permissions using a full page permission modification interface (Figure 7.2(a)), and the other half modified permissions using a popup dialog (Figure 7.2(b)). 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 17-person pilot.

7.2.1 Study conditions

This study had two experimental variables: proximity display, and permission modification interface. Both variables had two levels, resulting in four experimental conditions:

Condition name	Proximity display	Permission modification interface
Control Dialog	None	Dialog
Control Full	None	Full page
Under Dialog	Under photo/album thumbnail	Dialog
Under Full	Under photo/album thumbnail	Full page

Proximity display – Participants in the control condition see no permission information on the photo management interface (Figure 7.1(a)). To access the permission modification interface, control participants must select "edit permissions" from one of the options menus. Participants in the under photo condition had the option of placing their mouse over an album or photo thumbnail to see the proximity display (Figure 7.1(b)), or using the "edit permissions" link in one of the options menus.

Permission modification interface – When a participant clicks on any of the "edit permissions" links or uses the "manage permissions" link on a proximity display, they are taken to a permission modificaiton interface. Participants in the dialog condition see a permission modification dialog that allows them to view/modify permissions for this album only (Figure 7.2(b)). Permission information for other albums is not shown on the dialog. Participants in the full permission modification interface condition are taken to a new page where they can view/modify permissions for any album (Figure 7.2(b)). To assist users, the album they were previously viewing is highlighted.

7.2.2 Protocol

The lab study was a between-subjects design with a round robin assignment to experimental conditions. A think-aloud protocol was used. Participants in all conditions performed the same tasks and saw the same permission errors, however, tasks and errors were shown in a random order.

Participants were first asked to read and fill out a consent form. Next particpants were given the opportunity to interact with the eye tracker. The pre-study indicated that participants who understood where the eye tracker could and could not see them were more likely to stay in range during the study. After interacting with the eye tracker participants were trained in how to think-aloud followed by a short calibration of the eye tracker.

Participants were then presented with a training version of the Gallery site. The training version is identical to the website used in the primary section of the study with the exception that it has a different and smaller set of albums and photographs. Participants were verbally given a user name and password and asked to log in. They were then given a printed tutorial and asked to work through it, making changes to the website as they went. The tutorial clearly stated that this was a practice version of the website and made it clear that experimenting would not impact the main study. The tutorial covers how to navigate Gallery, move photographs between albums, change titles on photographs and change

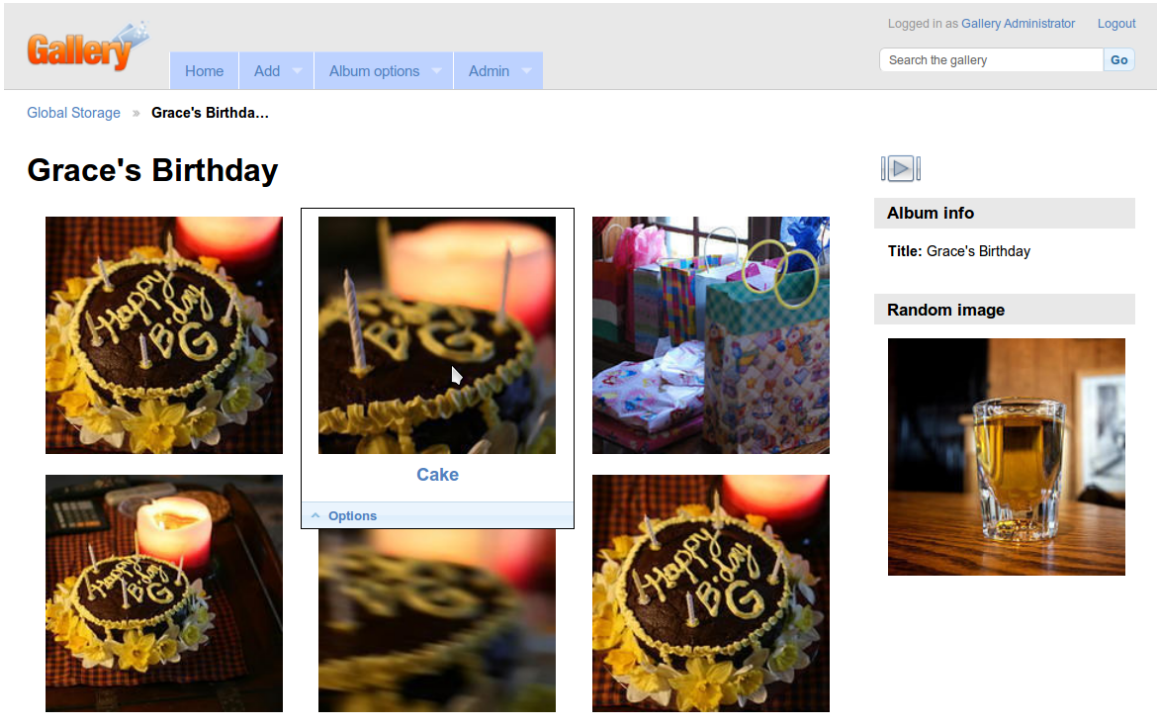

(a) Control

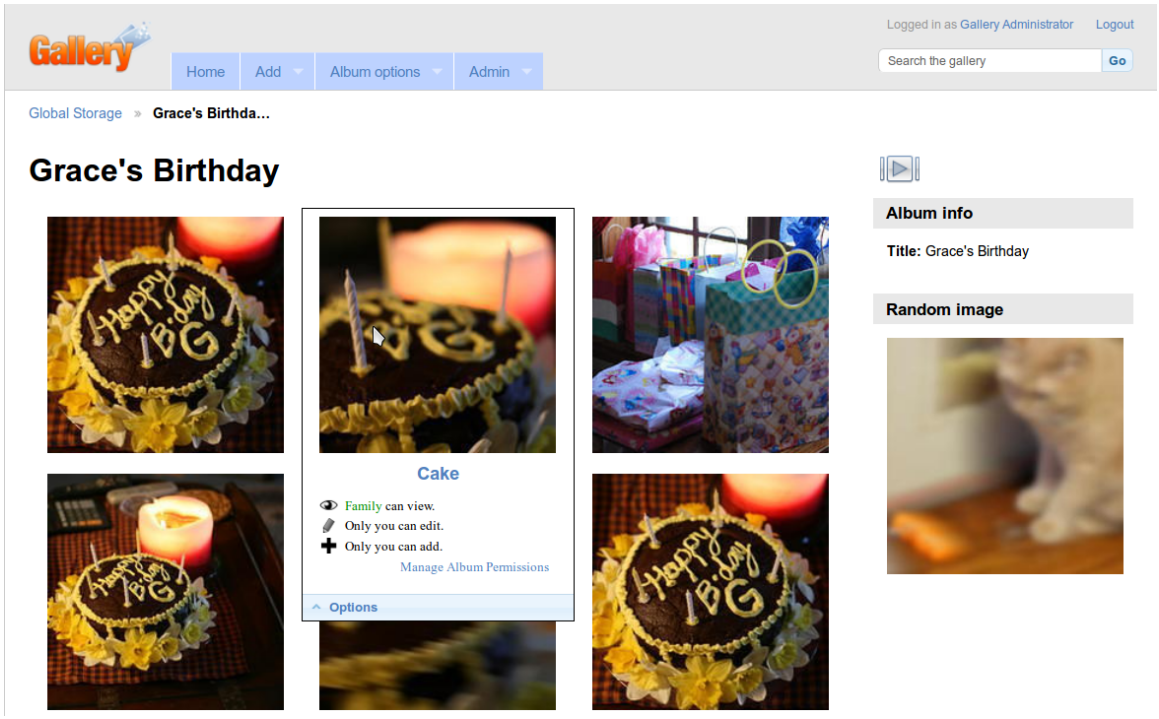

(b) Under photo

Figure 7.1: Gallery interface without a proximity display (a), and with a proximity display under every photo and album (b).

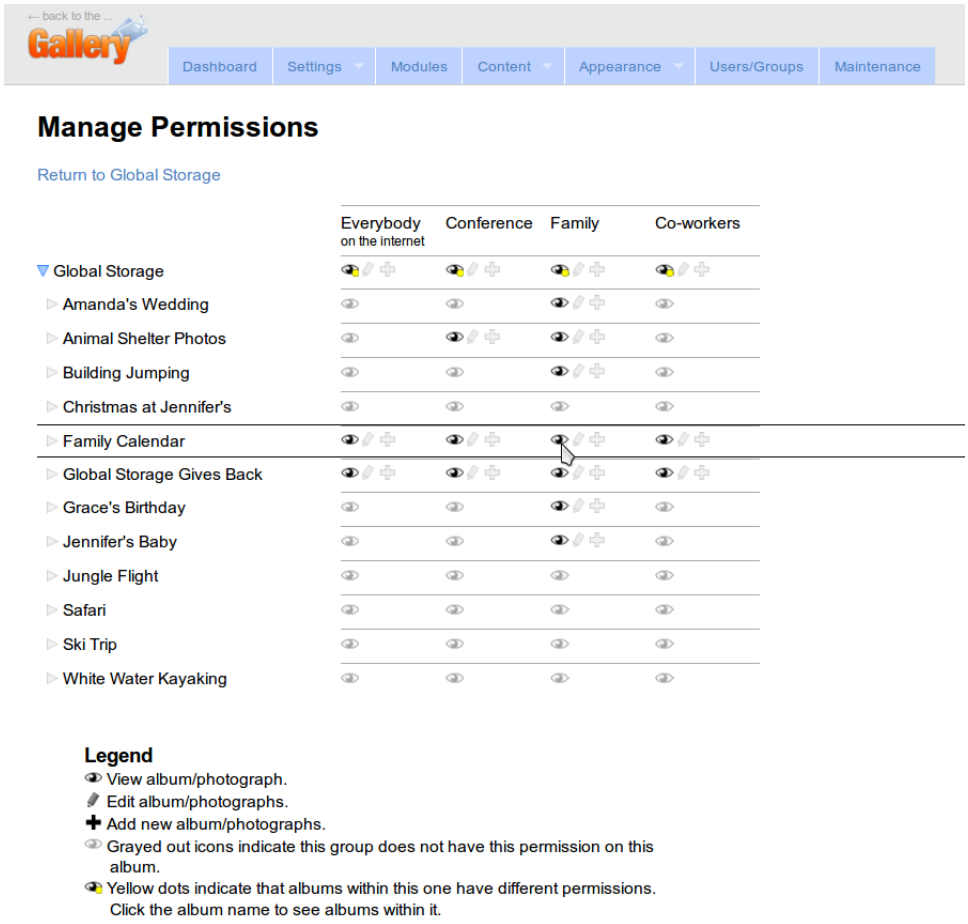

(a) Full page permission modification interface

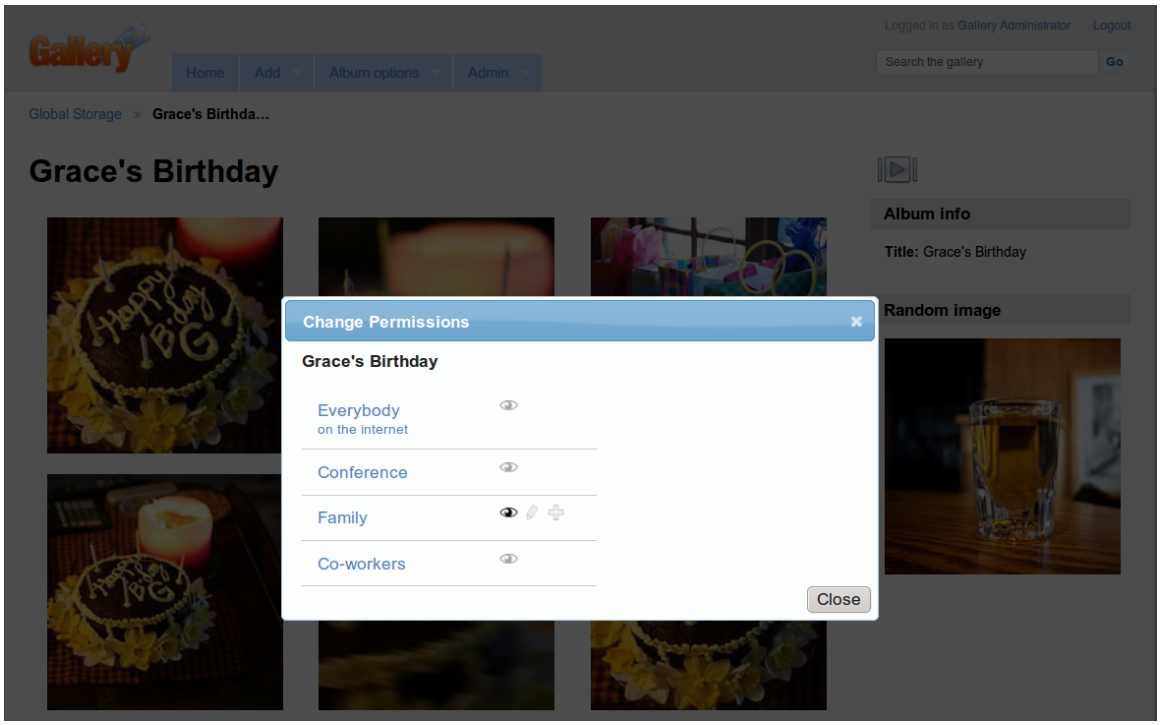

(b) Dialog permission modification interface

Figure 7.2: Full screen permission modification interface (a) and dialog permission modification interface (b).

1. Tutorial
2. Instructions about Pat and ideal policy
3. Training
   (a) Task 1 (Rotate, delete, cover, spelling, change permissions)
   (b) Task 2 (Move, read permissions)
4. Prompts – feedback on both training tasks
5. Three policy comprehension questions
6. Three tasks randomly drawn and removed from from the set of tasks with permission error drawn and removed from the set of permission errors and non-errors such that 50% of tasks have errors
7. Task 4 with permission error randomly drawn and removed from the set of permission errors
8. Task 4 task prompt
9. Eleven tasks randomly drawn and removed from the set of permission errors and non-errors such that 50% of tasks have errors
10. Memory and comprehension questions
11. Debriefing interview

Figure 7.3: Lab study protocol order.

Gerald's Photograph Policy

1. No photographs containing drugs, alcohol, or anything inappropriate.
2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.
3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.
4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.
5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

Figure 7.4: Gerald's Photograph Policy

permissions. Based on the pre-study we decided not to train participants how to rotate a photograph or change the cover image on an album as both of these features are easily found on the same menu as title manipulation which the participant is already trained to find. Participants took an average of 5 minutes 27 seconds to complete the tutorial.

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 emails from coworkers. These emails, delivered to participants in printed-out form by the researcher administering the study, were requestions from Pat's co-workers to perform various tasks with the online albums. The participant was allowed to look back through any piece of paper given to them including the tutorial and the instruction sheets. This was done to mitigate confounding issues with participants' memory.

The participant next opened the main Gallery website, and was given two instructional papers. The first paper described Pat Jones, an employee of Global Storage responsible for maining the online photo album. The second instructional paper explained that part of Pat's job involved helping other co-workers with the photograph system. However, it was also Pat's job to enforce the boss, Gerald's, photograph policy (Figure 7.4).

The bulk of the study was comprised of 2 warm-up tasks and 14 normal tasks. Tasks in this study are composed of a two paragraph email and an associated album. Each email started with a short paragraph from the sender describing the album, this paragraph clarified if the album was personal or professional. The second paragraph named the album the participant was to work with and stated the set of *explicit subtasks* the sender would like Pat to complete. The album also contained at least one photograph which conflicted with Gerald's rules, we refer to these conflicts as *implicit subtasks*. Each album contained either personal or professional photos and 50% of the tasks were associted with personal albums. Task and permission error orders were randomized so as to remove effect caused by task wording or ordering. The only exception was the fourth non-warmup task which always occured in the same location and always had a randomly selected permission error. In this way 50% of the personal tasks have permission errors and 50% of the professional tasks have permission errors, but the errors occur in a random order. Tasks begin when the participant is given the associated printed email and end when the participant requests the next email.

The participant was initially given two warm-up tasks, one-at-a-time, which matched the description of tasks in the prior paragraph in every way except that they had fixed permission errors that are never randomized. The two training tasks explicitly instructed the participant to perform every action needed in the study.

- Rotate a photograph
- Delete a photograph
- Change an album cover
- Move a photograph
- Change a title
- Change a permission

Permissions were never expressly mentioned in any of the emails, so for consistency they were not mentioned in the warm-up emails. Additionally, blurry photographs and photographs containing alcohol were never expressly mentioned in any email, instead of mentioning them in the warm-up we asked participants to delete a photo for another reason. If a participant had technical trouble completing an action during the warm-up then the researcher provided additional instruction. If a participant did not complete a subtask, then, after both training tasks, they were sent a prompting email from their boss. This email clearly stated that he noticed that the participant had not performed the action correctly and asked the participant to fix the issue. This prompt clearly stated what subtask the participant had failed to complete.

Participants next received three emails, one at a time, from a co-worker with three different photographs attached (Figure 7.5). The co-worker asked if the attached photograph was acceptable to post on Gallery and if so was there anything the co-worker needed to make sure and do. These emails were used to be certain that participants understood and could apply Gerald's policy. If the participant did not mention permissions when responding to the emails they were sent an email from the co-worker asking if they needed to do anything with the permissions. If the participant answered the question incorrectly they

Q1)  Q2)  Q3) 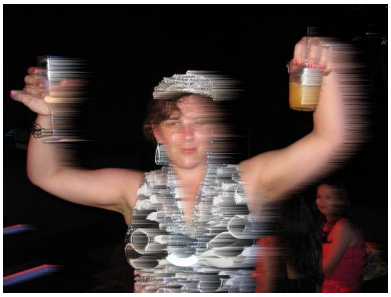

Figure 7.5: Participant was asked, by a co-worker, if each of the above images were acceptable to post on Gallery and if the co-worker needed to make sure to do anything when they put the photograph on Gallery. Q1 has no problems with the photograph but should be visible to Friends and Co-workers only. Q2 needs to be rotated and should be visible to Everybody on the internet. Q3 cannot be uploaded to Gallery because it is blurry and contains alcohol.

were sent an email containing Gerald's rules from their co-worker.

Participants were next given three tasks, as described above. These tasks were randomly ordered and *unprompted*, if a participant failed to complete any of the explicit or implicit subtasks no action was taken and they were allowed to move onto the next task with no prompting or reminders.

The fourth task was not randomized and every participant saw the same task at the same point in the study. The fourth task contained two rotation errors, two alcohol errors, and a permission error, which was randomly selected from the set of all personal permission errors. If a participant failed to complete any of the explicit or implicit subtasks they received a prompting email from Gerald, their boss. The email said that Gerald checked the album and was disappointed in Pat, the prompting email also contained a list of Gerald's rules, but did not specify what Gerald found to be incorrect. Based on the pre-study we found that the fourth task was the point where participants started realizing that no one was checking their work. It is also the point where they were experienced with the interface but may have forgotten about the rules. By providing a prompt here we ensured that participants remembered to enforce the rules and realized someone was going to be checking their work. If the participant failed to correct all the errors after the prompt no action was taken and they were allowed to move on to the next task.

The remaining ten tasks were presented to the user in random order. If the participant failed to fix any errors in these tasks no action was taken and the participant was allowed to move onto the next task without any intervention.

After completing all tasks the researcher asked the participant a set of recall questions for the last four albums the participant saw and the third and fourth albums the participant saw. The participant was asked to recall what the permissions were on those albums the last time the participant interacted with them, as well as what permissions the album should have had according to Gerald's policy. The participant was then asked to recall, in

their own words, what Gerald's policy was.

Finally, the researcher engaged the participant in an unstructured interview to better understand the decisions and behaviors the participant had engaged in. Many of the behaviors were short (1-2 seconds long), so we used a contextual interview approach [44] to help people remember what they were doing when they engaged in the behavior. For each question we had participants open the album they had been working with and the researcher explained the context in which the behavior occurred and asked the participant questions concerning what they were thinking or what they had been trying to do. The researcher also asked about prior experiences and opinions that might have impacted participant behavior.

7.2.3 Participants

Participants were all native English speakers who had previously uploaded photographs to an online social network or photograph sharing website. They were recruited using an existing pool of people interested in participating in behavioral research studies. This pool includes both students and members of the Pittsburgh community. They ranged in age from 18 to 53, with a median age of 22. Participants were predominately students (60%), currently enrolled in college (39.9%), and female (66%). All had previously shared photographs online using Facebook and other photo sharing sites. No lab study participant dropped out of the study, but one participant was excluded for not completing all the tasks in the time alloted.

7.2.4 Data collection and analysis

Participants were audio recorded and a screen capture program recorded a video of their web browser. Our custom version of Gallery recorded detailed logs of the participants' actions in a database. The researcher took detailed notes during the session including a timestamp whenever she handed the participant a paper. Data collected from participants falls into seven categories: permission correcting, permission checking behavior, action order, non-permission correcting, permission recall, rule comprehension, and interview data.

Permission correcting

In the Gallery system permissions are associated only with groups and albums. Neither photographs nor users can have permissions associated with them. Permissions are described as a triple of user, album, and action. Each user, album pair has three actions associated with it, view, add, and edit. In this study we tested seven different permission errors, which are shown in Table 7.2.

At the end of every study session, we archived the state of the final permission settings and automatically extracted them into a database. Each permission was compared to its initial state and marked as "changed" or "unchanged." The permission was also compared to its correct state and marked as "correct," or "wrong."

Professional/Personal	View	Add	Edit
Professional	Nobody can view	-	-
Professional	Family and Co-workers	Co-workers	Co-workers
Personal	Family, Co-workers, and Conference	Co-workers	Co-workers
Personal	Family only	Nobody	Nobody
Professional	Family and Co-workers	Co-workers and Family	Co-workers
Personal	Everybody on the internet	Co-workers	Co-workers
Professional	Family and Conference	Nobody	Nobody

Table 7.2: Possible default errors. Each participant experienced every error once during the primary 14 tasks. The first training task had a permissions error where everybody on the internet could see a personal album, so participants would have seen this error twice during the study session.

Permission checking behavior

The researcher recorded in their notes every time the participant *explicitly checked* a permission. To approximate when participants actionally notice permission errors we measured *explicit checking* behavior. Control participants are said to have *explicitly checked* permissions if they open the permission modification interface. Participants who were shown proximity displays are said to have *explicitly checked* permissions 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 (moving the mouse under the words while reading or circling the display with the mouse); or (4) pointed at the permission display with their hand while staring intently at the screen. The identification of explicit permission checking behavior in the under photo condition was done by the researcher during the study and recorded in the researcher's notes which were later coded in the database.

The majority of participants explicitly checked permissions during the task involving the album associated with that permission. However, some participants checked all permissions at once either by looking at permissions one-by-one on the full permission modification interface or by mousing over each album thumbnail one-at-a-time, reading the proximity display and correcting the permission when wrong. These participants were flagged in the database as having *checked permissions all-at-once*. Unless noted otherwise, the analysis's in the results section treat participants who are flagged as having checked permissions all-at-once as having checked permissions on every task.

The researcher's coding for checking permissions was compared with the action logs to determine approximate accuracy. For the control participants the action logs and the notes were perfect matches, control participants are forced to open a permission modification interface to check permissions and therefor all permission checking behavior is recorded in the action logs. Under photo participant action logs were also compared to researcher notes, the action logs were a subset of the researcher notes indicating that the researcher had not missed any measurable permission checking behavior.

Action order

The Gallery software was modified to record detailed logs of all participant actions. These logs include the type of action performed, what photograph/album it was performed on, any change in value, and a timestamp. After the study, the action log was modified to include what task the participant was engaged in when they performed each action. The task information was determined based on the timestamps the researcher manually recorded at the beginning of each task.

For each user and task, we analyzed the logs and created a list of the minimum and maximum timestamps for each action type (change cover, delete, move, permissions, rename, and rotate), an action type that did not occur for that user, task combination was excluded. We then ordered the actions based on the minimum timestamps and coded each action as *first*, *middle*, *last*, or *only*. We refer to this ordering as when the action was first engaged in. Similarly we ordered the actions based on the maximum timestamps and coded each action as *first*, *middle*, *last*, or *only*. We refer to this ordering as when the action was last engaged in. If a participant only engaged in one action on that task the action is marked *only* as there can be no order with only one action.

For example, suppose that a user performed three action events during a task: rotate, delete, and rotate. Then, the first engaged in ordering is rotate and delete, where rotate is first and delete is last. Conversely the last engaged in ordering is delete then rotate.

We chose the *first*, *middle*, *last*, and *only* codes based on the observation that in 70% of the tasks the participant engaged in three or less types of actions with an average of 2.8 different action types. Participants were free to take as long as necessary to complete tasks leading to a high variation amongst participants. We were chose to code the action order rather than normalize the timestamps because we felt it provided a more accurate picture of the order participants engaged in actions.

Non-permission error correction

Gallery stores the meta-data of all photographs and albums in a MySQL database, after each study session several scripts were run against this database to collect and code relevant information and then archive the database. The scripts compared the meta-data, on all albums and photographs to the default meta-data. If the default value was different from the final value then the album was marked as as having been rotated, re-titled, deleted, cover changed, or permissions changed depending the meta-data. A second script was run which coded each meta-data element as *error*, *error fixed*, or *no default error*. For example, some tasks required the participant to make any change to the title of a photograph, for these subtasks any change to the title was considered correct. Some subtasks required the participant to make a specific change, in these cases the applicable meta-data value had to match the required final state.

Permission recall

At the end of the study participants were verbally asked by the researcher to recall the final permissions of the last four albums they saw (tasks 10-14) and the albums from tasks

3 and 4. The last four albums were chosen because they were the most recent and therefore the most likely to be remembered. The album from task 3 was chosen because it was seen less recently, and before the prompt on task 4. The task 4 album was chosen because participants were prompted to change permissions on that tasks and were more likely to have seen and interacted with them. The participant's answers were coded such that they could be directly compared with the actual permissions in the database.

Rule comprehension

After completing the training tasks participants were asked, via an email from a co-worker, to determine if each of three photographs matched Gerald's policy and if not what changes needed to be made. Participant's answers to these questions were coded in terms of 1) if the photograph was suitable to be placed on Gallery at all, and 2) which, if any, of the existing issues the participant mentioned.

At the end of the study, in addition to being asked to recall permissions, participants were asked to apply Gerald's policy about permissions to each of six albums. Participants were asked what permissions Gerald would have wanted each album to have. These answers were coded and compared to Gerald's policy

Unstructured post-interview analysis

The post interviews with participants were transcribed, question and answer pairs were broken into topics. Each topic was printed on a slip of paper which was used in an affinity diagram [20].

7.3 Online study

The online study was an hour long within-subject study with 658 participants and 5 treatments. Each treatment was comprised of a control condition and an experimental condition, which were shown to the participant in serial. In the study, participants took part in two different role-playing scenarios in which they performed a variety of tasks, including various permissions-management tasks on a set of albums.

The goal of the online study (study 4) was to test the proximity display interface on more participants than could feasibly be brought into the lab. By using Mechanical Turk we were able to get a large number of users in a relatively short time frame.

This study was designed to test three hypothesizes:

H1: Correcting/checking permissions Users who see permission information on a proximity display notice errors more often than users who see permission information on a secondary page.

H2: Permission awareness Participants who see permission information on proximity displays can recall those permissions better than participants who see permission information only if they click to a second page.

H3: Negative effects Participants who see proximity displays take no more time, and correct no less non-permission errors than participants who see permission information on a secondary page.

7.3.1 Study conditions

Participants in this study were assigned round robin to one of five treatments, two condition orders, and two scenario orders, effectively assigning them to one of 20 possible treatment combinations. Each treatment was composed of two conditions: a control condition showing no access-control information, and an experimental condition displaying a version of the proximity display.

To prevent biasing either condition we ensured, via a round robin assignment to all 20 treatment combinations, that participants were equally likely to encounter the control condition first as the experimental condition. We also ensured that the home scenario was equally likely to appear first as the work scenario, and that they were equally likely to be paired with the different conditions.

We were concerned that the appearance/disappearance of the proximity display when participants switched conditions would draw unnecessary attention to the display, and bias our results. To counter this issue we created a similar proximity display which showed tag information instead of permission information. This display was placed on the control interface in the same location as permission information is shown in the experimental interface (Figure 7.6). For the audit, mixed, sidebar, and under conditions the tag display simply appeared in the same location as the permission display. For the Facebook condition a tag icon ( ) was displayed in the same place where the permission information icon was shown (Figure 7.6, Tables 7.3 and 7.4).

There were five experimental conditions in this study:

Audit – The audit condition showed, in the proximity display, who had recently accessed the album and what groups they were in. This display was visible under the album thumbnail, and when the album was opened it was displayed on the sidebar. If a user group had the ability to view an album then we setup the audit display to show at least one person in that group accessing the album recently. Similar with user groups who did not have access. For example the Animal Shelter album is only visible to Friends when it should be visible to Animal Shelter Employees. The audit display would show that several members of group Friends had viewed the Animal Shelter album but that no members of group Animal Shelter Employees had viewed the album.

Facebook – The Facebook condition was intended to simulate Facebook's access-control permission indicators as closely as possible. We decided to use Facebook's user interface design because it is both a very popular site for sharing photos and its user interface design is very different from our own. Facebook uses a set of icons to express the privacy policy associated with albums. An album can be publicly visible ( ), visible only to the owner ( ), visible only to friends ( ), or a custom settings ( ). Similar to Facebook's user interface we placed the relevant icons under each album thumbnail, and when the album was opened we placed the icons in the upper right hand corner. Mous-

ing over the icon resulted in a pop-up listing the groups who had the right to view the album. However, clicking on the icon resulted in our permission modification dialog rather than Facebook's drop down menu. Since we are testing people noticing errors rather than the impact of the permission modification interface design we felt it was more important that the permission modification interface be consistent across conditions than for it to be consistent with Facebook.

Mixed – The mixed condition showed the proximity display under the album thumbnail, and when the album was opened the proximity display was shown on the sidebar. This condition was selected based on the outcome of the two prior studies showing that participants check permissions at the beginning and ending of tasks, and that participants use the display under the album thumbnail to determine the presence of an error. This condition was expected to both support this behavior and take up less screen real estate than the Under Photo condition.

Sidebar – The sidebar condition shows the proximity display on the sidebar. No permission information is ever shown under the album photos. This condition setup is identical to the ones used in studies 1 and 2.

Under Photo – The under photo condition shows the proximity display under the album thumbnail, and when the album is opened it shows the display under every photo. This condition setup is identical to all three prior studies.

7.3.2 Participants

Participants in the online study were recruited using Amazon's Mechanical Turk. Participants coming from Turk were first shown a bulleted list describing what the study entailed, then they were shown the consent form, and if they agreed to it they were asigned a study group and began the study. We did not collect data from users prior to the consent form Table 7.5 shows the number of participants who completed the study in each condition (column 2), changed at least one permission without being told to do so (column 3), and the number of users who agreed to the consent form but did not complete the study (column 4).

Participants came from a wide range of professions and education levels. The most common profession was Student (26.3% of participants), followed by Unemployed (14.1% of participants). Only 5% of participants reported a technical profession. The most common education level was Some College (39.9% of participants), followed by Bachelors Degree (29.0% of participants). Participants ranged from 18 to 63 years of age with an average of 28 years old. 46.9% of the participants were male.

7.3.3 Protocol

This study was a within-subject online study conducted on Mechanical Turk. Participants were asked to read instructions, do a training, and complete eight tasks, for each of two conditions. After experiencing both conditions participants were asked to fill out a survey which asked memory questions about both conditions as well as demographics.

When participants visited the study from Mechanical Turk's website they were shown a page warning them that the study would take a full hour and they must complete at least 25% of the task components to be paid. If they chose to continue they were shown a consent form explaining that this was a photo management study.

The study web interface was divided into two frames, as pictured in Figure 7.7. The top frame showed instructions and emails to participants. The bottom of this frame contained a control bar that allowed participants to shrink the frame, obtain instruction on Gallery's features, and move to the next task. The lower interface showed the Gallery website the participant was currently working with.

Each participant went through the same set of tasks twice, once with each condition. So that the two sets of tasks did not appear identical we created two Gallery websites, one focused on Pat's personal life, and the other focused on Pat's professional life. We refer to whether the site concerned personal or professional photos as the site's *scenario*. The sites had different themes, titles, and photo content, but were otherwise identical. Care was taken that the albums in both sites had the same number of photos and the same type and number of permission and non-permission errors. The order in which participants encountered the two sites was assigned round robin.

Participants completed six training tasks on each site: open an album, rotate a photo, change a title, change a permission, change a tag, and move a photo. These tasks were expressed as stated directions. For example: "delete the blurry blue teapot." If a participant had trouble completing a training task they could select one of the instructional pages from the "show me how to" drop down. These instructional pages were viewable at any point in the study, but were most used during training.

After training participants were shown an instructional page telling them that their personal relations or boss expected Pat to assist in keeping the online photo albums to a certain standard, expressed by the set of rules shown in Figure 7.8. These rules differed only as necessary to accommodate the scenario. Similar to the lab study, this rule set included both permission and non-permission rules. There was one rule about tags, which states which tags need to be present. There were two rules about permissions which state who should and should not have the ability to view albums. The second permission rule is never associated with an error and primarily exists to support the logic of the scenario. Pat should always have access to the albums, we didn't want participants thinking they might remove their own access if they removed the group family.

Participants then went through eight tasks one by one. Tasks were introduced by a two paragraph email shown in the top frame of the browser window. Each email starts with a short paragraph from the sender talking about the album to be worked with, this paragraph makes it clear which group the album is associated with, the second paragraph names the album the participant is to work with and contains the set of *explicit subtasks* the sender would like Pat to complete. The album also contains at least one photograph which conflicts with the rules, we refer to these issues with photographs as *implicit subtasks*.

The first task was *prompted* in that if the participant missed any error, permission or otherwise, they were shown an promoting email pointing out the error and requesting that it be corrected. The remaining seven tasks were *unprompted*, no notification was given to the participant if they failed to correct an error. However, we were concerned that some

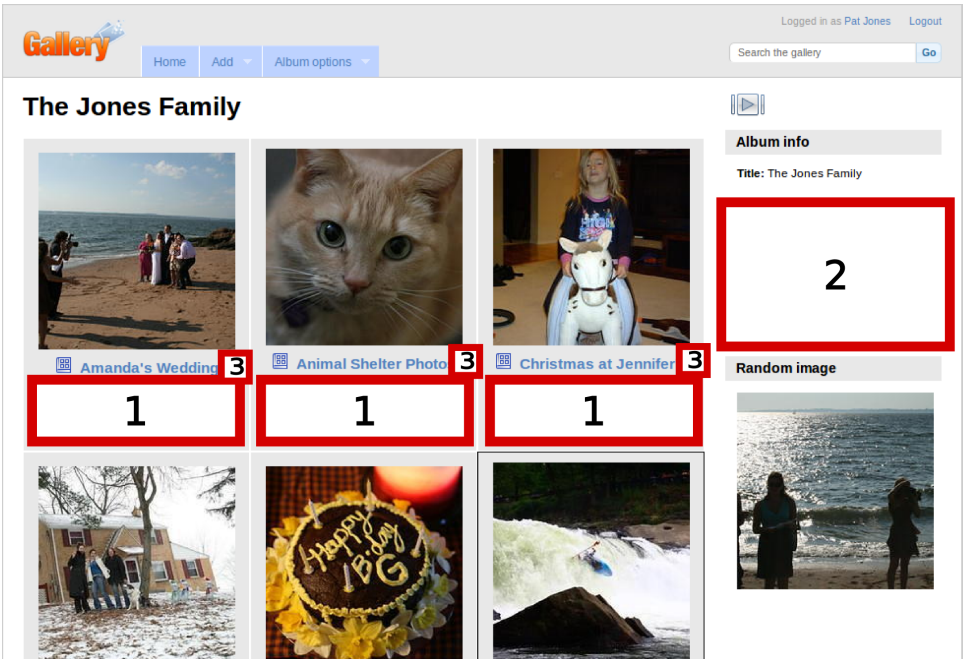

(a) Gallery albums page. Displays title and thumbnails for all albums.

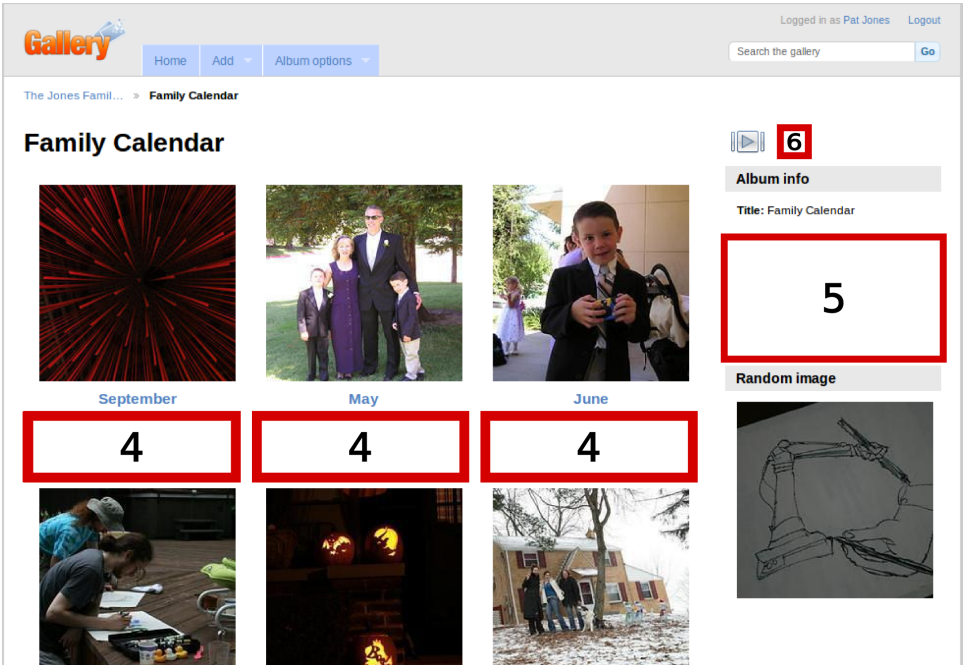

(b) Gallery photos page. Displays title and thumbnails for all photographs located inside a single album.

Figure 7.6: Gallery interface showing all the albums and their cover thumbnails (a), and the interface showing all the photos contained within a single album (b). Proximity display locations are marked with numbers 1-6 indicating the different locations where proximity displays were tested.

Condition	Figure 7.6(a) (all albums) Position	Figure 7.6(a) (all albums) Display	Figure 7.6(b) (opened album) Position	Figure 7.6(b) (opened album) Display
Audit	2 (sidebar)	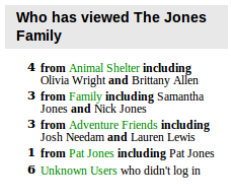	5 (sidebar)	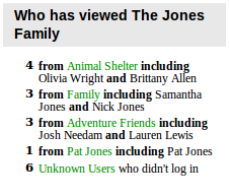
Facebook	3 (icon)		6 (icon)	
Mixed	1 (under)	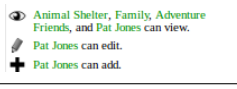	5 (sidebar)	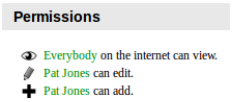
Sidebar	2 (sidebar)	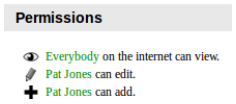	5 (sidebar)	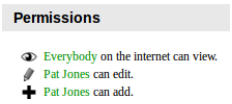
Under	1 (under)	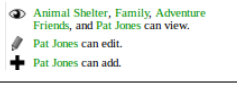	4 (under)	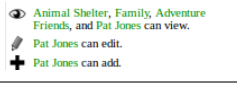

Table 7.3: Position and type of access-control proximity display shown for each condition and page.

Condition	Figure 7.6(a) (all albums) Position	Figure 7.6(a) (all albums) Display	Figure 7.6(b) (opened album) Position	Figure 7.6(b) (opened album) Display
Audit	2 (sidebar)	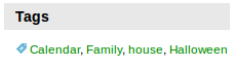	5 (sidebar)	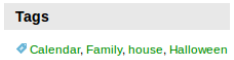
Facebook	3 (icon)		6 (icon)	
Mixed	1 (under)		5 (sidebar)	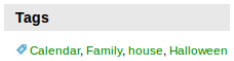
Sidebar	2 (sidebar)	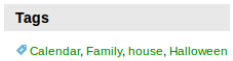	5 (sidebar)	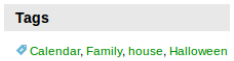
Under	1 (under)		4 (under)	

Table 7.4: Position and type of tag proximity display shown for each condition and page.

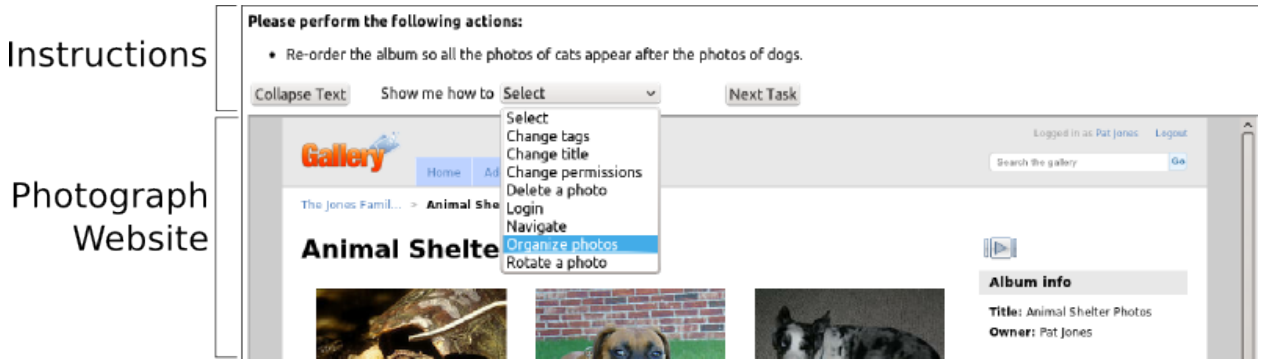

Figure 7.7: Screenshot from the online study showing the instructions and website frames. A control bar at the bottom of the instruction frame allowed participants to shrink the frame, obtain instruction on Gallery's features, and move to the next task.

In each task your friend or family member will ask you to perform a set of actions similar to the training. In addition to these actions you should also make sure the following statements are true:

- No spelling errors.
- Albums are tagged with the name of the friend or family member who took the photos.
- Photos are not sideways.
- Family albums can only be viewed by Family, and friend albums can only be viewed by Friends.
- No blurry photos.
- Pat can view, add, and edit all albums.

(a) Personal scenario ideal policy

In each task co-workers, who work with Starlight Phones and Purse Central, will ask you to perform a set of actions similar to the training. In addition to these actions you should also make sure the following statements are true:

- No spelling errors.
- Albums are tagged with the name of the co-worker who took the photos.
- Photos are not sideways.
- Starlight Phone's albums can only be viewed by contractors from Starlight Phone and Purse Central's albums can only be viewed by contractors from Purse Central.
- No blurry photos.
- Dezig Design co-workers can view, add, and edit all albums.

(b) Work scenario ideal policy

Figure 7.8: Ideal policy rules in the online study.

Study group	Completed study	Changed permissions	Dropped out
sidebar	139	114	33
under	131	116	39
mixed	124	114	44
facebook	128	112	43
audit	136	126	32

Table 7.5: The number of users in the online study who completed the study in each condition, the number of participants who made at least one change to permissions in either condition, and the number of participants who agreed to the consent form but did not complete the study.

Mechanical Turk participants would attempt to click through the study without making any changes. If a participant made no changes during a task they were prompted by an error message which indicated that they had not yet done the task. Participants were given a maximum of 2 minutes to complete each task, each prompting email reset the timer to 2 minutes, effectively giving the participant as long as necessary for task 1.

Tasks 3, 5, and 8 were all conducted on the same album. In task 3 there exists both a permission and tag errors associated with the album. When the user starts task 5 the script automatically adds three photos to the album, and the task email asks the participant to interact with these photos. No changes are made to the permissions or tags, effectively giving the participant a second chance to identify and correct the errors. When the user starts task 8, multiple errors, including permissions and tags, are introduced into the album by a script. The task email indicates that errors have been introduced but does not specify what the errors are or how many errors there are. Earlier studies showed that participants do not expect album content to change on such short notice, unless warned many participants simply do not believe that errors could be introduced so quickly or that the person emailing them would be that careless.

Task	Permission Error	Tag Error
1	Wrong group can view	Missing
2	No error	No error
3	Everybody on internet can view	Missing
4	No error	No error
5	Same as task 3	Same as task 3
6	Extra group can view	Wrong person
7	No error	No error
8	Wrong group can view	Missing

Table 7.6: Tasks and their associated permission and tag errors.

After the participant completed the training and tasks for both conditions they were asked to fill out an online survey. The survey asked them to recall permissions and tags from both Gallery sites. To test if participants were aware of the ideal policy we asked

them questions about the permissions both sites should have had. These questions were shown to the participant in random order to prevent ordering bias. The participant was also asked about past negative experiences, their own impressions about how many errors they found, and demographics.

7.3.4 Data analysis

All data from the study was placed in a MySQL database for analysis. During the study a log was kept of all the actions that the participant engaged in,as well as all the changes to photos and meta-data the participant made. After each participant completed the study, scripts automatically collected, graded, and archived all data from the session. Permission and tag information was collected, compared with the correct permissions, and compared with the default permissions. After completing the study participants used Survey Gizmo to fill out the end survey. This data was downloaded after all participants were finished and the data was put into the MySQL database.

Permission and tag correction

In this study we were unable to observe each participant's behaviors and were therefore unable to determine when they checked the permissions using the proximity display. Instead, for this study, we measured if the participant corrected the permissions/tags, as this is a strict subset of the number of participants who noticed a permission/tag error. In prior studies we have observed that some participants easily internalize that permissions are important and feel inclined to change them while other participants are unlikely to do so. In this study we observed that 37% of participants never made any change, correct or not, to the permission settings on either condition unless explicitly instructed to do so. Additionally all but 6 of those participants also never made any change to tags without explicit instruction.

Participants interacted with 5 tasks that had permission errors. Task 5 was, from a permission error standpoint, a second chance to correct the permission error in task 3, so we considered these two tasks together and only evaluated the permissions at the end of task 5. For each task we compared the resulting permissions to the default ones and marked them as either *correct* or *wrong*. We also recorded if permissions were changed but were still inaccurate. We then summed up the number of tasks where permissions were correctly changed, for each condition the participant saw. A Shapiro-Wilk test for normality showed the data to be non-normal so we used a paired Wilcox test to compare participants' permission and tag correcting performance on the control and experimental conditions. The permission correction were part of the planned tests, the remaining statistical tests were corrected using the Holm-Bonferroni method.

Non-permission error correction

Gallery stores the meta-data of all photographs and albums in a MySQL database, after each study session several scripts were run against this database to collect and code relevant

10. For the Family Calendar album which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Family Calendar Album	_ can currently view Family CalendarAlbum
Everybody on the internet	☐	☐
Adventure Friends	☐	☐
Animal Shelter	☐	☐
Family	☐	☐
Pat Jones	☐	☐
I don't know	☐	☐

Figure 7.9: Sample permission recall question from the post-survey. The question asks the participant to recall both what the permissions should have been and what the permissions were at the end of the study.

information and then archive the database. The scripts compared the meta-data, on all albums and photographs to the default meta-data. If the default value was different from the final value then the album was marked as as having been rotated, re-titled, deleted, cover changed, or permissions changed depending the meta-data. A second script was run which coded each meta-data element as *error*, *error fixed*, or *no default error*. For example, some tasks required the participant to make any change to the title of a photograph, for these subtasks any change to the title was considered correct. Some subtasks required the participant to make a specific change, in these cases the applicable meta-data value had to match the required final state.

Memory

As part of the post-study survey participants were asked recall questions about about six albums, five user groups, and one action (view). Participants were also asked what the permissions should have been according to the ideal policy. A sample question can be seen in Figure 7.9. The order of recall and ideal policy application questions was randomized for each participant.

Participants were asked about the album they used in the tutorial, the album that changed permissions (tasks 3, 5, and 8), and the album from task 6. We decided to use the album from the tutorial because all participants would have seen and modified the permissions on the album. The other two albums were selected because one of them is used on multiple tasks, and the other is used on only one task. While analyzing data we discovered a data collection error were we mistakenly asked participants about the album associated with task 7 instead of task 6 when asking about the work scenario.

At the end of the study we archive the final state of the permissions for all users. We download the answers to all survey questions from Survey Gizmo and load them into the database. We then compare the final permission state to what the user thought was the

correct permission. The result is coded as "correct," "wrong," or "I don't know." The answers to the ideal policy comprehension questions were compared to the correct settings for those albums and were also coded as "correct," "wrong," or "I don't know."

Chapter 8

Effectiveness of proximity displays

Proximity displays are intended to help end users identify permission errors, and increase awareness of their permission settings, without negatively impacting users' ability to complete their primary tasks. Theoretically, proximity displays do this by making it easy for people to notice permission settings during their normal interactions with their online photo albums. In this chapter we test these theories in controlled lab and online environments. We empirically test if different styles of proximity displays improve participants' ability to identify errors and become aware of the settings. We also examine how people interact with permissions when they are located on a secondary interface and how showing them on a proximity display changes participant behavior.

In the eye-tracker study we found that people see proximity displays, located under photo/album thumbnails, throughout the task, but tend to change permissions at the end of tasks, rarely interacting with them in the middle. In the lab study we found that participants are able to glance at proximity displays and quickly determine if there is an error. We also observed that some participants are better supported by proximity displays than other participants. Across conditions, participants tended to check permissions at the beginning and end of tasks leading us to design an interface to support this behavior. We ultimately showed, in the online study, that some proximity display designs do positively impact peoples' ability to identify and fix permission errors. However, we did not observe a difference in awareness. Additionally, none of the designs negatively impacted participants ability to fix other errors. We also observe behaviors in our online study that support the observations we made in our lab study.

We present the results of the eye-tracker, lab, and online studies in one chapter, to give the reader a more holistic understanding of how users interact with proximity displays, and the effect showing these displays has on user behavior. All the studies had similar methodologies and were all designed to answer the same set of hypotheses. We believe that, each study provides a unique view point on how people interact with permissions, and when shown together they provide a more complete picture than if described separately. To assist the reader in recalling the details of each study Table 8.1 provides an overview of the methodologies and Table 8.2 lists the conditions tested.

Name	Location	Type	Length	Tasks	Participants	Conditions
Pre-study	Lab	Between-subjects	1 hour	9	26	3
Eye tracker	Lab	Between-subjects	1.5 hours	12	34	3
Lab	Lab	Between-subjects	1.5 hours	16	33	4
Online	Online	Within-subjects	1 hour	16	658	5

Table 8.1: Methodologies used in each study.

	Control		Proximity display					Permission Modification	
Study	No Info	Tag Info	Under Photo	Sidebar	Mixed	Facebook	Audit	Full	Dialog
Pre-study	X		X	X				X	
Eye tracker	X		X	X				X	
Lab	X		X					X	X
Online		X	X	X	X	X	X		X

Table 8.2: The conditions tested in each study, details on each condition can be found in Section 7.3.1.

8.1 Hypothesis testing

In the online study we empirically tested our three main hypotheses:

H1: Correcting/checking permissions Users who see permission information on a proximity display notice errors more often than users who see permission information on a secondary page.

H2: Permission awareness Participants who see permission information on proximity displays can recall those permissions better than participants who see permission information only if they click to a second page.

H3: Negative effects Participants who see proximity displays take no more time, and correct no less non-permission errors than participants who see permission information on a secondary page.

The online study had the largest number of participants and was designed to quantitatively evaluate each of five proximity display designs showing permission information against the same designs showing tag information. Recall that the online study was a within-subjects study design so each participant saw both control and a proximity conditions. There were five different control condition designs to match the five proximity display designs, we refer to each pair of control and experimental conditions as a *treatment*.

8.1.1 H1: Correcting/checking permissions

We found in the online study that placing a proximity display with permission setting information under album thumbnails and photos (Wilcox, p=0.045), or under album thumbnails and on the sidebar (Wilcox, p=0.023), resulted in participants correcting statistically significantly more access-control permissions than they corrected using the respective control conditions. However, there was no statistically significant difference in the number

of permissions corrected between the control and experimental conditions on the sidebar treatment (Wilcox, p=0.052), the treatment which emulated Facebook's proximity icons (Wilcox, p=1.0), or the treatment that showed information about who had previously viewed the album (Wilcox, p=0.953).

To better understand the difference between how proximity displays impact the way people interact with permissions, as opposed to other types of settings, we listed the keywords with which albums were tagged on the proximity displays as the control condition. We saw no statistically significant difference in the number of tags corrected between showing tag information on a proximity display (control) and permission information on the proximity display (experimental). The largest difference between control and experimental was observed in the under-photo condition where participants corrected an average of 0.91 *fewer* tag errors if they saw the tag information on a proximity display (control) than on a secondary interface. In both the under and mixed conditions participants were more likely to correct tag errors if they saw permission information on the proximity displays then if they saw tag information.

	Permissions corrected out of 4				
	Wilcox	Control		Permissions on proximity	
Condition	p-value	Average	StDev	Average	StDev
under	0.045	0.924	1.316	1.176	1.444
sidebar	0.052	0.784	1.19	1.007	1.283
facebook	1	1.094	1.422	1.094	1.45
mixed	0.023	0.774	1.202	1.048	1.378
audit	0.953	1.14	1.394	1.147	1.443

Table 8.3: Average number of permissions corrected in the control and experimental conditions. Results of statistical significance Wilcox paired t-test (within-subjects).

8.1.2 H2: Permission awareness

After completing all tasks in both conditions, the online study participants were asked to fill out a post-survey. Participants were asked to recall the current permissions for three albums in each condition. Participants answered 30 permission recall questions about three albums, all five user groups, and one action (view). In the analysis of permission awareness we exclude results from the training albums which all participants were forced to edit permissions on, this leaves four albums, two from each condition. Additional details about the data analysis and question format can be found in Section 7.3.4.

The results from the two non-training albums show no statistically significant difference in the number of permissions remembered between conditions in any of the treatments (Table 8.4). Participants recalled an average of 6.5 permissions out of a total of 10. For comparison, participants recalled an average of 7.5 permissions out of 10, when asked to recall the permissions from the training album which they had all previously set.

While participants in the mixed treatment showed the most improvement between control and experimental conditions, the largest differences in memory were observed in the

	Permission settings recalled out of 10				
	Wilcox	Control		Proximity	
Condition	t-test	Average	StDev	Average	StDev
under	1.0	6.466	2.813	6.779	2.946
sidebar	1.0	6.345	2.807	6.525	2.793
facebook	1.0	6.586	2.912	6.414	2.822
mixed	1.0	6.242	2.73	6.742	2.814
audit	1.0	6.676	2.751	6.721	2.685

Table 8.4: The online study participants' ability to recall permission settings for two non-training albums (5 questions each). Reported p-values reflect a Holm-Bonferroni correction.

user group (Friends, Family, ect.) the participant was asked the question about. Memory questions about the users groups associated with the participant themselves having access (ie: the group Pat) showed the highest memory recall with 81.2% of experimental condition participants and 79.9% of control condition participants accurately recalling the permissions. These groups were unique in the study in that no error was ever associated with them, so their final, initial, and ideal states should be identical and participants should have no reason to ever change them. The group associated with the training task and the group associated with the album that changed showed the worst memory results. Participants recalled between 46% and 51% of permissions associated with those groups regardless of condition. These groups were also the most likely to be incorrectly answered as opposed to being marked "I don't know." Participants answered between 25% and 38% of these questions incorrectly.

In addition to recall questions we also asked participants what the correct permissions were for these albums. No treatment exhibited a statistically significant difference between conditions (Table 8.5). This was expected and shows' that participants in all treatments understood the correct permission state for the albums.

	Ideal policy recalled out of 10				
	Wilcox	Control		Proximity	
Condition	p-value	Average	StDev	Average	StDev
under	1.0	6.809	3.55	7.099	3.601
sidebar	1.0	6.813	3.564	7.108	3.629
facebook	1.0	7.25	3.514	7.133	3.249
mixed	1.0	6.855	3.414	7.113	3.45
audit	1.0	7.125	3.332	7.199	3.326

Table 8.5: The online study participants' ability to apply the permission rules in the ideal policy for the two non-training albums per condition (5 questions each). Participants were asked what permissions Pat/Pat's boss would have wanted to set. Reported p-values reflect a Holm-Bonferroni correction.

8.1.3 H3: Negative effects

As can be seen in Table 8.6, participants in the online study exhibited no significant difference in the number of non-permission tasks corrected. Because the online study was time limited we cannot make any claims about time required to complete the tasks. In the lab study, which was not time limited, the control condition and experimental conditions showed no significant difference in either time to complete the tasks or number of non-permission subtasks corrected. We therefor conclude that showing the proximity display does not cause a negative impact on time or accuracy of other tasks.

	Non-permission and non-tag errors corrected out of 37				
	Wilcox	Control		Proximity	
Condition	p-value	Average	StDev	Average	StDev
under	1.0	26.588	4.474	26.672	4.657
sidebar	1.0	27.079	4.188	26.698	4.995
facebook	1.0	27.68	3.669	27.102	4.819
mixed	1.0	26.75	4.121	26.823	3.984
audit	1.0	27.353	4.076	26.882	5.369

Table 8.6: In addition to tag and permission errors, the online study participants were asked to correct issues with the titles, organization, orientation, and content of photographs. This table reports the number of non-permission and non-tag errors the participant corrected out of 37 errors. Reported p-values reflect a Holm-Bonferroni correction.

8.2 How people notice and fix permission errors

The lab study was designed to collect a large amount of qualitative data to better understand how participants participants notice permission errors. In this section we primarily present the results from the lab study, where appropriate, we also present data from the online study which supports or contradicts our lab study conclusions, and the eye-tracker study to discuss when participants look at proximity displays.

We find, in the eye-tracker study, that participants see permissions on proximity displays throughout each task (Section 8.2.1) but they appear to correct the permission errors at the beginning and ending of tasks (Section 8.2.3). We observe that some participants check permissions rarely, these participants benefit the most from seeing proximity displays. Conversely, some participants check permissions frequently with little provocation. These participants tended to check permissions less often when shown a proximity display (Section 8.2.2).

8.2.1 Noticing permissions

One of the goals of proximity displays is to enable participants to notice, and check permissions, quickly and easily. In the lab study we observed that participants who see proximity

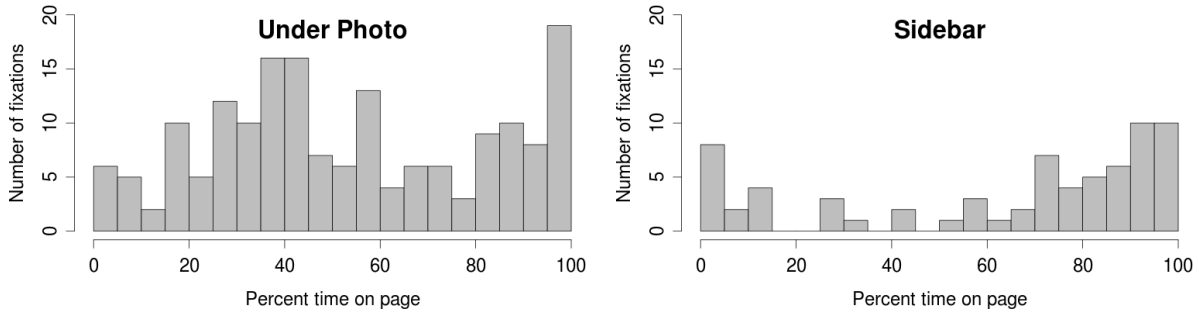

(a) Under-photo (b) Sidebar

Figure 8.1: Histogram of the number of fixations for all participants (y-axis) against the amount of time spent in the page, normalized (x-axis).

displays are observed to check permissions less often than control participants, but both groups corrected the same number of permissions. In this section, we explore how people become aware of an error in their permissions, specifically looking at the process people go through when finding permission errors using the displays.

To understand how people are noticing permission errors we use data from the eye-tracker study, where we used an eye-tracker to determine when participants are looking at proximity displays. Figure 8.1 is a histogram of all the instances where participants, in the under-photo and sidebar conditions, fixated on a proximity display. A *fixation* is an eye tracking term for when the participant's gaze rests on a single point on the interface. We looked only at the webpage the participants worked with in advance of modifying permissions, and we normalized participants' time on the page to make the time at which fixations occurred comparable across participants. The majority of participants stayed on a single album page for the length of a task. Time 0 on the graph represents the instant when participants opened the album page, and time 100 represents when participants opened the permission-modification interface which is located on a new page. Participants in all three conditions spent an average of 4.4 minutes on a page before opening the permission-modification interface. We observe that under-photo participants fixated on the proximity displays throughout the task, but explicitly checked and corrected permissions at the end of the task (Figure 8.1(a)). Sidebar participants looked at the display just before transitioning to the permission-modification interface, but looked at the display rarely before that point (Figure 8.1(b)). Participants tended to read the printed email at the beginning and ending of tasks so the slight decrease in fixations at the beginning and near the end for under-photo is likely an artifact of participants not looking at the screen.

Given that the under-photo condition places proximity displays all over the screen, it is difficult for a participant to not see or fixate on a display during the course of a task. So we have to ask if participants are in fact noticing the permissions or just staring at the proximity displays without really looking at them. To test if participants re really noticing the displays in advance of changing them we designed the lab study to have permission errors randomly ordered and assigned to tasks. This minimized the effect task wording and order had on when permissions were checked. Further description of the lab study

methodology can be found in Section 7.2.2.

To better understand whether participants were able to notice permission errors easily, we needed to measure noticing behavior separately from permission correcting behavior. We use the terms *checked* or *noticed* when discussing what the participant actually saw, the value we are attempting to measure (ground truth). In the lab study we approximated when participants notice permission errors by measuring *explicit checking* behavior. Control participants were said to have *explicitly checked* permissions if they opened the permission-modification interface. Participants who were shown proximity displays were said to have *explicitly checked* permissions 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 (moving the mouse under the words while reading or circling the display with the mouse); or (4) pointed at the permission display with their hand while staring intently at the screen. The identification of explicit permission-checking behavior in the under-photo condition was done by the researcher during the lab study.

We compared the number of times participants explicitly checked permissions on tasks with errors to tasks without errors (Figure 8.2(a)). Participants in the control condition were equally likely to explicitly check permissions on tasks with and without permission errors, this is expected because participants had no way of knowing if an error existed without opening the permission-modification interface. Consequently, for the control condition, we had a very accurate measurement of how often a permission was checked. Participants in the under-photo condition were more likely to explicitly check permissions on tasks with permission errors than on tasks without permission errors. On average, under-photo participants explicitly checked permissions on 3.2 tasks with errors and 1.7 tasks without errors. This suggests that participants are paying attention to the display more often then we are observing through measuring explicit checking behavior, because our measurement definition was unable to capture all the permission-checking events.

Why, in the lab study, are we not observing every permission-checking event? The eye-tracker data suggests that participants are fixating on the displays mid way through the tasks, but the permission-modification behavior suggests that they explicitly check permissions occasionally at the beginning and frequently at the ending of the tasks. While participants eyes may be fixating on a proximity display, we can not be sure that they are absorbing the information, so the eye-tracker is likely overestimating the number of times a participant paid attention to a proximity display. Similarly, think aloud data only captures information the participant processes sufficiently to articulate in a linear format. We know, from the lab study (Figure 8.2(a)), that participants are noticing the absence of an error and not sufficiently exhibiting this notice event in the think aloud verbally or through behavior for us to measure them explicitly checking. Think aloud theory tells us that this is because the information is either: 1) non-linear and therefore challenging to verbalize, or 2) in working memory for a very short time [104].

When asked, lab study participants reported noticing permissions while working on other issues. Because they were distracted they put fixing the permissions off until the end of the task, when they might or might not remember the error.

I was more just focused on getting this done first. I felt like if I looked at the

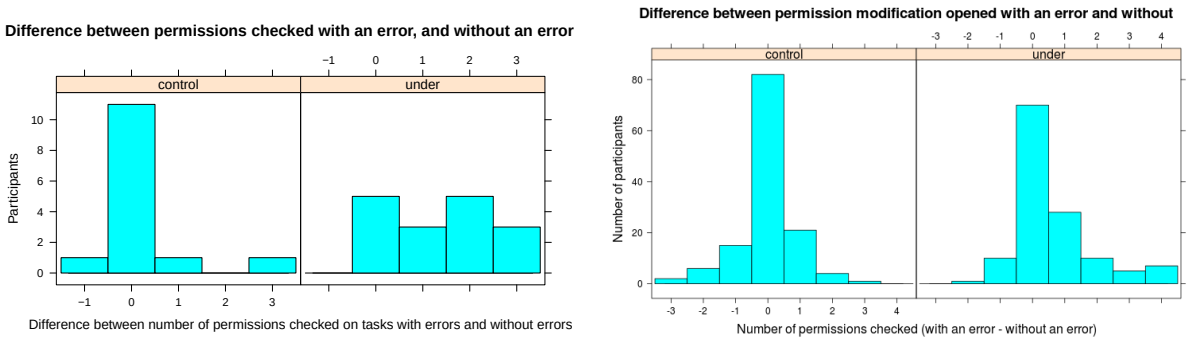

(a) Lab study (study 3)

(b) Online study (study 4)

Figure 8.2: Histogram of the number of tasks where the participant checked permissions and there was an error subtracted by the number of tasks where the participant checked permissions and there was not an error. For example, we can see in the lab study that 11 participants checked the same number of permissions in tasks with errors as they did in tasks without errors in the control condition. In the lab study (a) "checked" was defined based on observed behavior. In the online study (b) "checked" was defined as opening the permission-modification interface.

> permissions like or if I glossed it over, I just wanted to get this stuff done first, and thinking that I would go back to it but I never ended up doing that. – *Under Full*

If we look at participant behavior in terms of the C-HIP behavioral model [112] and the HITL framework [30] described in Section 2.4, we can better understand what is happening. The C-HIP model describes a set of states a user must go through between when a warning becomes visible and the warning actually effecting behavior. According to the model, once a user has looked at the warning (Attention Switch) they decide if the warning is worth focusing on (Attention Maintenance). If the warning is worth focusing on, they try and understand the warning, and comprehend what the warning is saying (Comprehension and Memory). Once they comprehend the warning, they decide based on their beliefs if the warning applies to them (Attitudes and Beliefs). Finally, if they consider the warning relevant, they may be motivated to do something (Motivation) that may result in a change in their behavior (Behavior). We hypothesized that our participants were making the attention switch (eye-tracker in the eye-tracker study reports fixations), then glancing at the proximity display to see if it "looked wrong" (Attention Maintenance). If the display looked wrong they would focus on it (Comprehension and Memory), which is when we observed them explicitly checking permissions. At this point our participants either made the permission change or waited until the end of the task and then made the change (Motivation and Behavior). However, if the display did not look wrong participants saw no need to continue focusing on the display (Attention and Maintenance), and moved on to other tasks before they tried to comprehend the content (Comprehension and Memory).

We hypothesize that because we gave participants a single access-control policy that was globally applicable, they were able to learn to quickly differentiate between proxim-

ity displays that showed correct and wrong permissions. The quick differentiation ability allowed them to glance at the displays and determine if they looked correct during the Attention Maintenance stage of the C-HIP Model without having to move to the Comprehension Memory stage where they would have 1) transformed the contents of working memory to a form that was easy to vocalize, or 2) had to keep the information in working memory for long enough to vocalize.

During the unstructured interview at the end of the lab study, the researcher asked participants how they had identified permission errors. Participants in the under-photo conditions talked about the heuristics they used. Instead of reading the whole policy a participant said he would just look to see if Everybody could view the album:

> If it was company related then it should say Everybody and if it didn't say Everybody then it was wrong, and I would know that just by looking. – *Under Dialog*

Another participant showed the researcher what correct and wrong policies looked like. His primary metric appeared to be the length and shape of the words on the display:

> [Indicates proximity display on the screen] if there is like a lot of things I will look because there is only one a few things you should have up as the permissions. – *Under Full*

In the online study we measured the number of times the participant opened the permission-modification interface (Figure 8.2(b)). We observed that under, sidebar, and mixed condition participants were statistically more likely to open the permission-modification interface if there was a permission error than if there was no error (all three p-values $<$ 0.001). Participants in the control conditions were equally likely to open the permission-modification interface whether a permission error existed or not. This shows that participants in the under, sidebar, and mixed conditions were able to use proximity displays to identify errors and avoid needlessly opening the permission-modification interface when an error did not exist.

These results tell us that proximity displays need to provide sufficient information for a user to determine if an error likely exists without interacting with the display. Participants feel that they can determine if an error exists quickly and if they do not notice one they will likely move on without focusing sufficiently to realize their mistake.

Checking permissions all at once

Another reason we decided to provide participants with a single ideal policy in the lab study was to provide a more natural environment. People typically know their own policies. By providing a single policy we allowed participants to choose when to make changes instead of forcing them to make changes during a specific task. We found that some participants, regardless of condition or permission-modification interface, have a tendency to take a single pass through the whole policy, correct all the permission errors, and never look at the permissions again. We term this behavior *checking all at once*.

Participants who check permissions all at once never check any permission again after doing so, even though some of the decisions they made in their permission-correcting pass

were wrong. As part of the unstructured interview the researcher asked these participants whether they had considered looking at the permissions again, or if they had been concerned about errors. Some participants reported briefly debating whether a permission was wrong but had decided to do nothing about it.

This observation tells us that some participants want to correct permissions in a single pass and not think about them again. This tendency appears to be of the treatment and condition, though some display designs might assist the behavior more than others. The disinclination to check permissions again indicates that though these participants correct more permissions in their one pass they are susceptible to missing errors introduced after the checking took place.

The online study was time limited, which discouraged participants from checking all at once. However, we still observed 92 participants (14%), from all treatments, correct permissions on more than one album during at least one task. Twenty of those users corrected permissions on more than two albums during a single task (3% of all users).

8.2.2 Participants' tendency to check permissions

We observed a high variance between lab study participants in their behavior towards permissions. Some participants completely ignored permissions, and some participants took checking permissions very seriously. This made it difficult to determine if the permissions were being checked because of the interface or because the participant was pre-disposed to check them. To account for the difference, we made the online study a within-subjects study so that the same participant would experience both the control condition and an experimental condition.

The lab study participants in both under-photo dialog and control-dialog conditions appear to check permissions either frequently or not at all (Figure 8.3(a)). We hypothesize that some unobserved variable causes some participants to frequently check permissions and some participants to rarely check permissions. We refer to these two groups as *infrequent permission checkers* and *frequent permission checkers*. In the lab study we define frequent permission checkers as those who check permissions on more than half of tasks.

Infrequent permission checkers checked permissions on an average of 1.8 tasks, and frequent permission checkers explicitly checked permissions on an average of 8.3 tasks, if they saw the proximity display (under), or 12.5 tasks if they did not see a proximity display (control). As discussed in Section 8.2.1, our measurement for explicit checking behavior underestimates the true number of times an under-photo participant checks permissions. Underestimating permission checking is the reason for the difference between under-photo and control participants who checked permissions frequently.

In the online study we looked at the number of tasks where the participant opened the permission-modification interface, as a way to measure how often they were checking permissions. As previously mentioned, this is an accurate measure of the number of times participants check permissions during the control condition, but is a lower bound on the number of times experimental condition participants checked permissions. Similar to the results from the lab study, participants in the online study frequently never checked permissions (Figure 8.3(b)). However, those participants who do check permissions tend to

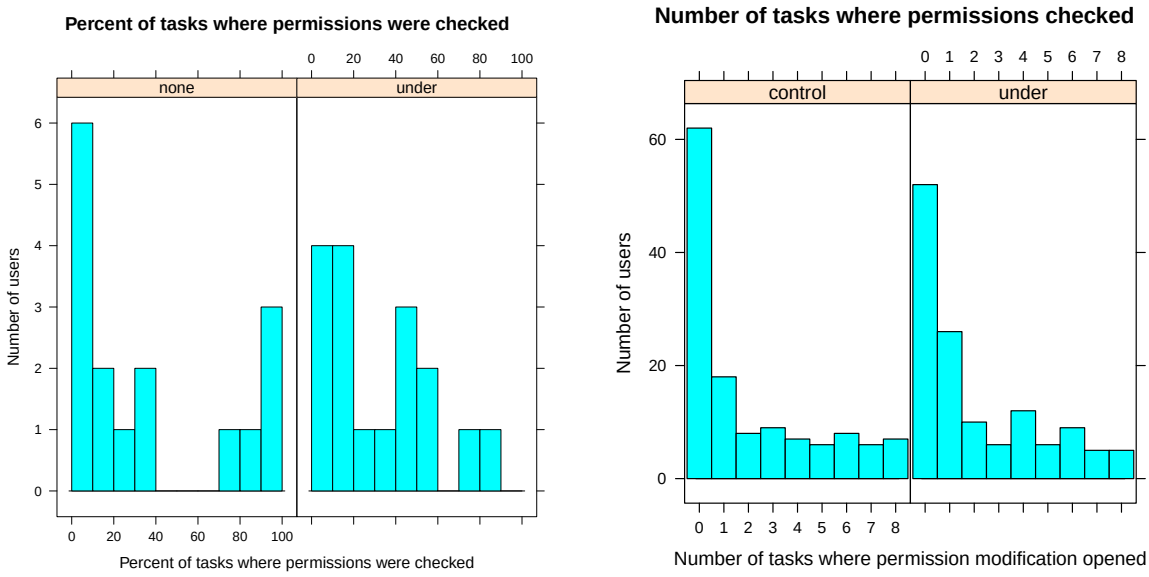

(a) Lab study (study 3)

(b) Online study (study 4)

Figure 8.3: Graph *a* shows the percentage of the 14 non-training tasks where the participant checked the permissions by the presence or absence of a proximity display. Graph *b* shows the number of tasks where the permission-modification interface was opened for participants in the under-photo condition. Graphs of other conditions are nearly identical.

check a similar number in both conditions. There is a positive linear relationship between the number of tasks where the permission information display is opened in the control and experimental conditions regardless of the current treatment (linear model, p ¡ .001).

The decision to check permissions or not check permissions appears to be an individual choice based on an unknown variable. In the lab study post interview participants were asked why they did or did not check the permissions. When asked why they checked permissions, most lab study participants responded that checking permissions was part of their job, or that three of the rules mentioned permissions. Participants who checked permissions in the majority of tasks could not even seem to understand why the researcher had asked the question. From their perspective the fact that permissions should be checked was obvious.

> But it said it was your job. You know what I mean, if you could loose your job because you screwed it up then why wouldn't you .... – *Control Full*

Some participants mentioned that managing permissions was important to them in their own lives, or that the rules resonated with them.

> There were guidelines explicit in the instructions that had to do with what the boss Gerald wanted like access to certain albums. And also personally with my privacy settings on the internet I want to make sure that my albums are only available to people I want. – *Control Dialog*

Across conditions, 18 participants (54.5%) checked permissions in less than half of the

tasks. When asked why they did not check permissions participants came up with a wide range of answers, but the primary answer is well summarized by the following quote:

> I think I may have forgotten about the permissions. – *Control Full*

In the unstructured interview, researchers asked participants why they were not checking or changing the permissions despite the fact that they were making sure the non-permission requirements were being met.

> With all the other errors since they were right in front of me I could just see them and they kinda triggered my memory that way. I guess without the permissions error being there I couldn't... it didn't just pop in my head. – *Control Dialog*

Despite the lab study post-interview, we were not able to determine what variable caused some participants to forget to check the permissions and others to check permissions frequently. The post-survey in the online study included multiple questions to determine if attitudes, opinions, or past experiences had an effect on participants' tendency to check permissions during the control condition. We created these questions based on the work by Wang et al. on Facebook regrets [108], and Tsai et al. on online shopping privacy [103]. The only question to show any correlation was: "Do you agree or disagree with the following statement: Most businesses handle the personal information they collect about consumers in a proper and confidential way." Participants answered using a five point Likert scale. Participants who agreed with this statement were more likely to change permissions in the control condition (corrected ANOVA, p=0.043).

Proximity displays impacted participants differently depending on their permission-checking behavior. In the lab study, participants who checked frequently appeared to be negatively impacted by the proximity display. The numbers were too small for meaningful statistics but the researcher observed several cases where an under-photo participant, who was clearly checking permissions on every task, forgot to check on one or two tasks and during the post-survey was certain that they had checked permissions on every task. We believe that by encouraging glancing at permissions we decreased the amount of focus participants gave to the act of checking permissions and thereby increased the number of errors missed. Social psychology tells us that tasks that receive less focus are more likely to be forgotten [28]. We have seen this effect before in error identification work [105], where participants felt that they would notice if the screen showed an error, failed to notice the error because it was insufficiently "obvious," and were therefore very confident that no error existed regardless of the truth.

Participants who checked infrequently showed the opposite trend: under-photo participants tended to check permissions on about one more album than control participants. This trend was visible in both the lab and online studies (Figure 8.3). Looking at the online study we see that participants in the under-photo condition corrected only 0.25 more permission errors in the experimental condition than the control condition, and mixed saw only a 0.77 improvement on average (Table 8.3). These numbers are not large and reflect the fact that most participants either check permissions on 1-2 albums more when they see proximity displays than they would on control, or they check permissions on neither.

These numbers indicate that proximity displays are being used as passive displays. Recall that proximity displays are intended to be passive and only be noticed by users occasionally. They are intended to assist infrequent permission checkers by enabling them to easily check for errors at any time. The result that participants in the under-photo and mixed conditions check a permission more than when they are in control, shows that the displays are fulfilling their intended role of occasionally assisting people. Proximity displays may be less beneficial in environments where checking the permissions frequently is important since participants may be more likely to miss errors when glancing than when explicitly checking.

8.2.3 When do people change permissions

One of the purposes of proximity information displays is to provide permission information to end users in a way that naturally fits into their normal workflow. We want to introduce permission information to people at the time and place when they will most likely have need of it and be receptive to it. We wanted to know 1) where/when do people naturally become interested in permission information, and 2) how can we manipulate the proximity display design to best support this behavior?

When we began testing proximity display effectiveness, we expected that the proximity displays' spatial proximity to the users' main focus (the photos), would closely approximate the integration with their workflow. Hence, we expected that the under-photo condition would outperform the sidebar condition, in which displays were not spatially located near the participants' primary focus point. The eye-tracker study showed this to be true, with the under-photo condition outperforming the sidebar condition [106], and later the online study showed the same thing (Table 8.3). However, putting the proximity display under every photo takes a large amount of screen real estate and potentially distracts users, so we wanted to use our understanding of how permission errors are identified to find a more appropriate solution. We observed the following permission checking behaviors: 1) participants check permissions at the beginning and ending of tasks, 2) participants tend to view permission errors as dissimilar to the other errors they are looking for, and 3) when proximity displays are shown under album/photo thumbnails, participants tend to check permissions using the display located under the album thumbnail.

Participants check permissions at the beginning and the end of tasks

During the pre-study and eye-tracker study we observed that the majority of the participants explicitly checked and changed permissions at the beginning and end of tasks. This was observed across conditions and across tasks. At the time we hypothesized that the behavior was due to participants' need to go to a separate page in order to change the permissions. To test this hypothesis we introduced a permission-modification-dialog condition into the lab study. Half of the participants were given the full-page permission-modification interface used in the eye-tracker study that required the participant to switch web pages. The other half of the participants were given the dialog permission-modification interface, which did not require switching pages and therefore changing a title, which also brings up

a dialog.

Participants in the lab study engaged in a wide number of activities, and because they were not time limited, they took a wide range of times to complete tasks. Because of the wide variation in times, we analyzed the order participants chose to engage in the different actions. The analysis of the order (Figure 8.4) showed that participants predominately engaged in a permission-modification action as the last action they engaged in . To determine this, for each user and action we ordered the possible action types based on the first timestamp associated with each action type (which action was done first). We also ordered the possible action types based on the last timestamp associated with each action type (which action was finished last). We refer to the last time an action was engaged in as when the action was *completed*. A more detailed explanation of this analysis can be found in Section 7.2.4. There was no significant difference in action order between conditions.

In the lab study post interview, participants who were obviously checking permissions at the end or beginning of tasks were asked why they were doing so. Most people did not know why they were checking at the end and instead talked about how they had approached the tasks in general. The following quote describes typical user behavior:

> I think maybe because in the beginning I was jumping around just exploring the whole thing. And not really paying as much attention. Then I methodically went through and just, and it is pretty easy to just mouse over stuff, so it did not hurt to check. – *Under Dialog*

Participants reported being very focused on the distractor subtask and the non-permission errors. They talked about permissions as a different type of task from the non-permission tasks.

> I guess I read the task and did everything that was required of me and left monitoring, personal monitoring for, you know, the last stuff. It was just easier to do what I had to do first, or just perform the request and then make sure that the policy was followed.... Seemed more intuitive the way I did it. – *Under Full*

A control-full user explained this tendency to put permissions last the best. She likened setting the permissions to remembering to turn off the oven and then later decided it was closer to locking the door at night. Both were tasks that she always had to explicitly remember to do before going to bed.

> Just because it comes at the end doesn't mean that it is unimportant to me. It probably means that ... That is like the final, this is it now. You do everything you are supposed to do before you go to bed then you make sure you lock the door. So that is like locking the door, checking those permissions, that is like the final security piece. – *Control Full*

Participants deliberately decided to not modify permissions until they were done checking the other requirements (Figure 8.4. Visually obvious actions such as sideways photographs are engaged in first or in the middle, and are the first subtask to be completed. Less visually obvious actions such as renaming, which includes both spelling errors and

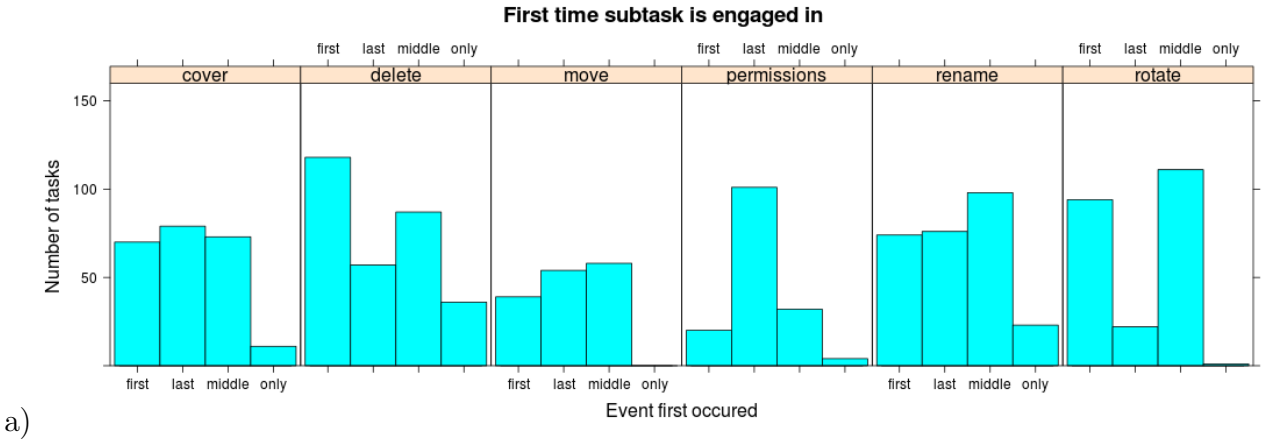

a)

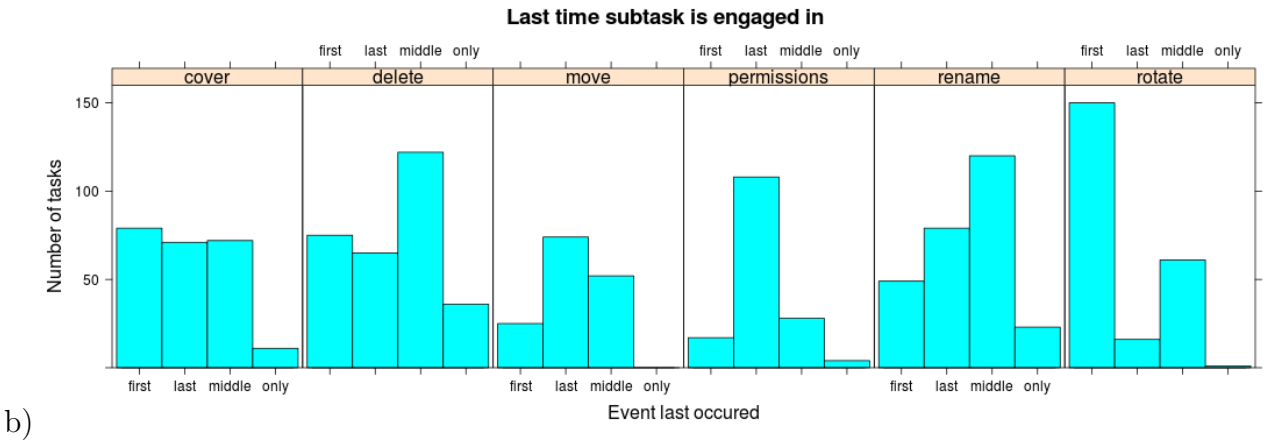

b)

Figure 8.4: While working on a task participants were free to engage in actions in any order, including interleaving actions. For example: a participant could rotate a photo, delete a photo, then rotate a photo. Graph *a* shows the first time an action of that type was engaged in during a particular task and whether that action was the first action, the last, neither first nor last (middle), or the only action engaged in. The height of the bars indicates the total number of tasks across all users; the summation of all bars in a subgraph is the number of tasks, across all users, where that subtask was engaged in at least once. Graph *b* is similar to graph *a* except that it shows the last time a subtask is engaged in during a task.

changes specified by the email, are initiated at any time and are rarely the first action to be completed. Actions that require a large dialog and focused attention, such as organizing photos (move), are engaged in at any point with a minor bias towards later and tend to be completed last.

After the lab study we decided that we needed to better support permission-checking at the beginning and ending of tasks. We were also concerned that putting information under every photo was causing participants to naturally ignore it as information they did not need right now. Finally, we were concerned that when participants see a display everywhere they just assume that they will spot an error if it exists and therefore don't focus on the displays enough to actually check them.

To address these concerns we added the *mixed* condition to the online study. This condition shows the proximity display under the album thumbnail, and when the album is opened the proximity display is shown on the sidebar instead of under every photo. The intention was to encourage participants to notice permissions as the album was opened, or at the end of the task after the album was closed. Participants who choose to check permissions in the middle could easily do so with a quick glance at the sidebar, but information was not pushed at users while they were actively engaged in other tasks.

As can be seen in Table 8.3, the mixed condition outperformed all other conditions in terms of number of permission errors corrected. We also wanted to know whether the mixed interface fit better into participants' work flows than other interface designs. To answer this question we plotted all the times a permission-modification interface was opened (Figure 8.5).

Permissions perceived differently from other errors

Participants talked about how challenging it was to have to think about both permission and non-permission errors at the same time. Participants were given five rules (Figure 8.6) that they were supposed to enforce when interacting with albums. However, they appear to have considered the permission rules to be different from the others.

> It was hard it was kinda balancing two aspects, it was either like maintaining the policy like the whole alcohol thing and at the same time making sure it was like a ... it is not just open to everybody it was exclusive to some people who can see it and understand it. If it was selective in that sense that people could see it that you wanted them to see it, you know, the alcohol might have been ok or the policy might not have applied as strictly. But it was like trying to balance. – *Under Dialog*

As part of the verbal lab study post-survey, participants were asked to recall the boss's rules (ideal policy) in their own words. Figure 8.6 shows how many participants remembered the policy rules and the order in which they recalled the rules. The majority of participants first recalled rules 1 and 5, which have to do with alcohol, blurriness, rotations, and spelling errors, and then recalled rules 3 and 4, which concerned permissions. Rule 4, which specified who could add to or edit albums, was rarely recalled. There was no significant difference in the rules remembered amongst conditions.

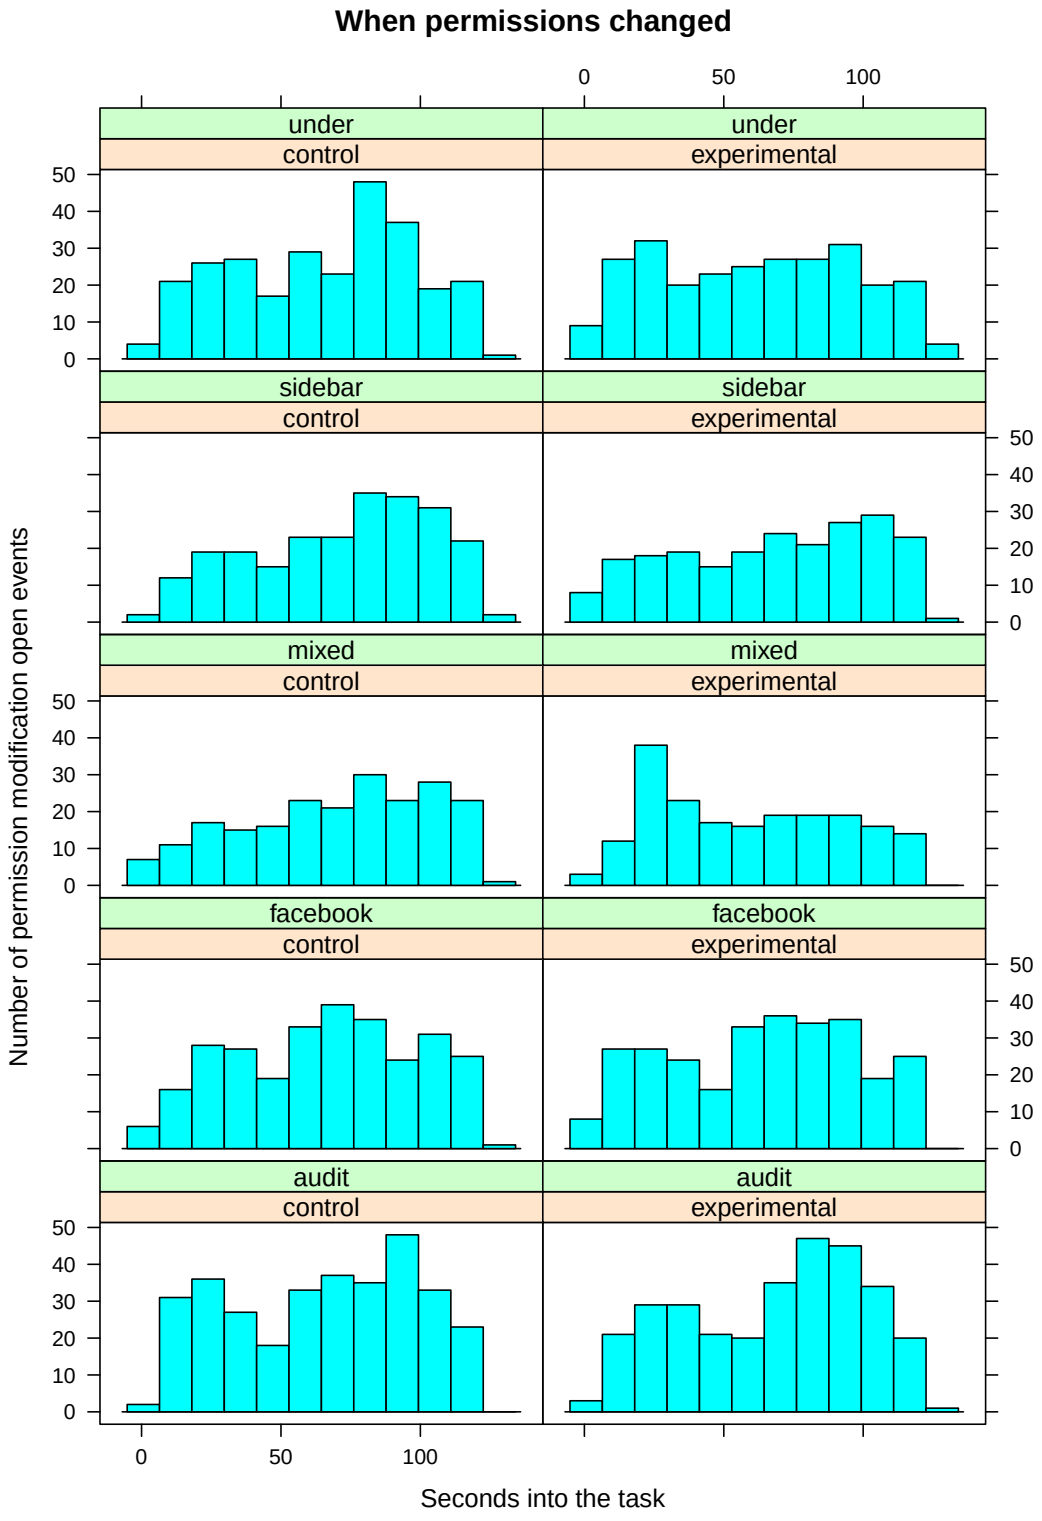

Figure 8.5: The number of seciond into the task when the permission-modification interface was opened by participants in each condition. Events from task 1 and the training are excluded to remove bias caused by prompting the participant.

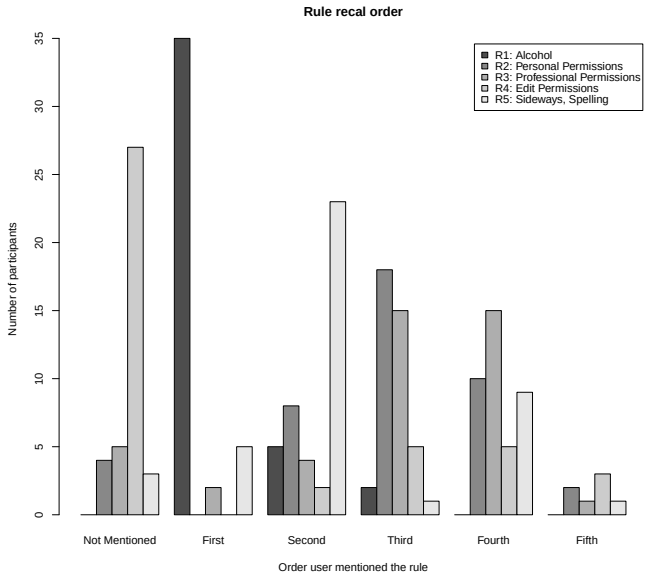

Gerald's Photograph Policy

1. No photographs containing drugs, alcohol, or anything inappropriate.
2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.
3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.
4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.
5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

Figure 8.6: As part of the verbal post-survey participants were asked to recall Gerald's rules, in their own words. The above graph shows the order in which participants recalled the rules. As the graph shows the majority of participants recalled the rules in the following order: R1, R5, R2, R3 and forgot to mention R4.

Combined with post-study interviews, the information from Figure 8.6 suggests that participants are mentally grouping permission rules as different from the non-permission requirements. Additionally, participants appear to think of permissions after thinking about the other type of requirements. The way people are grouping types of errors is important because it may help explain why participants are changing permissions first or last. If permissions really are perceived as different than the other actions, then checking them may require participants to swap out working memory. People would not want to change what they are thinking about multiple times in a task, so they wait until the end and change what they are thinking about then.

Participants in all treatments tended to open the permission-modification interface at the end of the task when they were experiencing the control condition (Figure 8.5, column 1). However, participants who saw permission information in the mixed condition (Figure 8.5, column 2, row 3) did the opposite and tended to open the permission-modification interface at the beginning of the task. This is particularly notable since the same participants had the opposite behavior in the control condition.

Participants in the lab study talked about how permission errors were different then the other types of errors they were looking for. One potential difference between the types of errors might be the participants' pre-study understanding of "correct" and "wrong" action states look like. Participants entered the study with a well-practiced ability to identify spelling errors and sideways photos. We did not have to impart what correct and wrong states were for these error-identification subtasks. Even the rule about no alcohol

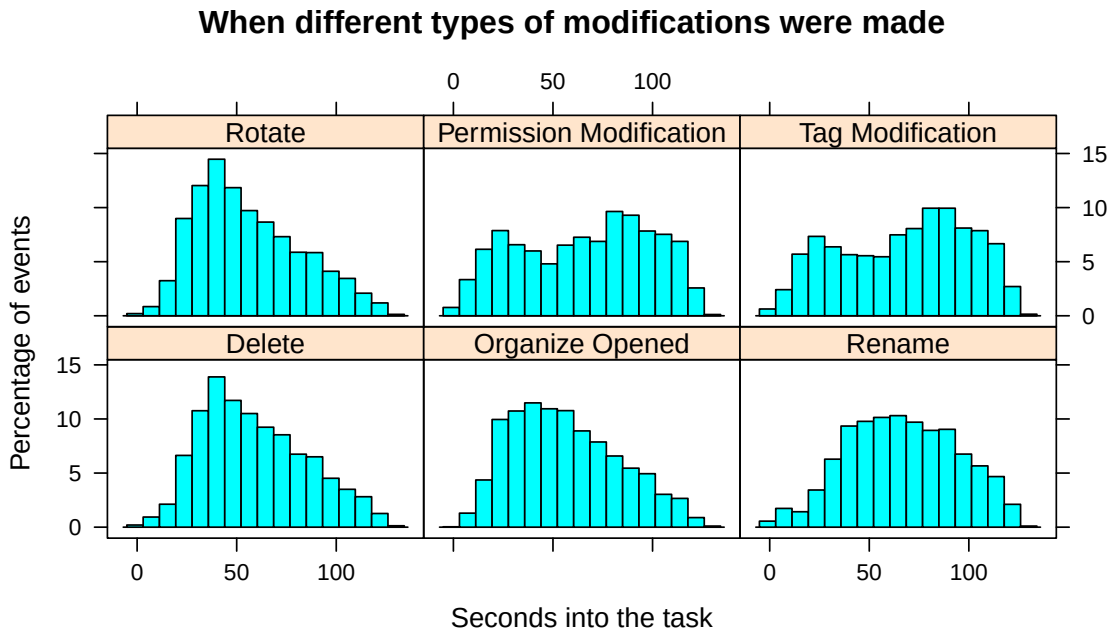

Figure 8.7: Number of seconds into a task that an action was done. Histograms show all participants across all conditions, both with and without permission proximity displays. Events from task 1 and the training are excluded to remove bias caused by prompting the participant.

in photos was reasonably familiar to users and several commented how they normally do not post that type of photo. However, participants had to be told what the "correct" and "wrong" permission states were. So, permission errors might have been different because participants did not have prior experience identifying those types of errors.

We wanted to know whether permissions are modified first and last because people intrinsically modify them then, or because participants had to learn to identify new "correct" and "wrong" states for this study. In the online study we accounted for this by introducing tags as a new action type that could have an error. Similar to the permissions, we told participants what "correct" and "wrong" tag states looked like. Because the correct tag state was artificial, participants would have had no prior experience looking for these errors. We also put tag information on proximity displays during the control condition so participants received the same type and amount of exposure to tags as they did to permissions.

Figure 8.7 shows at what point during tasks participants engaged in each type of action. Rotate, delete, organize, and rename subtasks look similar in that they resemble skewed normal distributions with the majority of participants selecting similar times to make changes. Rotate and delete actions are typically done first, while the more complex rename and organize actions were done in the middle. These results are similar to the ones we observed in the lab study (Figure 8.4). However, if we look at the graphs for permission and tag modifications we see that participants modify both tags and permissions at the beginning and ending of tasks. This suggests that the reason participants are modifying

permissions first and last is that the "correct" state is one participants have no prior experience identifying. Therefore, it may be that participants needed to focus more attention on the permissions, than on the other actions, to determine if an error existed.

While familiarity with identifying correct and wrong permissions states may have been a factor in participants' tendency to check permissions last, it is unclear if familiarity would significantly alter behavior. In the lab study we gave participants 16 tasks and observed participants constructing heuristics and exhibiting pattern-matching behavior as methods of detecting correct and wrong permission states. In theory, these participants should have had enough experience, near the end of the study, to accurately notice permission errors with minimal effort, as evidenced by them noticing permissions in advance of explicitly checking them. Despite noticing permission errors quickly, the lab study participants still performed permission-modification as the last action. Additionally, they self reported doing this purposefully. While the fact that the ideal permission policy was not their own likely effected when permissions were changed, we believe that the way people think about permissions was also a major factor.

Where permissions were checked

In the lab study, under-photo condition participants had several places where they could check permissions: under the album thumbnail, under any of the photo thumbnails, or by opening the permission-modification interface from the album thumbnail view or the photos view. The researcher noticed a behavior where participants in the under-photo conditions would exit the album when finished with a task, mouse over the album thumbnail, explicitly check that the permissions were correct, and then declare themselves finished with the task. One of the participants had this built into such a routine that he did not believe that the proximity display was even visible once an album was opened.

> So once I open then I finish what I'm doing with the task and go back and look at the album and look at the permissions because you can't just see it right away. Um. Or you have to do it right away and then perform the task. – *Under Full*

Conversely, some participants would click on the album thumbnail to open the album and appear to glance at the permissions while the album page was loading. These participants did not explicitly check permissions according to our definition but the eye-tracker indicates that they were looking at the display. Our data shows that the majority of lab-study participants explicitly checked permissions at the end. However, not all lab study participants changed permissions as the last thing they did. Three participants in the under-photo condition, who checked permissions on more than 50% of the tasks, made a point of checking permissions first on nearly every task.

In the online study, we observed (Section 8.2.3) that the under-photo and sidebar treatement participants tended to check permissions at the beginning of tasks in the experimental condition, as opposed the control condition, where they tend to check at the end of tasks. We originally theorized, based on the lab study, that mixed participants would check permissions at the end of the task and use the options menu below the album thumbnail to open the permission-modification interface. Instead, what we see is that

participants, regardless of treatment or condition, use the options menu inside an opened album. We anticipate that mixed and under-photo condition participants, are clicking on the album and looking at the proximity display. By the time the album opens they have decided whether there is an error. If they think a permission error might be present, they open the permission-modification interface; if not, they go on to other actions.

8.3 Proximity display designs

We tested five different proximity display designs in the lab and online studies. We use our observations from lab and online studies to better understand why some conditions performed well and others did not.

8.3.1 Under photo

The under-photo condition kept permission information in close spatial proximity to the users' main focus (the photos) but also used the most screen real estate. It is therefore unsurprising that participants viewing this condition corrected statistically significantly more permissions than they did in the control condition. However, we observed in the lab study that seeing the proximity display in so many places caused some participants to start ignoring the display to the point where they could not remember seeing it. We are concerned that if this design were to become commonly used people might become habituated to ignoring it. We hypothesize that habituation is one of the reasons that the under-photo treatment showed a smaller difference between experimental and control than the mixed condition.

8.3.2 Sidebar

The sidebar condition showed no statistically significant difference between conditions in either the eye-tracker study or the online study, but in both studies it was close to significance. The online-study participants corrected 0.22 more permission errors on average when viewing the sidebar condition than when in the control condition. If we consider both eye tracking (Figure 8.1) and observed behavior (Figure 8.5) we see that, unlike the under and mixed conditions, the sidebar condition does not impact when the participants notice or correct permission errors. This condition makes the permission information easier to find, but does not place the display directly in the users' visual path. Consequently, it only helps participants who are looking for permission information. Its advantage over control is that the permissions are easier to find and checking them takes less time and effort. If participants are not looking for the information, or viewing the sidebar for some other reason, they are unlikely to encounter the information. Our results indicate that this condition is less assistive than other conditions in helping users identify errors, but it does appear to give more assistance than control and future work should not dismiss it as ineffective.

8.3.3 Mixed

The mixed condition was designed to combine the best parts of the under and sidebar conditions based on how users interacted with those conditions. Permission information was placed under album thumbnails so it could be noticed as the user entered or closed the album. Under-photo participants in the lab study were observed to primarily check under the album thumbnail, as opposed to under the photo thumbnails. When the album was opened, the proximity display was moved to the sidebar, where it would not interfere with the participants' primary activities. In the lab study we observed that even when participants noticed a permission error mid-way through the task they would wait to the end to correct it, potentially forgetting about the error in the process. Eye-tracker-study participants in the sidebar condition also checked permissions at the end of the task, a behavior which is enabled by keeping permissions on the sidebar when the album is open. Figure 8.5 suggests that this approach worked well, in that participants who saw the displays tended to open the permission-modification interface early in the task. The shift in when the permission-modification inteface was opened may have been because they saw the permission information as they opened the album.

8.3.4 Facebook

The Facebook condition showed no statistically significant difference between conditions in the online study and was not tested in the other studies. Even the average number of permissions corrected for the two conditions was identical. The lack of difference was likely caused by participants failing to notice or comprehend the proximity display icons. Due to the nature of the errors tested, three of the four tasks with errors displayed a  icon, indicating custom permissions. Facebook uses this icon whenever a user allows a set of groups other than Public, Friends, Friends of Friends, or Private to view the album. A single task showed the  icon, which indicates the album is public. The predominance of the  icon may have put the Facebook interface at a disadvantage. However, the task showing the  icon only saw one (0.4% of participants) participant correct the permissions when they saw the Facebook icon as opposed to the tag icon in the control condition. For comparison: the mixed condition, on this task, had six (2.4% of participants) participants correct permissions when they saw permissions on the proximity display. It is likely that participants simply did not understand the meaning of the icons or may not have noticed them at all. Also notable is that 89% of participants in the Facebook treatment had previously used Facebook to share photos and should have been previously exposed to the icons.

8.3.5 Audit

We discussed displaying audit information with focus group participants, piloted several display designs on participants in the eye-tracker and lab studies, and finally conducted a full evaluation in the online study. Our final evaluation showed no statistically significant difference between control and showing audit information. With only 0.007 more permis-

sion errors being corrected when the participant saw audit information on the display, we are confident that displays are not assisting users in identifying errors in this study. Focus group participants voiced concern that audit data would be unhelpful to them. The lab study's pilot participants were never observed to explicitly check permissions using the audit information on the proximity display. When asked why, participants responded that the information was irrelevant to their goals and not helpful. Audit information is primarily intended to give people feedback about how their permission settings are being used so they can re-evaluate their existing permission decisions and adapt the policy over time. Because this was not our participants' actual policy, they had no need or interest in re-evaluating, making the data potentially irrelevant to them. We tried to account for this in the online study by including the names of all groups who could view the albums in the displayed information, but it appears to have not helped. However, we feel that the audit proximity display may be more effective in other domains, or when participants' actual data and policies are used. Future researchers and designers should further explore this display.

8.4 Limitations

While we feel that these studies provide valid results, that can be applied in other contexts and domains, there are some important limitations the reader should consider.

Role playing

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.

We chose to use role play with contrived policies because it ensured all participants had a similar experience and that we knew which albums had errors. However, this choice meant that participants were not previously familiar with the ideal policy, and had no real investment in it. It is possible, even likely, that participants might have behaved differently if given an opportunity to work with their own albums.

Perceived risk could also have been an issue. If a participant failed to protect a study album, no real harm came to the participant. If they did not protect their own albums, there is the potential for actual harm. It is possible that participants might have taken the tasks more seriously, and corrected more permission errors, if the albums had been their own.

The audit condition in the online study was intended to assist users in both identifying errors and reassessing their prior policy decisions. We showed in the online study that the audit condition did not help participants find setting errors. However, with artificial policies, participants could not really adjust or change the policy as it suited them, so we do not know the impact this condition would have had on actual policies. Additionally, this condition displays individual names of people who have accessed the album in the past. While participants in our study were informed of the names of their friends, family, and

co-workers, they may not have internalized the names sufficiently to use them to identify errors.

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.

Photo sharing

Our study used an open-source photo-sharing website software called Gallery and asked participants to conduct photo-manipulation tasks. We selected Gallery because it is relatively unknown to the general populous and could be easily modified. We used version 3.1 which was released in October of 2010 and was a significant departure from the style and user interface of the 2.x versions. We are confident that few, if any, of our participants had prior experience interacting with the 3.1 version of Gallery, guaranteeing all participants received an equal amount of training and experience. However, this also meant that participants were working in an unfamiliar environment. It is possible that, given more time to become familiar with the Gallery interface, participants might have behaved differently.

Our photo-manipulation tasks were designed to simulate the user spending time working with their photo albums. The tasks we chose were selected to be plausible and represent tasks an average user might engage in. However, we made no attempt to accurately replicate a typical online photo-editing experience. Our goal was instead to create a scenario that was sufficiently compelling that participants could easily role play it, and which required the participant to have the album page open for a similar amount of time across tasks. We feel that our role-playing scenarios are a reasonable approximation of the mindset of a user interacting with their online photo albums. However, it is possible that issues such as the length of time spent on the page, or the exact parts of the interface which drew the user's eye, could impact the results of our studies.

Priming

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 overly biasing them towards fixing permissions. In our studies participants were directly informed that permission modification and upkeep were a component of the study. By thus informing participants, we effectively primed them to look for permissions, thereby increasing their likelihood of doing so despite our efforts to present the information in a group with similar, irrelevant, information. We anticipate that if we had not primed participants to look for permission errors, we would have seen a lower number of participants finding and correcting errors. We may also have seen a higher difference in the number of permissions checked between the control and experimental conditions.

8.5 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 and mixed conditions, where access-control information was located under each album thumbnail and under every photo (under), or under the album thumbnail and on the sidebar (mixed), performed statistically significantly better at noticing and fixing errors in albums associated with tasks. We also observed that participants in all conditions tended to change permissions at the beginning and end of tasks, and that some participants were inclined to check all the permissions at once in a single pass. Finally, we observed a high variance between users. Some participants were very inclined to check and correct permissions and others simply forgot about them.

We believe our studies have 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. 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.

Chapter 9

Conclusion

This thesis addresses the issue of helping users to become aware of and better identify errors with their privacy policies. We find that people use a variety of methods to control the security of their resources and data. Many of these methods are predicated on the user being aware both of their current access-control settings, i.e., of what could happen, and of how those settings have been used, i.e., what has happened. To improve the policy awareness and maintenance we proposed the use of proximity displays — small interface components spatially located near the data elements (or near a representation of data, e.g., file name in a file manager or thumbnail photo in a photo album) that contain information about who has or who could access the data. We applied the concept to a photo-sharing website. We then tested the following hypothesis:

> Users of a system that includes proximity information displays of access control-information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface.

To test the hypothesis we conducted focus group studies to gauge user reactions, and empirical evaluations to test the effectiveness of the different proximity displays at improving users' error identification and awareness of policy settings. In the focus groups we found that, for the personal photo domain, users liked the idea of making privacy policy settings appear in close proximity to the photos. However, participants had a strong association between seeing who had viewed photos in the past and stalking behavior. The evaluation studies showed that participants who saw proximity displays with comprehensive permission information that could be easily glanced at were better able to identify access-control policy errors. Participants who saw displays that were overly coarse-grained, on the sidebar, or showed information about who had previously viewed the photos, showed no improvement over users who saw permission settings only on a secondary interface. Our studies suggest that using proximity displays to show access-control settings can significantly help users identify permission errors.

While the proximity displays appear to help people find permission errors they seem to have no effect on permission awareness. We observed no difference between conditions in the number of permissions participants recalled.

The hypothesis is partially supported by our work. People who see proximity displays are more likely to notice errors than when they have to visit a secondary interface to see the permission information. However, we have been able to find no effect on policy awareness. We conclude that proximity displays are a promising approach to error detection in privacy policies, but have minimal impact on awareness, at least in the form we investigated.

9.1 Contributions

The original contributions of this thesis provide a guide for both researchers trying to design lab studies in this space and interface designers wanting to create interfaces that increase policy awareness and enable people to decrease policy errors.

Security professional interviews: Through an interview study, we identified that policy professionals can be categorized into two roles: policy implementer and policy maker. These two roles exist in organizations in which the ability to change the enforced policy is limited to a small set of people. These roles will likely be different in environments such as social networking where the same person is expected to make and implement the policy.

Showing proximity displays assists users in identifying errors: Showing the display under the album/photo thumbnails, or under the album thumbnail and on the album sidebar, appeared to have the most effect on participants. They appeared to check the display as they opened the album or after they had closed the album. Conditions that placed displays under the album thumbnail showed a statistically significant improvement in participants' ability to identify errors.

Providing sufficient information on a proximity display for a user to determine the presence or absence of an error by glancing is important to assist users in identifying errors. People who do not notice an error when they glance at the display are unlikely to dedicate more attention to identifying the error [112]. Placing sufficient information on the proximity display allows participants to more accurately determine if an error exists or not.

Methodology: studying security as a secondary task: Designing a methodology that both keeps participants treating security as a secondary task, while at the same time, imparting the goal permission state, is a challenging problem. Our analysis of the methodological issues we encountered, their causes, and how to overcome them, is a valuable tool for future researchers in this domain.

Understanding user behavior and sentiment: We observed in both our focus group studies and our evaluation studies that some users are more concerned about permission settings than others. In our focus group studies some participants strongly felt that their privacy settings did not really matter because websites would likely loose or expose their photos anyway. This difference caused them to view permission information as unimportant.

We observed in the lab and online studies that people in the control condition check permissions primarily at the end of tasks and rarely the beginning. Participants

seeing proximity displays under the album thumbnail tend to check permissions at the beginning of tasks rather than the end.

Audit information: Showing people information about who has previously seen their personal photos was not well liked by users. Participants felt that making people privy to so much information encouraged stalking. Online study participants did not show an improved ability to notice and correct permissions over control when shown audit information. However, when asked about a document sharing system participants liked the idea of seeing audit information, and considered it a useful component of online document management.

9.2 Future work

Future work falls into three categories. First, developing more effective proximity display designs. Second, understanding what causes people to look for errors in their access-control policies. Third, exploring domains beyond personal photo management.

9.2.1 Proximity display design

In our studies we have explored several designs and spatial placements for proximity displays but we have only taken an initial look at the space of possible designs.

- **Additional privacy settings:** We limited our analysis to showing permission information related to what other people could do, or had done. However, there are many other privacy settings that could be placed on proximity displays. For example, on Facebook a user can control what information is available to their friends when their friends use apps, as opposed to what their friends can see normally.
- **Display designs:** We explored only a small portion of the possible designs for proximity displays. There are many different ways to display privacy information in a way that can be easily glanced at [101]. In particular, we would like to explore the effectiveness of using different styles of icons, and other compact policy representations.
- **Error detection at a glance:** In this and other work [105, 112] we see that people glance at information displays and if they do not detect an issue they assume there is no issue and move on. A proximity display needs to show people enough information that they can accurately identify an error at a glance. If too little information is shown a user may inaccurately decide that there is no error. The question is what data best assists users in identifying errors and how much of it is necessary. The displays we proposed use a non-trivial amount of screen real estate. We would like to know how compact the display can be made before its effectiveness begins to decrease.
- **Display locations:** We showed that placing displays under every photo and album, or under every album thumbnail and on the sidebar, helped participants identify errors. However, we anticipate that with more participants (power) we might see a

significant effect when permissions are displayed just on the sidebar. Additionally, there may be other places and times when showing the proximity display might be more effective.

- **Habituation:** Many of our proximity display designs were unfamiliar to our participants. As we have seen previously, displays that work initially may stop working when participants become habituated to their use [95, 100]. Proximity displays should be tested for an extended time period to detect issues with habituation.

9.2.2 Understanding policy error identification behavior

We have shown that participants tend to check permissions at the end of tasks, and that exposing them to proximity displays causes them to check at the beginning of tasks more frequently. We also observed that some people are more inclined to check permissions for errors than other people. Post-study survey answers, and information from the focus groups, suggest that peoples' assumptions about whether privacy settings on websites will be effective at protecting their content, impacts their permission checking behavior. We would like to further explore this observation and determine if people's mental models of website behavior really do impact permission checking.

9.2.3 Exploring proximity displays in other domains

In this work we looked at proximity displays in the domain of online photo sharing. However, we believe that proximity displays could be effective in helping end users manage their access-control policies in a variety of domains.

Social networking A clear extension to this work is to test proximity displays in a social networking site context. Social networking sites, such as Facebook, are creating increasingly more complex privacy policies which users can configure. In future work we would like to explore how these settings could be incorporated into the proximity display design.

There is also the issue of awareness: while our studies did not indicate that proximity displays improved participants' ability to remember permissions, they did make it easier for participants to find, and check, permissions when they were interested. Placing setting information on the proximity display may improve peoples' understanding of what settings are available to be manipulated. In addition to not being aware of their own permission settings people are not always aware of all the settings that are available [59]. For example, a user may not attempt to opt-out of marketing data being sold if they are not aware that opting out is an option.

Document sharing Managing document sharing in an organization is an issue corporate IT departments are struggling with [105]. Documents are easy to create and share and if the corporate document sharing system has too unusable or restrictive of an interface people resort to email and USB drives to share documents. While convenient, these technologies

are less secure and are more likely to be lost or compromised than the company's servers. Additionally, as the internal structure of an organization changes, the access-control policy does not always change with it leaving people with too much or too little access than is necessary to do their jobs. Proximity displays could help people keep up with the changes, by helping them identify permission errors and enabling them to easily determine who can see each document.

Healthcare The domain of healthcare is interesting in that emergency personnel need immediate access to health care records for safety reasons, and the data in medical files is generally considered privacy sensitive. Proximity displays could be used to help health care professionals maintain security on the files through *ex-post* control, allowing anyone access to any file, but also showing that access attempt to anyone else interacting with the file.

Chapter 10

Bibliography

[1] Gallery 3. Accessed on: July 2012, http://gallery.menalto.com/. 6.1.1

[2] Generation Y online security survey. Technical report, TRU Research, 2010. 4.3.1

[3] Facebook settles FTC charges that it deceived consumers by failing to keep privacy promises, November 2011. Accessed on: July 2012, http://ftc.gov/opa/2011/11/privacysettlement.shtm. 4.3.1

[4] Martin Abadi. On SDSI's linked local name spaces. *Journal of Computer Security*, 6:3–21, 1998. 3.3.5

[5] Anne Adams and Martina Angela Sasse. Users are not the enemy. In *Communications of the ACM*, volume 42, pages 40 – 46, 1999. 2.2.4

[6] Shane Ahern, Dean Eckles, Nathaniel S. Good, Simon King, Mor Naaman, and Rahul Nair. Over-exposed?: Privacy patterns and considerations in online and mobile photo sharing. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2007. 2.2.1

[7] Andrew W. Appel and Edward W. Felten. Proof-carrying authentication. In *Proceedings of the ACM conference on Computer and Communications Security*, 1999. 3.3.5

[8] Dawn M. Appelli, Akash G. Desai, Andrew P. Moore, Timothy J. Shimeall, Elise A. Weaver, and Bradford J. Willke. Management and education of the risk of insider threat (MERIT): Mitigating the risk of sabotage to employers' information, systems, or networks. Technical Report CMU/SEI-2006-TN-041, CERT, Software Engineering Institute at Carnegie Mellon University and CyLab, 2007. 2.2.4, 3

[9] Dixie B. Baker. Fortresses built upon sand. In *Proceedings of the workshop on New security paradigms*, 1996. 2.2.3

[10] Rob Barrett, Eser Kandogan, Paul P. Maglio, Eben Haber, Leila A. Takayama, and Madhu Prabaker. Field studies of computer system administrators analysis of system management tools and practices. In *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 2004. 3.6

[11] Lujo Bauer, Scott Garriss, and Michael K. Reiter. Distributed proving in access-control systems. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2005. 2.3, 3.1.1, 3.6

[12] Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, and Kami Vaniea. Lessons

learned from the deployment of a smartphone-based access-control system. In *Proceedings of the Symposium on Usable Privacy and Security*, 2007. 2.2.6, 2.3, 2.3, 2.3, 3.7

[13] Lujo Bauer, Lorrie Faith Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. A user study of policy creation in a flexible access-control system. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2008. 2.2.3, 2.3, 2.3, 3.4.4, 3.5.2, 3.7, 4.1, 6.1, 6.5.1

[14] Lujo Bauer, Lorrie Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. Real life challenges in access-control management. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2009. 2.2.6

[15] Lujo Bauer, Lorrie F. Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. Effects of access-control policy conflict-resolution methods on policy-authoring usability. Technical Report CMU-Cylab-09-006, Cylab, Carnegie Mellon University, March 2009. 2.3, 5.3.3

[16] Allan Beaufour and Philippe Bonnet. Personal servers as digital keys. In *Proceedings of the IEEE International conference of Pervasive Computing and Communications*, 2004. 2.3, 3.1.1

[17] Victoria Bellotti and Abigail Sellen. Design for privacy in ubiquitous computing environments. In *Proceedings of the conference on European conference on Computer-Supported Cooperative Work*, 1993. 2.2.2

[18] Michael Benisch, Patrick Kelley, Norman Sadeh, and Lorrie Faith Cranor. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Journal of Personal and Ubiquitous Computing*, pages 1–16, 2010. 2.2.6

[19] Andrew Besmer and Heather Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2010. 1, 2.2.1

[20] Hugh Beyer and Karen Holtzblatt. *Contextual Design: Defining customer-centered systems*. Morgan Kaufmann Publishers, 1998. 3.1.4, 4.3, 7.2.4

[21] Bob Blakley. The Emperor's old armor. In *Proceedings of the workshop on New security paradigms*, 1996. 2.3

[22] David Botta, Rodrigo Werlinger, André Gagné, Konstantin Beznosov, Lee Iverson, Sidney Fels, and Brian Fisher. Towards understanding IT security professionals and their tools. In *Proceedings of the Symposium on Usable Privacy and Security*, 2007. 3.1.4, 3.6

[23] Carolyn A. Brodie, Clare-Marie Karat, and John Karat. An empirical study of natural language parsing of privacy policy rules using the sparcle policy workbench. In *Proceedings of the Symposium on Usable Privacy and Security*, 2006. 2.3, 3.6

[24] Achim D. Brucker and Helmut Petritsch. Extending access control models with break-glass. In *Proceedings of the ACM symposium on Access Control Models and Technologies*, 2009. 2.2.6

[25] A.J. Brush and Kori Inkpen. Yours, mine and ours? Sharing and use of technology in domestic environments. In *Proceedings of the international conference on Ubiquitous computing*, 2007. 2.2.3

[26] L. Jean Camp, Cathleen McGrath, and Alla Genkina. Security and morality: A tale

of user deceit. *Models of Trust for the Web*, 22, 2006. 4.3.1

[27] Sunny Consolvo, Ian E. Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. Location disclosure to social relations: Why, when, & what people want to share. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2005. 2.2.2, 6.1, 6.5.1

[28] Fergus I.M. Craik and Robert S. Lockhart. Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 11:671–684, 1972. 8.2.2

[29] Lorrie Faith Cranor. Privacy policies and privacy preferences. In L. F. Cranor and S. Garfinkel, editors, *Privacy and Usability*. O'Reilly, 2005. 1

[30] Lorrie Faith Cranor. A framework for reasoning about the human in the loop. In *Proceedings of the conference on Usability, Psychology, and Security*, 2008. 2.4, 2.4, 8.2.1

[31] Justin Cranshaw, Jonathan Mugan, and Norman Sadeh. User-controllable learning of location privacy policies with gaussian mixture models. In *Proceedings of the AAAI conference on Artificial Intelligence*, 2011. 2.2.6

[32] Brinda Dalal, Les Nelson, Diana Smetters, and Nathaniel Good. Ad-hoc guesting: When exceptions are the rule. In *Proceedings of Usability, Psychology, and Security*, 2008. 2.2.4, 2.2.6

[33] Rogério de Paula, Xianghua Ding, Paul Dourish, Kari Nies, Ben Pillet, David Redmiles, Jie Ren, Jennifer Rode, and Roberto Silva Filho. Two experiences designing for effective security. In *Proceedings of the Symposium on Usable privacy and security*, 2005. 2.3

[34] Paul Dourish, E. Grinter, Jessica Delgado de la Flor, and Melissa Joseph. Security in the wild: user strategies for managing security as an everyday, practical problem. *Journal of Personal Ubiquitous Computing*, 8:391–401, 2004. ISSN 1617-4909. 1, 2.2.2, 2.2.3, 3.6

[35] W. Keith Edwards, Erika Shehan Poole, and Jennifer Stoll. Security automation considered harmful. In *Proceedings of the IEEE New Security Paradigms workshop*, 2007. 2.2.6

[36] Serge Egelman. *Trust Me: Design Patterns for Constructing Trustworthy Trust Indicators*. PhD thesis, Carnegie Mellon University, 2009. CMU-ISR-09-110. 1.1, 2.1

[37] Serge Egelman, A.J. Brush, and Kori Inkpen. Family accounts: A new paradigm for user accounts within the home environment. In *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 2008. 2.2.3

[38] Serge Egelman, Andrew Oates, and Shriram Krishnamurthi. Oops, i did it again: Mitigating repeated access control errors on facebook. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2011. 6.2

[39] Willis D. Ellis. *A Source Book of Gestalt Psychology*. Psychology Press, 1999. 5.2.1

[40] David Ferraiolo, Janet A. Cugini, and D. Richard Kuhn. Role-based access control (rbac): Features and motivations. In *Proceedings of Annual Computer Security Application conference*, 1995. 2.2.4, 2.3

[41] David F. Ferraiolo, Dennis M. Gilbert, and Nickilyn Lynch. An examination of federal and commercial access control policy needs. In *National Computer Security*

conference, 1993. 2.2.4

[42] Shirley Gaw, Edward W. Felten, and Patricia Fernandez-Kelly. Secrecy, flagging, and paranoia: adoption criteria in encrypted email. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2006. 2.2.5, 3.6, 4.3.1

[43] Virgil Gligor. Characteristics of role-based access control. In *Proceedings of the ACM workshop on Role-based access control*, 1996. 2.3

[44] D Godden and A.D Baddeley. Context-dependent memory in two natural experiments: on land and under water. *British Journal of Psychology*, 66:325–331, 1975. 6.6.3, 7.2.2

[45] Grant Gross. Former gov't worker sentenced for passport snooping, March 2009. Accessed on: August 5th 2009, http://www.networkworld.com/news/2009/032309-former-govt-worker-sentenced-for.html. 2.2.4

[46] Joshua B. Gross and Mary Beth Rosson. Looking for trouble: understanding end-user security management. In *Proceedings of the Symposium on Computer Human Interaction for the Management of Information Technology*, 2007. 2.2.4

[47] Alejandro Gutierrez, Apeksha Godiyal, Matt Stockton, Michael LeMay, Carl A. Gunter, and Roy H. Campbell. Sh@re: negotiated audit in social networks. In *Proceedings of the IEEE International conference on Systems, Man and Cybernetics*, 2009. 2.3

[48] Joerg M. Haake, Anja Haake, Till Schümmer, Mohamed Bourimi, and Britta Landgraf. End-user controlled group formation and access rights management in a shared workspace system. In *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 2004. 6

[49] Cormac Herley. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *Proceedings of New Security Paradigms workshop*, 2009. 2.2.4

[50] Airlie House and Virginia Warrenton. CRA conference on grand research challenges in information security & assurance, November 2003. Accessed on: July 2012, http://www.cra.org/Activities/grand.challenges/security/home.html. 2

[51] Elaine M. Huang and Khai N. Truong. Breaking the disposable technology paradigm: opportunities for sustainable interaction design for mobile phones. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2008. 3.1.4

[52] Trent Jaeger, Antony Edwards, and Xiaolan Zhang. Managing access control policies using access control spaces. In *Proceedings of the ACM symposium on Access control models and technologies*, 2002. 2.3, 2.3

[53] Lukasz Jedrzejczyk, Blaine A. Price, Arosha K. Bandara, and Bashar Nuseibeh. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. In *Proceedings of the Symposium on Usable Privacy and Security*, 2010. 2.2.2, 2.3

[54] Xin Jin, Ram Krishnan, and Ravi Sandhu. A unified attribute-based access control model covering dac, mac and rbac. In *Proceedings of the Annual WG 11.3 conference on Data and Applications SEcurity and Privacy*, 2012. 2.2.6

[55] Maritza L. Johnson, Steven M. Bellovin, Robert W. Reeder, and Stuart E. Schechter. Laissez-faire file sharing: Access control designed for individuals at the endpoints.

In *Proceedings of the workshop on New security paradigms workshop*, 2009. 2.3

[56] Sara Kaemer and Pascale Carayon. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. In *Applied Ergonomics*, 2007. 2.2.4

[57] Clare-Marie Karat, John Karat, Carolyn Brodie, and Jinjuan Feng. Evaluating interfaces for privacy policy rule authoring. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2006. 2.3

[58] Timothy Kelly, Suzanne Lien, L. Jean Camp, and Douglas Stebila. Self-identified experts lost on the interwebs. In *Proceedings of the Learning from Authoritative Security Experiment Results (to appear)*, 2012. 2.1

[59] Jennifer King, Airi Lampinen, and Alex Smolen. Privacy: Is there an app for that? In *Proceedings of the Symposium on Usable Privacy and Security*, 2011. 2.2.1, 9.2.3

[60] Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Lorrie Faith Cranor, Nitin Gupta, and Michael Reiter. Tag, you can see it!: using tags for access control in photo sharing. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2012. 2.2.2, 2.2.6

[61] Brian Krebs. Court rules against teacher in myspace 'drunken pirate' case, December 2008. Accessed on: July 2012, http://voices.washingtonpost.com/securityfix/2008/12/court_rules_against_teacher_in.html. 1

[62] B.W. Lampson. Protection. *ACM SIGOPS Operating Systems Review*, 8(1):18–24, 1974. 2.3

[63] Barry Leibowitz. Facebook blunder invites 15,000 to teen's 16th birthday party; 100 cops show up, too, June 2011. Accessed on: July 2012, http://www.cbsnews.com/8301-504083_162-20069457-504083.html. 1

[64] Ninghui Li and John C. Mitchell. Understanding SPKI/SDSI using first-order logic. *International Journal of Information Security*, 2004. 3.3.5

[65] Eric Lieberman and Robert C Miller. Facemail: showing faces of recipients to prevent misdirected email. In *Proceedings of the Symposium on Usable Privacy and Security*, 2007. 2.1

[66] Linda Little, Elizabeth Sillence, and Pam Briggs. Ubiquitous systems and the family: thoughts about the networked home. In *Proceedings of the Symposium on Usable Privacy and Security*, 2009. 2.2.3

[67] Mary Madden and Aaron Smith. Reputation management and social media. Technical report, Pew Internet & American Life Project, 2010. 2.2.2

[68] Michelle Madejski, Maritza Johnson, and Steven M. Bellovin. The failure of online social network privacy settings. Technical Report CUCS-010-11, Department of Computer Science, Columbia University, 2011. 1, 2.2.1, 4.1

[69] Roy A. Maxion and Robert W. Reeder. Improving user-interface dependability through mitigation of human error. *Int. J. Hum.-Comput. Stud.*, 63:25–50, 2005. 2.3, 5.3.3

[70] Alain Mayer, Avishai Wool, and Elisha Ziskind. Fang: A firewall analysis engine. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2000. 2.3

[71] Michelle L. Mazurek, J.P. Arsenault, Joanna Breese, Nitin Gupta, Iulia Ion, Christina

Johns, Daniel Lee, Yuan Liang, Jenny Olsen, Brandon Salmon, Richard Shay, Kami Vaniea, Lujo Bauer, Lorrie Faith Cranor, Gregory R. Ganger, and Michael K. Reiter. Access control for home data sharing: Attitudes, needs and practices. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2010. 2.2.2, 2.2.3, 3.7

[72] Michelle L. Mazurek, Peter F. Klemperer, Richard Shay, Hassan Takabi, Lujo Bauer, and Lorrie Faith Cranor. Exploring reactive access control. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2011. 2.2.2, 2.3, 6.1, 6.5.1

[73] Sylvie Noël and Jean-Marc Robert. Empirical study on collaborative writing: What do co-authors do, use, and like? *Comput. Supported Coop. Work*, 13:63–89, 2004. 2.2.4

[74] D.A. Norman. *The design of everyday things*. Basic Books New York, 2002. 2.3

[75] National Academy of Engeineering. Grand challenges for engineering: Secure cyberspace. Accessed on: July 2012, http://www.engineeringchallenges.org/cms/8996/9042.aspx. 2

[76] Antti Oulasvirta. Finding meaningful uses for context-aware technologies: The humanistic research strategy. In *Proceedings of Computer Human Interaction*, 2004. 2.2.2

[77] Sameer Patil and Jennifer Lai. Who gets to know what when: configuring privacy permissions in an awareness application. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2005. 2.2.2

[78] Erika Shehan Poole, Marshini Chetty, Tom Morgan, Rebecca E. Grinter, and W. Keith Edwards. Computer help at home: methods and motivations for informal technical support. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2009. 2.2.3

[79] Dean Povey. Optimistic security: A new access control paradigm. In *Proceedings of New Security Paradigms workshop*, 1999. 2.2.6, 2.3, 2.3

[80] E. Rader, R. Wash, and B. Brooks. Stories as informal lessons about security. In *Proceedings of the Symposium on Usable Privacy and Security*, 2012. 2.2.2, 2.2.3

[81] Marisa R. Randazzo, Michelle Keeney, Eileen Kowalski, Dawn Cappelli, and Andrew Moore. Insider thread study: Illicit cyber activity in the banking and finance sector. Technical report, Carnegie Mellon University Software Engineering Institute, 2005. 3

[82] Robert W. Reeder, Clare-Marie Karat, John Karat, and Carolyn Brodie. Usability challenges in security and privacy policy-authoring interfaces. In *Human-Computer Interaction*, 2007. 2.3

[83] Robert W. Reeder, Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, Kelli Bacon, Keisha How, and Heather Strong. Expandable grids for visualizing and authoring computer security policies. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2008. 2.3, 3.6, 4.3.4, 5.2.1, 5.3.1, 6.1.1

[84] Robert W. Reeder, Lujo Bauer Bauer, Lorrie Faith Cranor, Michael K. Reiter, and Kami Vaniea. More than skin deep: Measuring effects of the underlying model on access-control system usability. In *Proceedings of the annual SIGCHI conference on*

Human factors in computing systems, 2011. 5.2.1, 6.1.1

[85] E. Rissanen, B. Firozabadi, and M. Sergot. Towards a mechanism for discretionary overriding of access control. In *Security Protocols*, 2002. 2.2.6, 2.3, 2.3

[86] Norman Sadeh, Jason Hong, Lorrie Faith Cranor, Ian Fette, Pattrick Gage Kelley, Maduh Prabaker, and Jinghai Rao. Understanding and capturing people's privacy policies in a mobile social networking application. *Journal of Personal and Ubiquitous Computing*, 13(6), 2009. 2.2.6

[87] Brandon Salmon, Frank Hady, and Jay Melican. Learning to share: A study of sharing among home storage devices. Technical Report CMU-PDL-07-107, Carnegie Mellon University Parallel Data Lab, October 2007. 2.2.3

[88] Brandon Salmon, Steven W. Schlosser, Lorrie Faith Cranor, and Gregory R. Ganger. Perspective semantic data management for the home. In *Proceedings of the USENIX conference on File Storage Technolgies*, 2009. 2.2.2, 2.2.3

[89] Jerome H. Saltzer and Michael D. Schroeder. The protection of information in computer systems. *IEEE, Proceedings*, 63:1278–1308, 1975. 2.3, 3.4.4

[90] Bruce Schneier and Marcus Ranum. Schneier-ranum face-off: Is perfect access-control possible?, September 2009. Accessed on: July 2012, http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1365957,00.html. 2.2.4

[91] Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs. Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2010. 6.2

[92] Sara Sinclaire, Sean W. Smith, Stephanie Trudeau, M. Eric Johnson, and Anthony Portera. Information risk in the professional services - field study results from financial institutions and a roadmap for research. Technical report, Dartmouth College, 2007. 2.2.5

[93] Supriya Singh, Anuja Cabraal, and Gabriele Hermansson. What is your husband's name?: sociological dimensions of internet banking authentication. In *Proceedings of the Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments*, 2006. 3.6

[94] D. K. Smetters and Nathan Good. How users use access control. In *Proceedings of the Symposium on Usable Privacy and Security*, 2009. 2.2.5

[95] Andreas Sotirakopoulos, Kirstie Hawkey, and Konstantin Beznosov. On the challenges in usable security lab studies: Lessons learned from replicating a study on ssl warnings. In *Proceedings of the Symposium on Usable Privacy and Security*, 2011. 2.1, 6.1, 9.2.1

[96] Hanna Stelmaszewska, Bob Fields, and Ann Blandford. The roles of time, place, value and relationships in collocated photo sharing with camera phones. In *Proceedings of the British HCI Group Annual conference on People and Computers: Culture, Creativity, Interaction - Volume 1*, 2008. 2.2.2

[97] Gunnar Stevens and Volker Wulf. A new dimension in access control: Studying maintenance engineering across organizational boundaries. In *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 2002. 2.2.6

[98] Gunnar Stevens and Volker Wulf. Computer-supported access control. *ACM Trans. Comput.-Hum. Interact.*, 16(3):1–26, 2009. 2.3, 2.3, 2.3

[99] Oliver Stiemerling and Vulker Wulf. Beyond" Yes or No"-Extending Access Control in Groupware with Awareness and Negotiation. *Group Decision and Negotiation*, 9 (3):221–235, 2000. 2.2.6, 2.3

[100] Joshua Sunshine, Serge Egelman, Hazim Almuhimedi, Neha Atri, and Lorrie Faith Cranor. Crying wolf: An empirical study of SSL warning effectiveness. In *Proceedings of the conference on USENIX security symposium*, 2009. 2.1, 6, 6.1, 9.2.1

[101] Jennifer Tam, Robert W. Reeder, and Stuart Schechter. I'm allowing what? disclosing the authority applications demand of users as a condition of installation. Technical Report MSR-TR-2010-54, Microsoft, May 2010. 5.2.1, 5.2.2, 9.2.1

[102] Janice Y. Tsai. *The impact of salient Privacy information on decison-making*. PhD thesis, Carnegie Mellon University, 2009. 12-1-2009. 2.1

[103] Janice Y. Tsai, Serge Egelman, Lorrie Faith Cranor, and Alessandro Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Journal of Information Systems Research*, 22(2):254–268, 2011. 2.1, 8.2.2

[104] Maarten W. van Someren, Yvonne F. Barnard, and Jacobijn A.C. Sandberg. *The Think Aloud Method: A practical guide to modelling cognitive processes*. Academic Press, 1994. 6.6.3, 8.2.1

[105] Kami Vaniea, Clare-Marie Karat, Joshua B. Gross, John Karat, and Carolyn Brodie. Evaluating assistance of natural language policy authoring. In *Proceedings of the Symposium on Usable Privacy and Security*, 2008. 2.3, 8.2.2, 9.2.1, 9.2.3

[106] Kami Vaniea., Lujo Bauer., Lorrie Faith Cranor, M. K. Reiter, and Mike K. Reiter. Out of sight, out of mind: Effects of displaying access-control information near the item it controls. In *Proceedings of Privacy Security and Trust*, 2012. 6.2, 7, 8.2.3

[107] Yang Wang. *A Framework for Privacy-Enhanced Personalization*. Ph.D. dissertation, University of California, Irvine, 2010. 2.1, 6

[108] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. I regretted the minute i pressed share?: A qualitative study of regrets on facebook. In *Proceedings of the Symposium on Usable Privacy and Security*, 2011. 1, 8.2.2

[109] Ryan West. The psychology of security. *Communications of the ACM*, 51:34–40, 2008. ISSN 0001-0782. 6.1, 6.7

[110] Tara Whalen, Diana Smetters, and Elizabeth F. Churchill. User experiences with sharing and access control. In *Proceedings of the extended abstracts on Human Factors in Computing Systems*, 2006. 1, 2.2.1, 2.2.4, 2.3

[111] Alma Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proceedings of USENIX Security Symposium*, 1999. 6.1, 6.2

[112] Michael S. Wogalter. *Communication-Human Information processing (C-HIP) Model*, pages 51–61. Lawrence Erlbaum Associates, 2006. 2.4, 5.1.2, 8.2.1, 9.1, 9.2.1

[113] Allison Woodruff, Sally Augustin, and Brooke Foucault. Sabbath day home automation: "it's like mixing technology and religion". In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2007. 3.1.4

[114] Volker Wulf and Bjorn Golombek. Direct activation: A concept to encourage tailoring activities. *Journal of Behaviour and Information Technology*, 20(4):249–263, 2001. 2.3

[115] Xia Zhao and M. Erik Johnson. Information governance: Flexibility and control through escalation and incentives. In *Proceedings of the workshop on the Economics of Information Security, Dartmouth College, June*, 2008. 2.3

Appendix A

Focus group study

A.1 Focus Group Script

Welcome. We want to thank you for your participation in our focus group on how people share pictures and other documents using the internet. My name is Kami and this is Veda - we are both students at Carnegie Mellon University. I will be moderating today and Veda will be taking notes.

This session will be audio recorded for latter review. Please try and stay on topic during the discussion and try not to say anything that you wouldn't want to be recorded. The topics we will be discussing today should not be of a sensitive nature. However, if at any time you want to say something that you do not want to be recorded please just let me know and I will temporarily turn off the audio recorder.

Your opinions are very important to us, and we want you to feel free to tell us exactly what you think - and we hope, that your ideas will create discussion.

Today we will be talking about sharing documents and pictures online using websites such as Flicker, Facebook, Picasa, YouTube or Myspace. (Icebreaker) To start I want everyone to tell us your first name and a web site you use to share information such as pictures. I'll start, I use a photo sharing software called Gallery to share picutures with friends and co-workers.

I'd like to continue this session with a discussion about your past experiences with sharing electronic files like photographs, music, videos and documents with other people using the computer. I'd like to go around the table again and have everyone tell us about the last time you posted a file to an online sharing site. When I say "file" I mean anything from a Microsoft Word document to a photograph. What were you sharing? Who were you sharing it with and why did you chose that particular way to share it?

If anyone says something interesting ask a question but this section should have limited conversation. Prompts

- *Why did you choose that web site?*
- *I'm less interested in Facebook posts and more interested in Photographs, video or documents such as Word documents.*

Thank you. Now that we have all heard about how the other people at this table share information with people they want to share with. Can anyone tell me about an experience where you discovered that someone you didn't want to see your shared files either could or did see them?

(If no one answers: How likely is it that your shared files can be seen by someone who you don't want to see them?)

Prompts:

- *How did you find out that they saw your files?*
- *Were you able to solve the problem?*
- *Why were they able to see your files?*
- *How did you feel about that person seeing your files?*
- *Do you still feel comfortable sharing files online?*
- *Did you alter how you post files online. For example did you choose to not post some files because of this experience or did you change your privacy settings?*
- *How are you preventing this from happening in the future?*
- *Does anyone else want to share a different experience where sharing files online didn't go as you expected?*

Has anyone had the opposite problem where you tried to share a file with someone and they couldn't see the file?

Prompts:

- *I'm less interested in email and technical issues and more interested in situations where your settings prevented them from seeing the file. For example if I shared pictures on Facebook and I only wanted my friends to see it not my Mom so I only shared with friends and latter realized that my sister, who I wanted to see the pictures couldn't see them.*
- *How did you find out that they couldn't see your files?*
- *Were you able to solve the problem?*
- *Why couldn't they see your files?*
- *How are you preventing this from happening in the future?*

Hand out comics.

Now I would like to move on. You talked about sharing information using *[insert example from prior conversation]*. Now imagine a photo sharing web site had a feature where you could see who has been looking at your shared photos and who could look at your photos. I've handed you comics about two people named Alice and Joe who use a web site like this. Please read their stories.

Can you imagine an instance where you or a friend might experience a situation like those Alice and Joe encountered?

Prompts:

- *Can you see yourself or a friend using information about who has seen your pictures to reconnect with a friend?*
- *Can you see yourself or a friend using information about who could see your pictures to identify people who can see your pictures but shouldn't?*

I'm now going to give each of you a packet with some example photo sharing websites.

We are going to go through each page of the packet together so please do not look ahead in the packet.

Hand out packets

Please open the packet like this (*demonstrate opening so both the websites are visible.*) so you can see two pages at once.

The first two pages of the packet are screen shots of a potential photo sharing web site that lets you organize and share your photos. I'd like you to imagine that this is your favorite photo sharing website and that it already has all the features you are used to seeing. If I click on one of the photo albums it will open and show the pictures inside.

It has a feature where the owner of a photo album can see information about who has and and who could see their photographs. We are going to use the projector to show you how this website might work. Your comments and opinions are extremely important to us so feel free to write on any of the pages in the packet including the pictures. I'm going to collect the packets at the end so if there is anything you thought was important but didn't get to say please write it down.

In this first example (*describe the interface*)

Allow participants to make comments at this point. If they ask questions about things covered in the written description answer them, if not ask the participant what they think it would look like or what they think it should do.

On the next page there are several questions about this website. Its important to remember we are testing our sample website designs, our vocabulary and layout choices not you.

The questions on this page are designed to represent several different questions people might try to answer if they had information about who could see their pictures and who has seen their pictures. They are supposed to assist you in understanding how you might use this webpage so you can give more informed opinions about it as well as compare it to the other websites I will be showing you. Not all the questions can be answered and some have ambiguous answers. If you feel that it is impossible to answer a question just write down that it can't be answered. We are testing the webpage layout not you. There are no wrong answers to these questions. Also, if anything seems particularly confusing about the website design I would like you to circle it so we can discuss it latter.

Do you have any questions?

Please try and answer the questions on your own right now.

Wait for the majority to answer the questions

I'd like to move on to a discussion of this website design now. Its all right if you haven't finished answering the questions. Feel free to write any additional comments you have during the discussion. After interacting with this interface do you think it is something you would like to use as part of your favorite photot sharing website?

[use prompts below]

For each pair of information display pages in the packet repeat the following script.

Please turn to the next page. (*describe the interface*)

Allow participants to make comments at this point. If they ask questions about things covered in the written description answer them, if not ask the participant what they think it would look like or what they think it should do.

Please look at the website and answer the questions in the provided space on the second page. Feel free to draw on the webpage screenshot and point out anything you think is confusing. I'll give you a few minutes to do so.

Wait for the majority to answer the questions

Now that everyone has looked at the website can someone start us out by saying what they think the best and worst thing about this website is?

Prompts:

- *Would this website be useful for Alice?*
- *Would this website be useful for Joe?*
- *What do you think the feature shown in this webpage would be useful for?*
- *Was any of the language on this page confusing?*
- *If you saw this information display next week how confident are you that you could use it?*
- *What did you like or find confusing?*
- *Which was more useful in this interface: who could see the pictures or who did see the pictures?*
- *If you could change the way the website looks, what would you change?*

After going through the whole packet

Now that you have seen several different ways of showing information about who has and who could see pictures in an online photo album, I'd like to go around the table and have each person say what their favorite and least favorite website was and why.

Prompts:

- *Of the different types of information you have seen presented today which do you find to be the most useful?*
- *Are you more interested in who looked, what was looked at or how often it happened?*

I would like to thank everyone for coming. Please leave your packets on the table.

A.1.1 Information visualization explanations

Website A Information about who has and who could see each of these albums is listed below the album name. For example Alex, Jane and four other people, who's names are not listed, have seen Halloween 2009 photos. A total of six people have the ability to view the album.

- Would you prefer to see who the "potential viewers are"?
- What else might you want to find out about your photo use that this application isn't showing you?

Website B This webpage shows information about who has and who could view the albums shown on this page as well as anything inside those albums.

On the left side of this webpage there is a grid of people across the top and albums down the left side. The colors in the grid indicate if that person can see that album, green means they can see anything in the album, yellow means they can some pictures in the album but not all and red means they can't see anything.

The numbers indicate how often they have looked at the album. At the bottom left there is a small bar graph showing how often people have looked at any of the albums

over a long time. The small window is showing the time period where the numbers are coming from.

For example if I were to look at Nicole I can see that she can see the Niagara Falls pictures and that she has looked at one picture in the album in the last three months. If I select Nicole the albums she could see are all highlighted. The highlight color indicates how often Nichole looked at the pictures in that album. Dark blue means the most and light blue means the least.

- Do the colored frames around the pictures make sense?
- Can anyone tell me what the bar graph in the bottom left means?
- Why is "Around Pittsburgh" colored yellow? What does that mean?

Website C On the top of this webpage there is a list of people and a graphic showing information about who could view these albums. The list of people on the left shows who has been looking at pictures. People above the dotted line have looked at some of Alice's pictures in the last month.

Albums at this website can contain other albums inside of them. For example "Around Pittsburgh" may contain another album called "The Strip." The graphic shows all the albums including some of the albums inside of other albums.

The graphic also shows what albums the highlighted person has or could see. If an album is green than the highlighted person can see anything in that album. If the color is yellow then that person can see some of the pictures in that album and red indicates that the person can see nothing in that album.

The bigger the rectangle that represents the album the more times that person has looked at that album. If I were to click on one of the names the graphic would change to show what albums that user has and could see.

- This website shows you the policies of every album and subalbum that instead of just the albums in this folder. Is this useful to you?
- The list on the left shows at a glance who has been recently looking at photographs. Is a name with no context sufficient to understand what is going on.

Website D In this website information about who can and has seen pictures in any of the albums is shown on the left. There is a list of people in this box. On the left of each person's name is a colored box, if it is green they can see any of the albums, yellow means they can see some of the albums and red means they can't see any albums.

The small graph to the right of the person's name shows when they saw pictures. The start and end dates for this graph are indicated by the labels on the top. In this case they go from January to April.

- Would you think to click on the names on the left to determine what they looked at?
- Is it clear how long the graph next to the names is for?
- Is it easy to understand the re-sizing of the images?

- Is it easy to understand why some of the album pictures are greyed out?

Website E Information about who has and who could see each of these albums is listed below the album name. For example the Phipps photos were seen by 6 people and could be seen by Alice, Kate and 6 other people. The small graph indicates when the album was viewed over the last month.

- Can anyone tell me one person who recently viewed the Phipps photos? Is it clear that the names are people who could view not people who have viewed?
- Do the small graphs make sense?
- Do you think you would casually look at this information when viewing your online photo albums?

Website F On the left of this application are several sections each labeled with a person's name. Below each label are several albums that person has seen over the last month. The bigger the name the more often they saw the album. Black albums have been seen recently and they fade to grey as time passes. After a month they completely disappear.

- Is a month long enough?
- Is it clear what the names of the albums are in the information display?
- Would you expect to see names of subalbums here?

Website G On the left of this application is a list of the people who have and can see any of the albums. "Who has seen my pictures" is ordered starting with the person who most recently viewed an album. Next to each name is the last time they saw a picture and how long they looked at the pictures on that occasion. Below is a list of all the people who can see at least one picture in these albums.

- Is the length of time they looked at your albums insteresting?
- Is the list of who could see pictures intersting even though you don't know what they can see?

A.2 Focus group 1 packet

Alice is looking through her online photographs for ones she can use for her screensaver.

Alice is delighted to discover that her good friend Sue found time to look at the pictures of their trip to Chicago together.

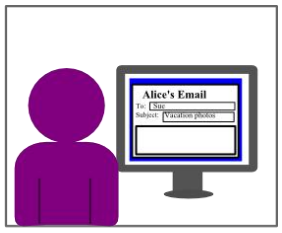

Alice hasn't had a chance to talk to Sue much since they got back so she decides this is the perfect time to reconnect. Alice writes Sue an email saying how much fun she had and she can't wait to see Sue again at Christmas.

Joe is posting pictures from his vacation on his favorite photo sharing website. He is looking through the pictures to make sure they are all ok when he notices that his ex-girlfriend, Erin, can still see his vacation pictures.

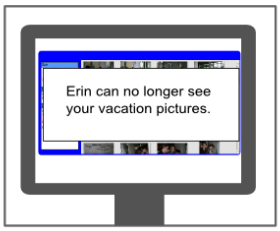

Joe is very upset and quickly removes Erin from the list of people who can see his vacation pictures.

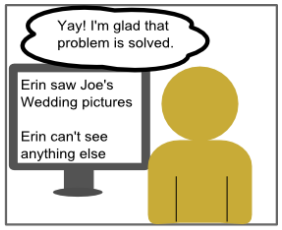

Joe is very concerned that Erin may have seen some pictures he doesn't want her to see. So he looks through the remainder of his photo album to make sure she can't see anything else and to check that she hasn't seen anything she shouldn't.

Website A

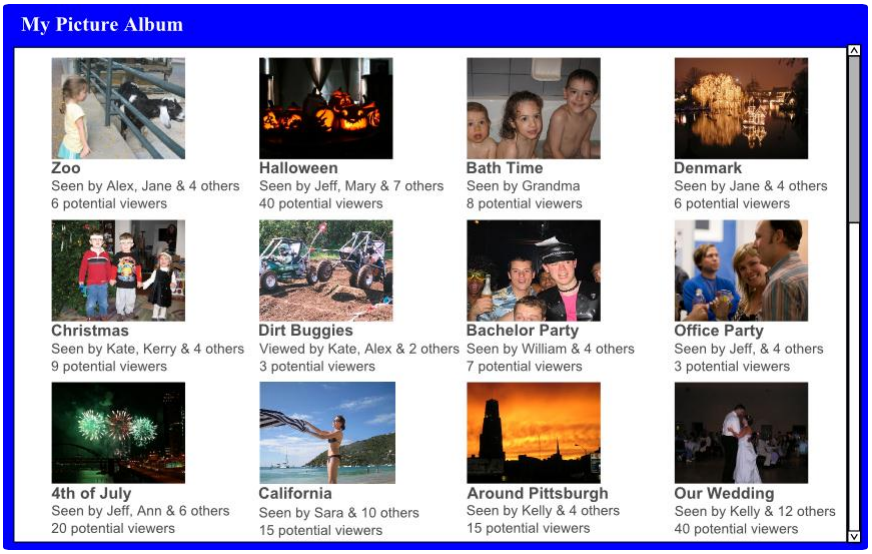

1. Name a person who can see the "Around Pittsburgh" pictures.

2. Name one album Jeff can see.

3. Name one album that Sara has viewed.

4. Name one person who has viewed the Denmark photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album that was viewed today?

8. Name a person who viewed an album today?

9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B

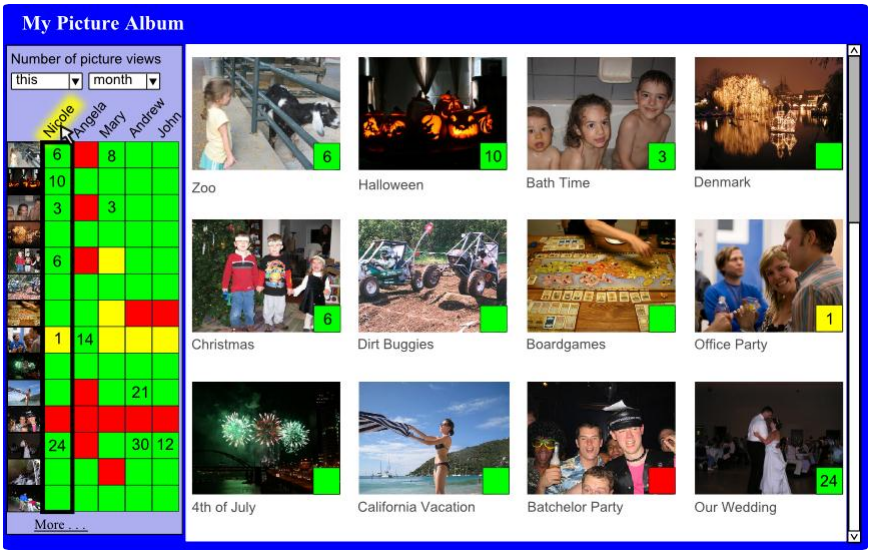

1. Name a person who can see the "4th of July" pictures.

2. Name one album Mary can see.

3. Name one album that Andrew has viewed.

4. Name one person who has viewed the Boardgames photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed this month.

8. Name a person who viewed an album this week.

9. Nicole claimed that she looked at “Our Wedding” pictures over Christmas vacation. Is this true?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C

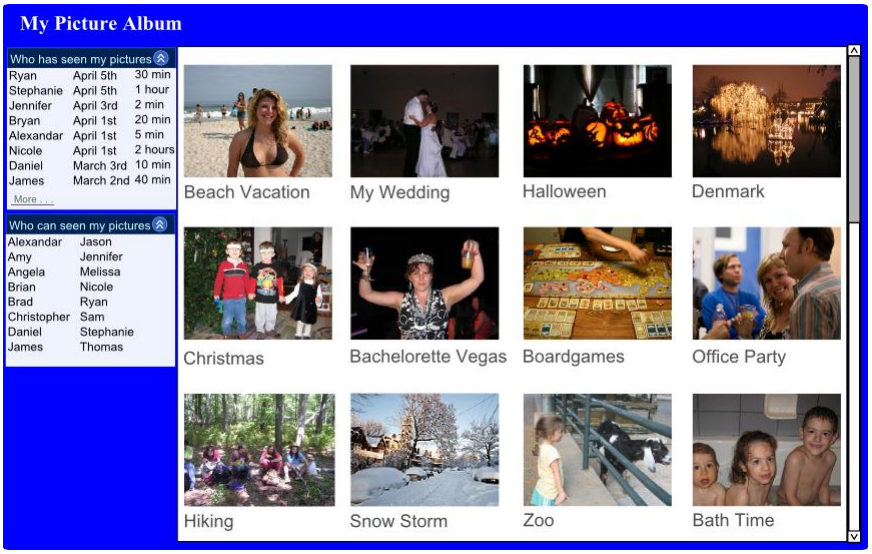

1. Name a person who can see the "Snow Storm" pictures.

2. Name one album James can see.

3. Name one album that Alexandar has viewed.

4. Name one person who has viewed the Denmark photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed in March?

8. Name a person who viewed an album in the first week of April?

9. Of the people who recently looked at pictures who spent the most time?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D

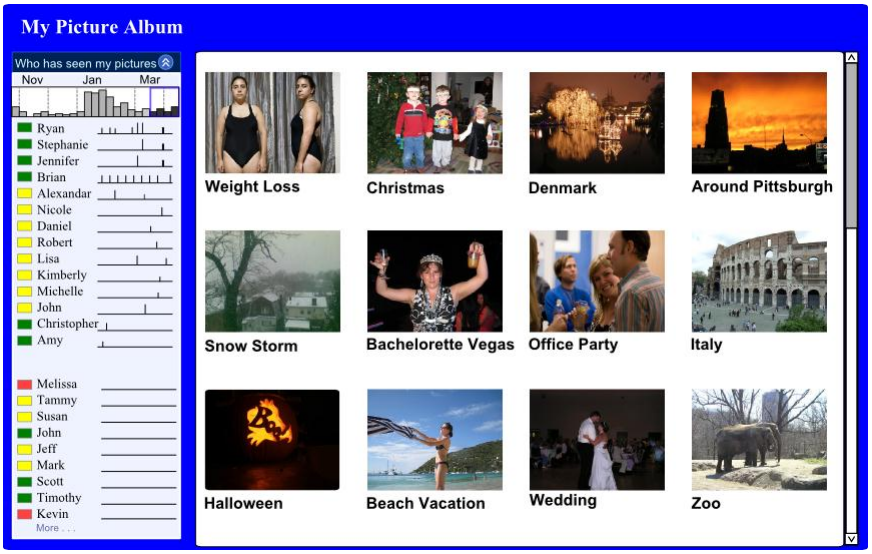

1. Name a person who can see the "Weight Loss" pictures.

2. Name one album that Nicole can see.

3. Name one album that Stephanie has viewed.

4. Name one person who has viewed the “Beach Vacation” photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. What is the most recently viewed album?

8. Who was the last person to view a picture?

9. Which friend views your pictures regularly?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E

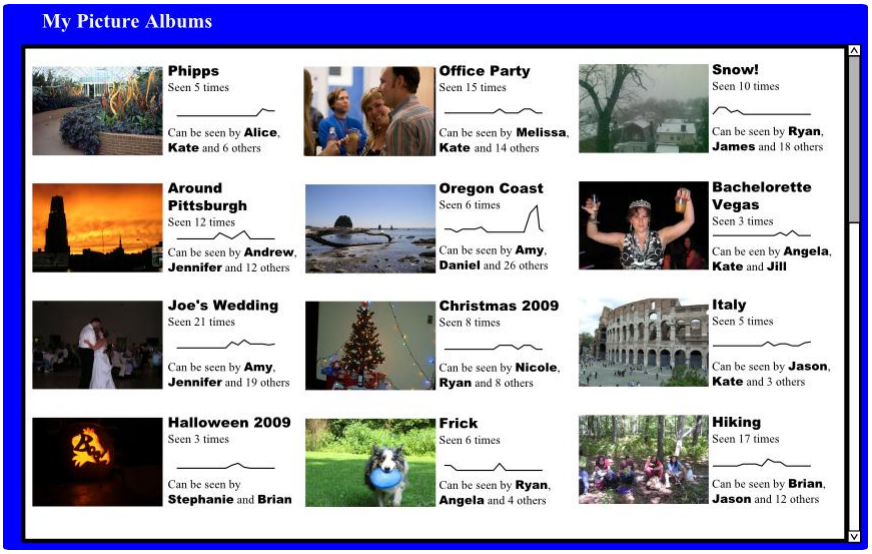

1. Name a person who can see the "Frick" pictures.

2. Name one album Brian can see.

3. Name one album that Jennifer has viewed.

4. Name one person who has viewed the Christmas 2009 photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed this month.

8. Name a person who has looked at an album this month.

9. You recently emailed out a link to one of her albums to a large number of her friends. Which album was it?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F

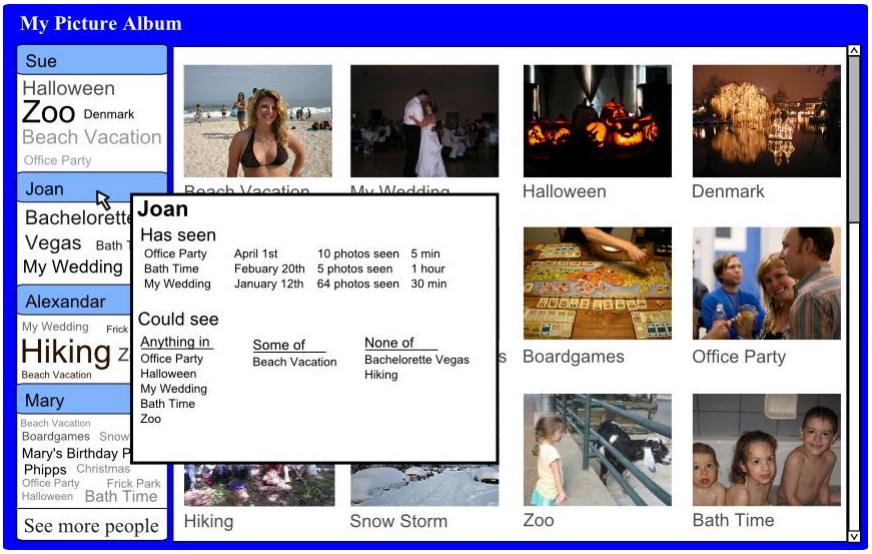

1. Name a person who can see the "Hiking" pictures.

2. Name one album Sue can see.

3. Name one album that Mary has viewed.

4. Name one person who has viewed the Zoo photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album that was recently viewed.

8. Name a person who recently viewed an album.

9. Which of your friends likes to glance at lots of your pictures?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website G

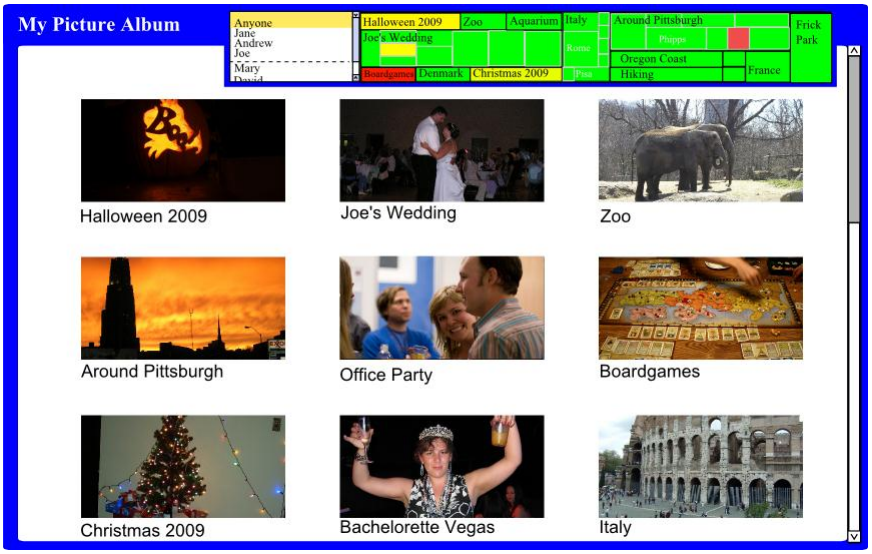

1. Name a person who can see the "Joe's Wedding" pictures.

2. Name one album Joe can see.

3. Name one album that Andrew has viewed.

4. Name one person who has viewed the Italy photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. What is the most recently viewed album?

8. Who was the last person to view a picture?

9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

A.3 Focus group 2 packet

Website A

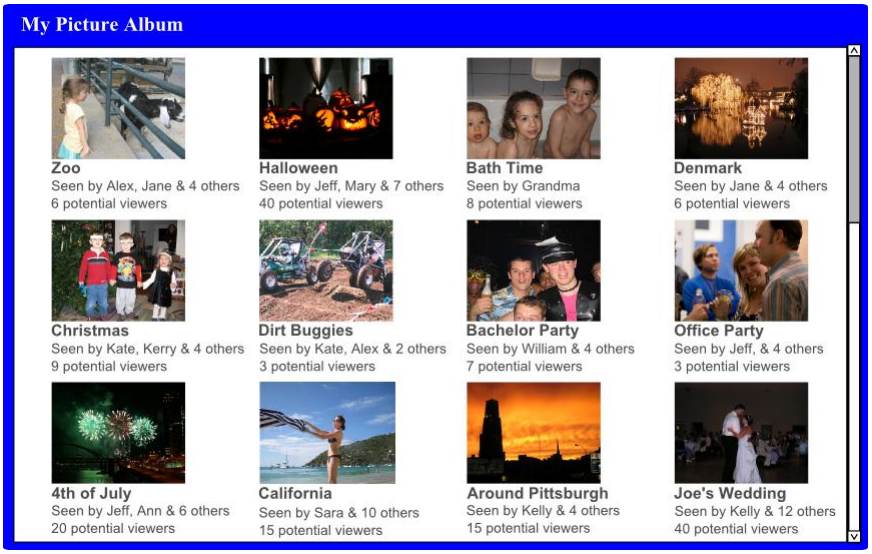

1. Name a person who can see the "Around Pittsburgh" pictures.

2. Name one album Jeff can see.

3. Name one album that Sara has viewed.

4. Name one person who has viewed the Denmark photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album that was viewed today?

8. Name a person who viewed an album today?

9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B

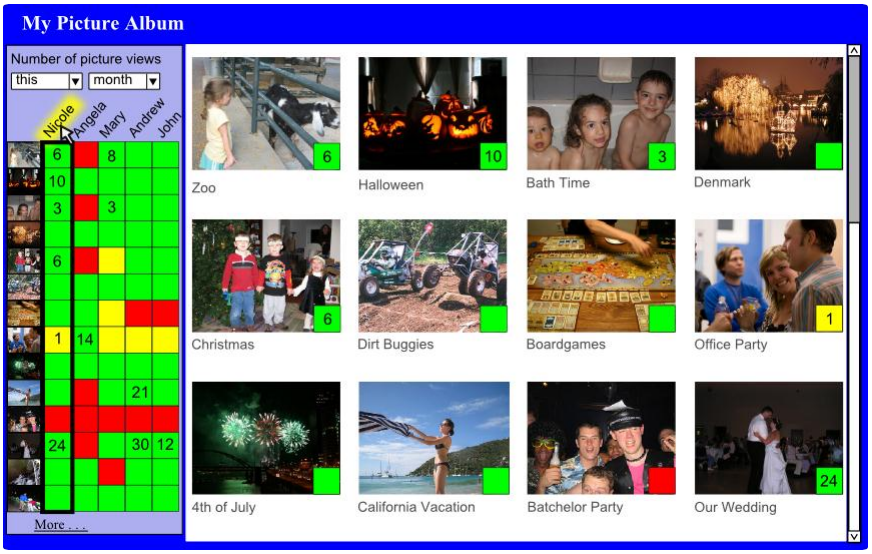

1. Name a person who can see the "4th of July" pictures.

2. Name one album Mary can see.

3. Name one album that Andrew has viewed.

4. Name one person who has viewed the Boardgames photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed this month.

8. Name a person who viewed an album this week.

9. Nicole claimed that she looked at “Our Wedding” pictures over Christmas vacation. Is this true?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C

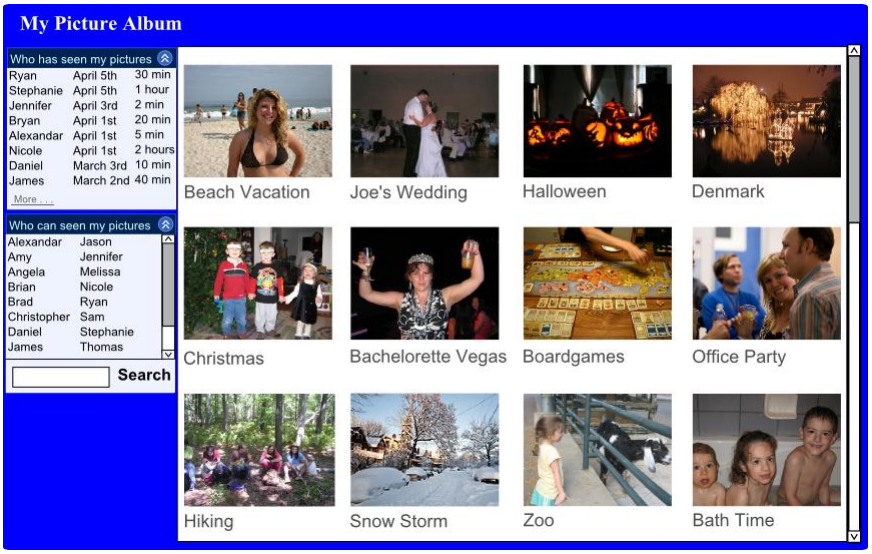

1. Name a person who can see the "Snow Storm" pictures.

2. Name one album James might be able to see.

3. Name one album that Alexandar has viewed.

4. Name one person who has viewed the Denmark photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed in March?

8. Name a person who viewed an album in the first week of April?

9. Of the people who recently looked at pictures who spent the most time?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D

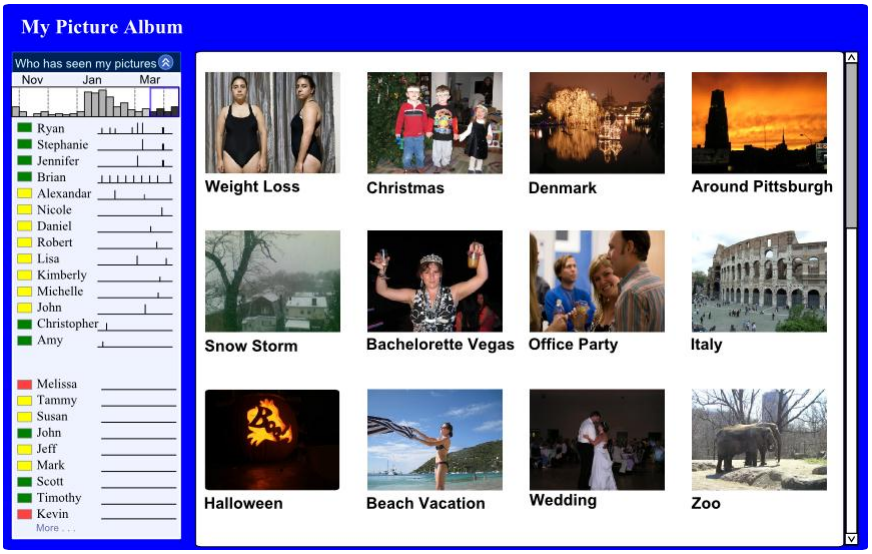

1. Name a person who can see the "Weight Loss" pictures.

2. Name one album that Nicole can see.

3. Name one album that Stephanie has viewed.

4. Name one person who has viewed the “Beach Vacation” photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. What is the most recently viewed album?

8. Who was the last person to view a picture?

9. Which friend views your pictures regularly?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E

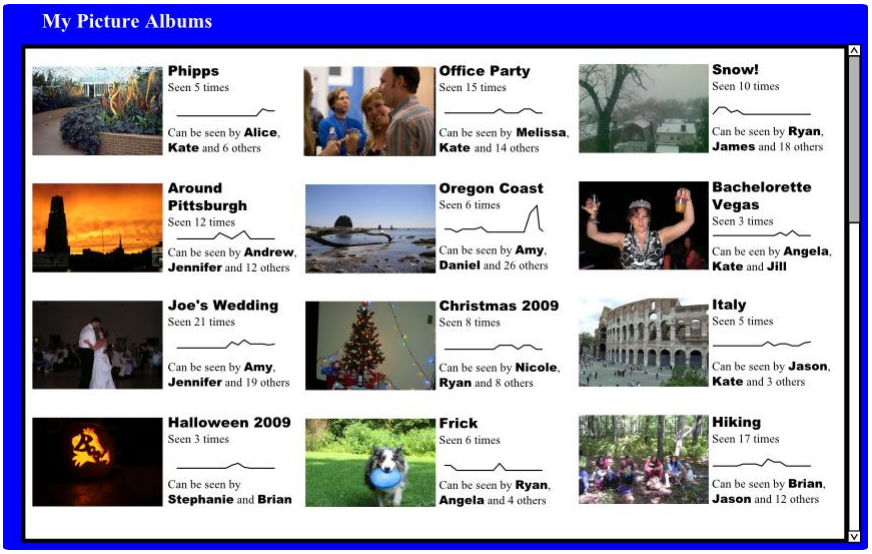

1. Name a person who can see the "Frick" pictures.

2. Name one album Brian can see.

3. Name one album that Jennifer has viewed.

4. Name one person who has viewed the Christmas 2009 photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed this month.

8. Name a person who has looked at an album this month.

9. You recently emailed out a link to one of her albums to a large number of your friends. Which album was it?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F

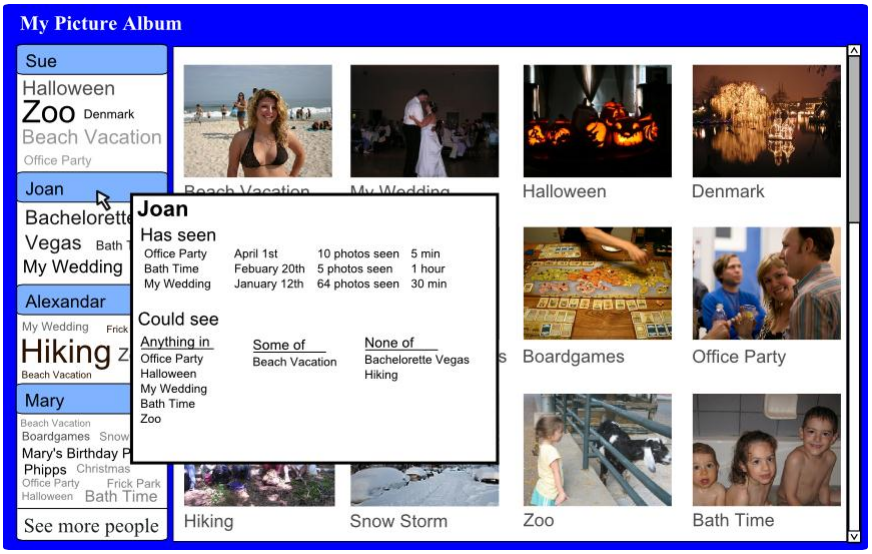

1. Name a person who can see the "Hiking" pictures.

2. Name one album Sue can see.

3. Name one album that Mary has viewed.

4. Name one person who has viewed the Zoo photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album that was recently viewed.

8. Name a person who recently viewed an album.

9. Which of your friends likes to glance at lots of your pictures?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website G

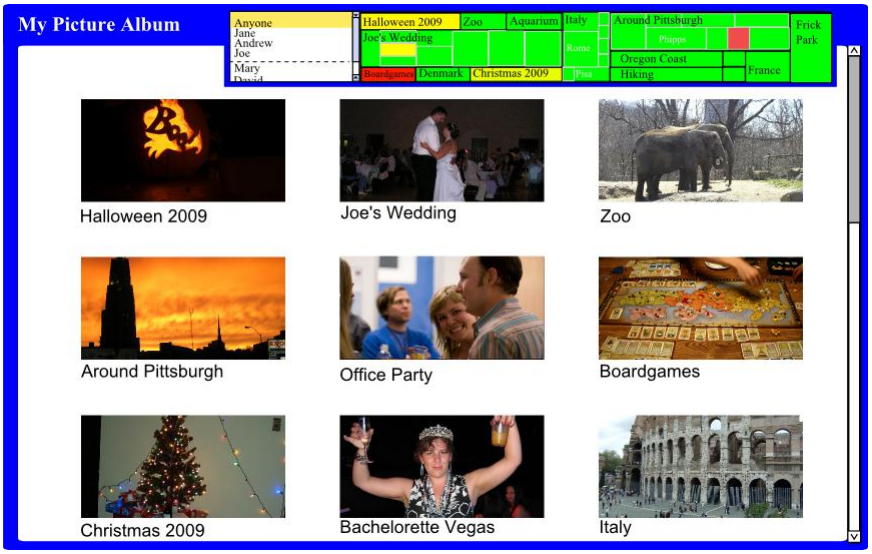

1. Name a person who can see the "Joe's Wedding" pictures.

2. Name one album Jason can see.

3. Name one album that Alexandar has viewed.

4. Name one person who has viewed the Christmas photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. What is the most recently viewed album?

8. Who was the last person to view a picture?

9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

A.4 Focus group 3 packet

Alice is looking through her online photographs for ones she can use for her screensaver.

Alice is delighted to discover that her good friend Sue found time to look at the pictures of their trip to Chicago together.

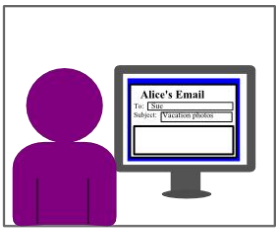

Alice hasn't had a chance to talk to Sue much since they got back so she decides this is the perfect time to reconnect. Alice writes Sue an email saying how much fun she had and she can't wait to see Sue again at Christmas.

Joe is posting pictures from his vacation on his favorite photo sharing website. He is looking through the pictures to make sure they are all ok when he notices that his ex-girlfriend, Erin, can still see his vacation pictures.

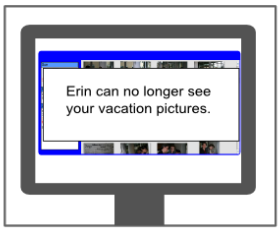

Joe is very upset and quickly removes Erin from the list of people who can see his vacation pictures.

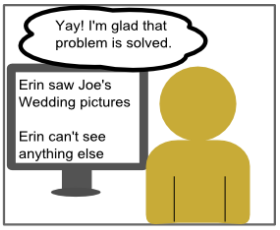

Joe is very concerned that Erin may have seen some pictures he doesn't want her to see. So he looks through the remainder of his photo album to make sure she can't see anything else and to check that she hasn't seen anything she shouldn't.

Website A

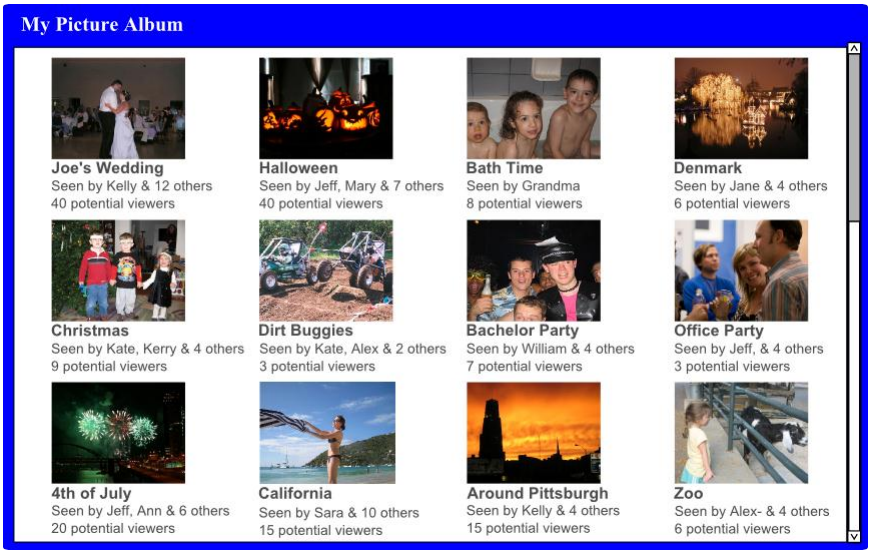

1. Name a person who can see the "Around Pittsburgh" pictures.

2. Name one album Jeff can see.

3. Name one album that Sara has viewed.

4. Name one person who has viewed the Denmark photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album that was viewed today?

8. Name a person who viewed an album today?

9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B

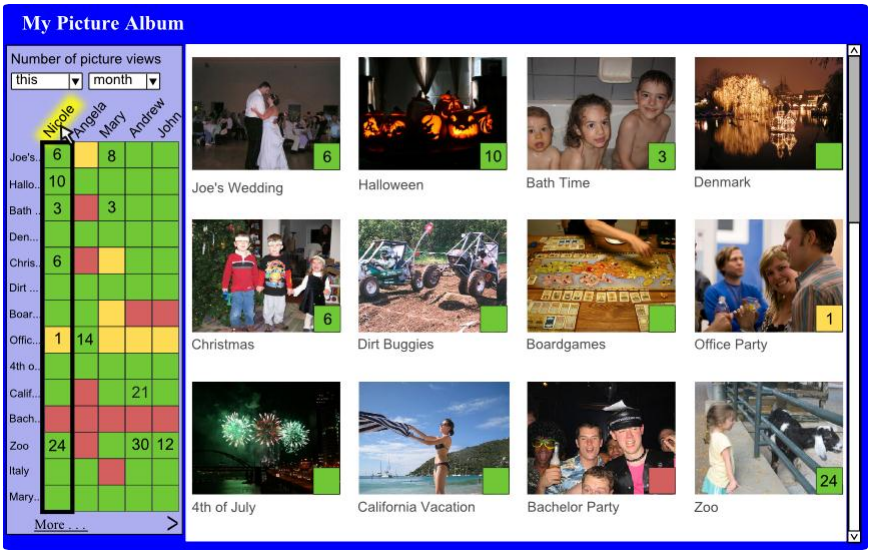

1. Name a person who can see the "4th of July" pictures.

2. Name one album Mary can see.

3. Name one album that Andrew has viewed.

4. Name one person who has viewed the Boardgames photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed this month.

8. Name a person who viewed an album this week.

9. Nicole claimed that she looked at “Our Wedding” pictures over Christmas vacation. Is this true?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C

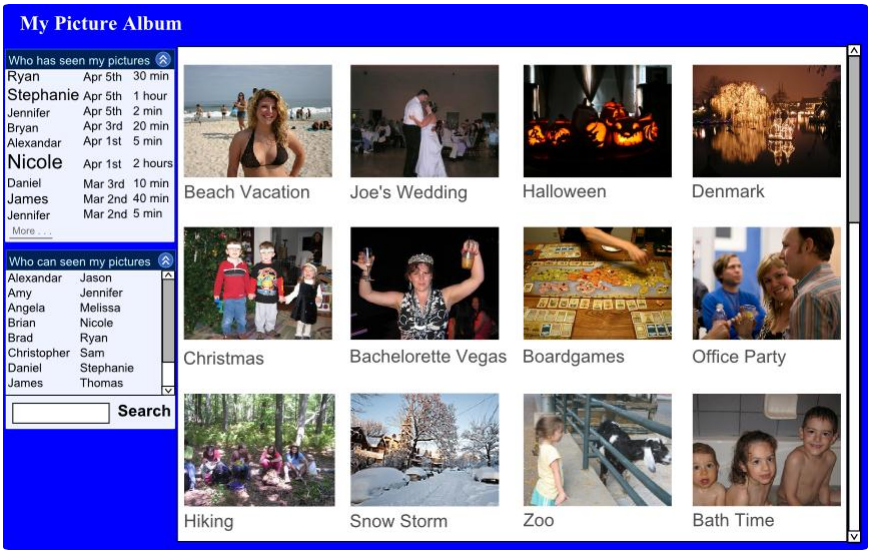

1. Name a person who can see the "Snow Storm" pictures.

2. Name one album James might be able to see.

3. Name one album that Alexandar has viewed.

4. Name one person who has viewed the Denmark photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed in March?

8. Name a person who viewed an album in the first week of April?

9. Of the people who recently looked at pictures who spent the most time?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D

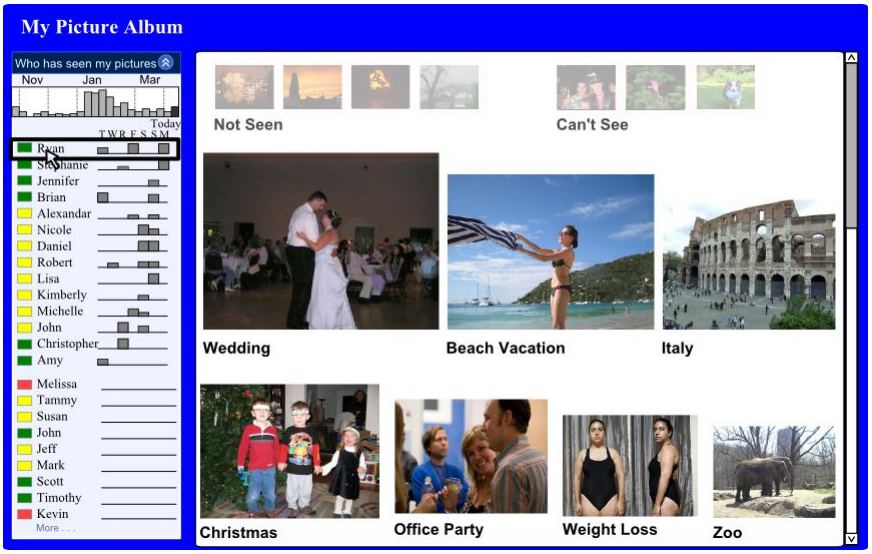

1. Name a person who can see the "Weight Loss" pictures.

2. Name one album that Nicole can see.

3. Name one album that Stephanie has viewed.

4. Name one person who has viewed the “Beach Vacation” photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. What is the most recently viewed album?

8. Who was the last person to view a picture?

9. Which friend views your pictures regularly?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E

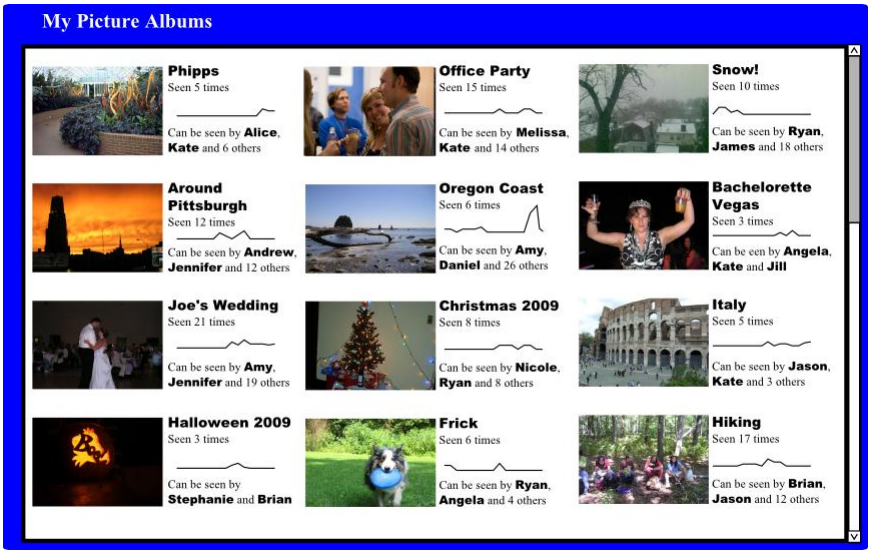

1. Name a person who can see the "Frick" pictures.

2. Name one album Brian can see.

3. Name one album that Jennifer has viewed.

4. Name one person who has viewed the Christmas 2009 photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album viewed this month.

8. Name a person who has looked at an album this month.

9. You recently emailed out a link to one of her albums to a large number of your friends. Which album was it?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F

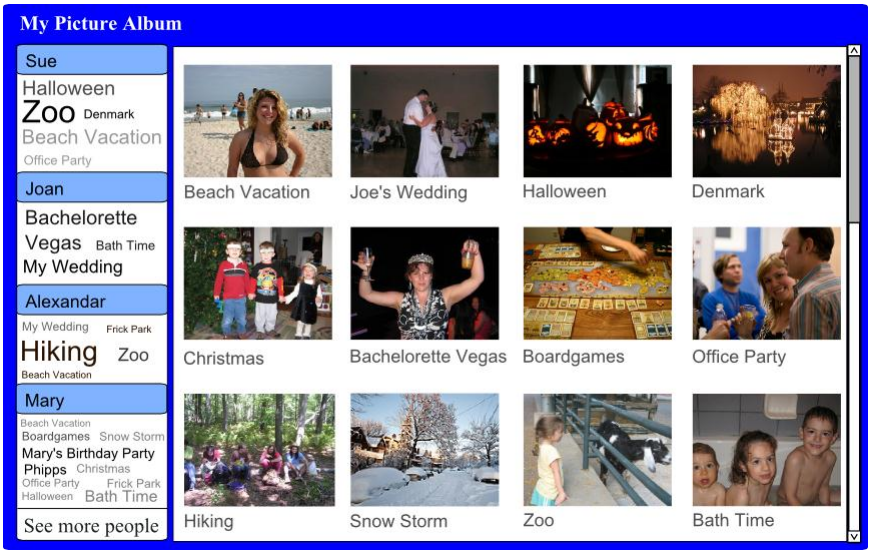

1. Name a person who can see the "Hiking" pictures.

2. Name one album Sue can see.

3. Name one album that Mary has viewed.

4. Name one person who has viewed the Zoo photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. Name an album that was recently viewed.

8. Name a person who recently viewed an album.

9. Which of your friends likes to glance at lots of your pictures?

10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website G

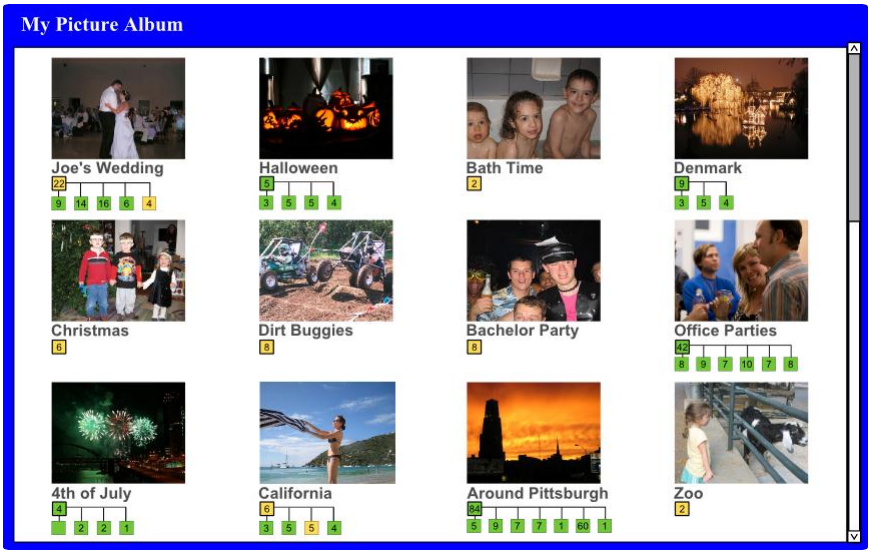

1. Name a person who can see the "Joe's Wedding" pictures.

2. Name one album Jason can see.

3. Name one album that Alexandar has viewed.

4. Name one person who has viewed the Christmas photos.

5. What is the most frequently viewed album?

6. Who viewed the most photos?

7. What is the most recently viewed album?

8. Who was the last person to view a picture?

9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

A.5 Focus group 4 and 5 packet

Alice is looking through her online photographs for ones she can use for her screensaver.

Alice is delighted to discover that her good friend Sue found time to look at the pictures of their trip to Chicago together.

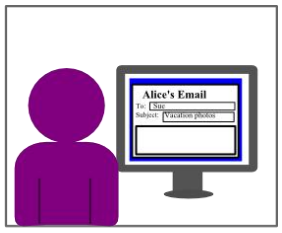

Alice hasn't had a chance to talk to Sue much since they got back so she decides this is the perfect time to reconnect. Alice writes Sue an email saying how much fun she had and she can't wait to see Sue again at Christmas.

Joe is posting pictures from his vacation on his favorite photo sharing website. He is looking through the pictures to make sure they are all ok when he notices that his ex-girlfriend, Erin, can still see his vacation pictures.

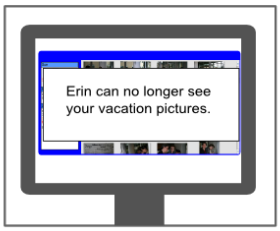

Joe is very upset and quickly removes Erin from the list of people who can see his vacation pictures.

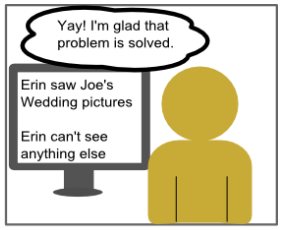

Joe is very concerned that Erin may have seen some pictures he doesn't want her to see. So he looks through the remainder of his photo album to make sure she can't see anything else and to check that she hasn't seen anything she shouldn't.

Website A

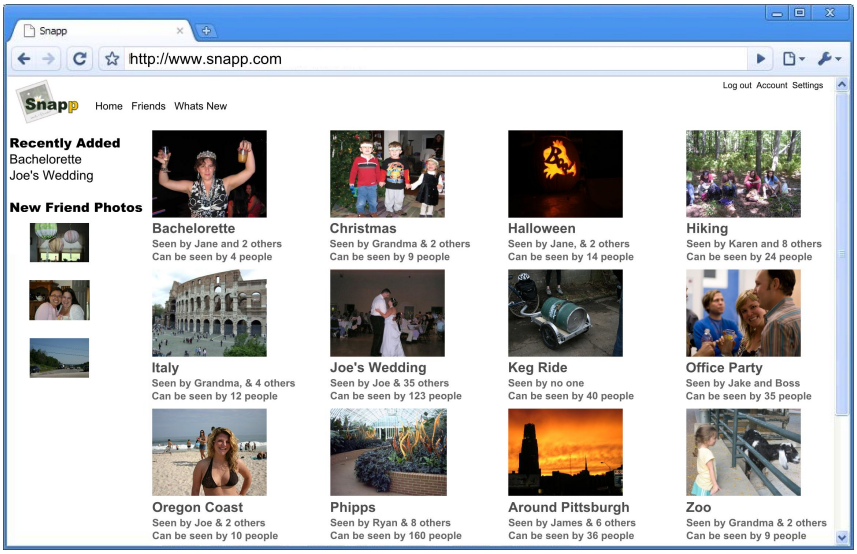

1. Name a person who can see the "Around Pittsburgh" pictures.

2. Name one album Grandma has seen.

3. What is the most frequently viewed album?

4. Name an album that was viewed today?

5. Would this website have helped Alice? Why?

6. Would this website have helped Joe? Why?

7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B

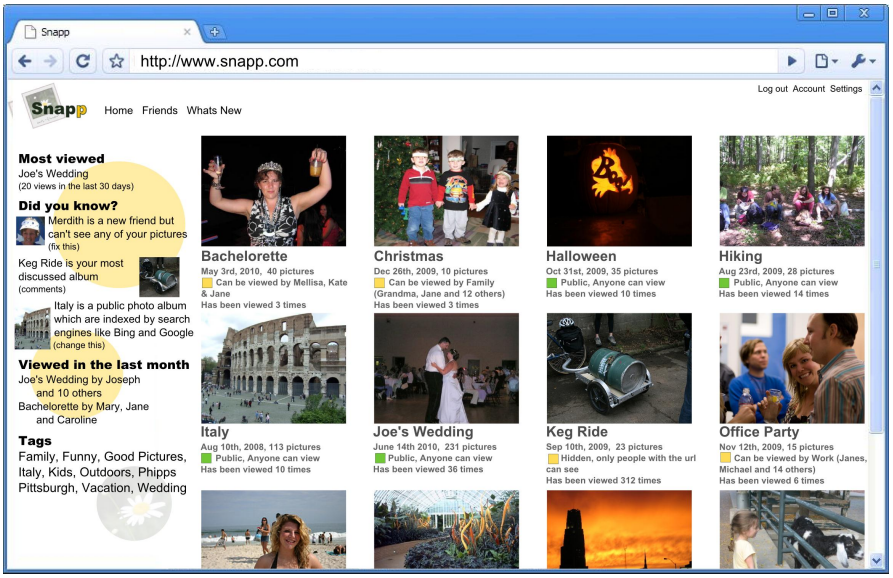

1. Who can see Joe's Wedding pictures.

2. Name one album Merdith **cannot** see.

3. What is the most frequently viewed album?

4. Name an album that was viewed today?

5. Would this website have helped Alice? Why?

6. Would this website have helped Joe? Why?

7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C

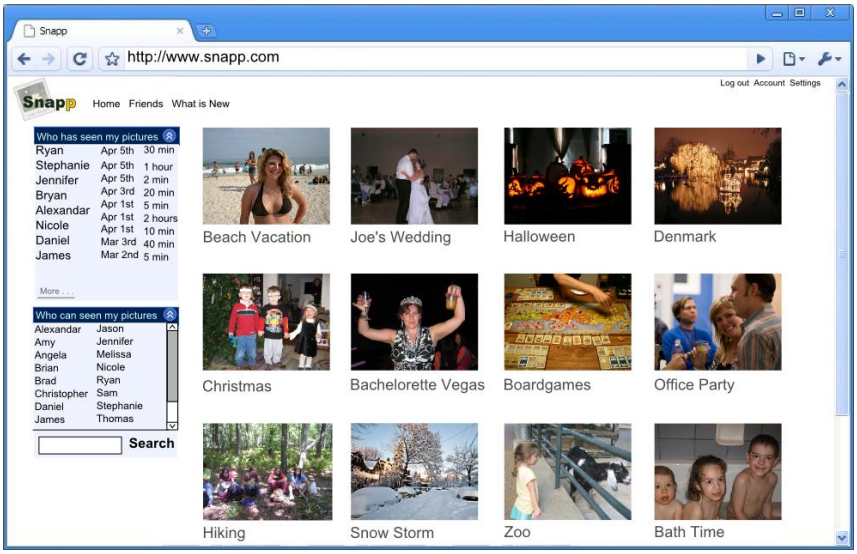

1. Name a person who can see the "Snow Storm" pictures.

2. Name one album Alexandar can view.

3. What is the most frequently viewed album?

4. Name an album that was viewed today?

5. Would this website have helped Alice? Why?

6. Would this website have helped Joe? Why?

7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D

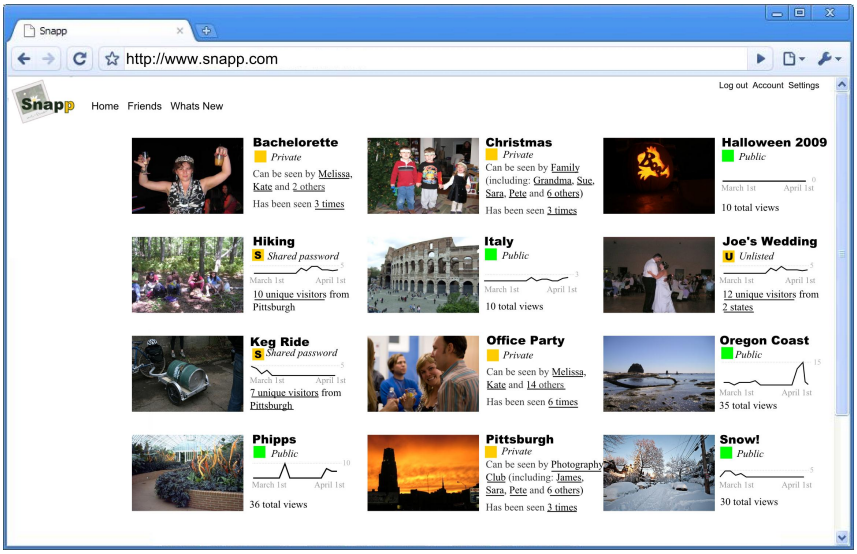

1. Name a person who can see the Office Party pictures.

2. Name one album that Nicole can see.

3. What is the most frequently viewed album?

4. Name an album that was viewed today?

5. Would this website have helped Alice? Why?

6. Would this website have helped Joe? Why?

7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E

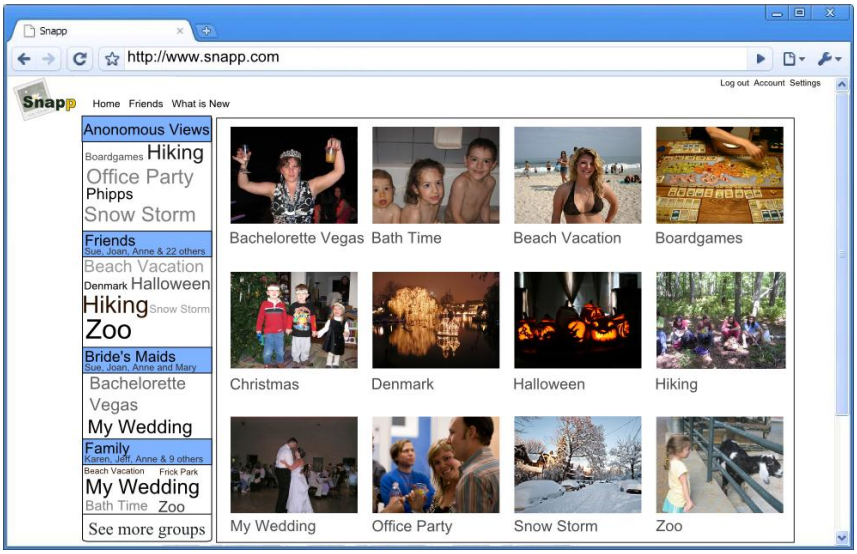

1. Name a person who can see the "Hiking" pictures.

2. Name one album Sue can see.

3. What is the most frequently viewed album?

4. Name an album that was viewed today?

5. Would this website have helped Alice? Why?

6. Would this website have helped Joe? Why?

7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F

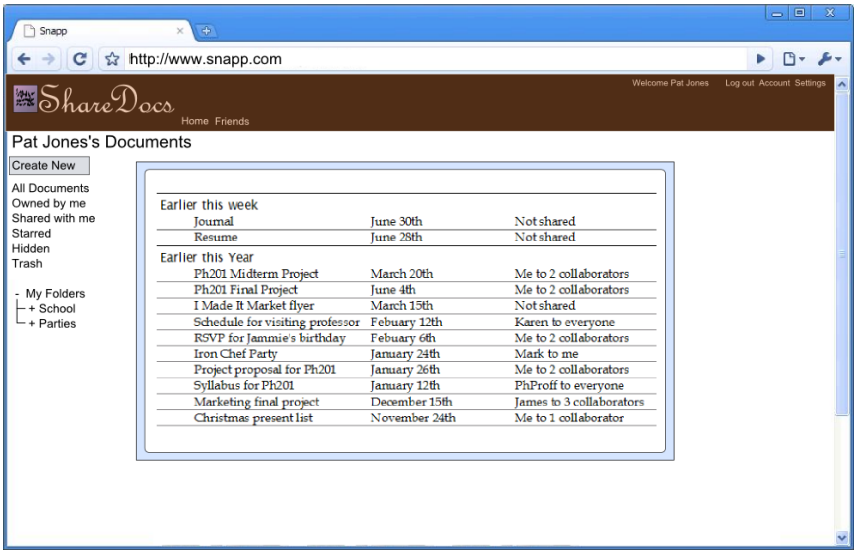

1. Name a person who can see the “Sillabus for Ph201”.

2. Name one document Karen can see

3. Name one document created by a friend.

4. Name a document last seen this week.

5. Would this website have helped Alice?

6. Would this website have helped Joe?

7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website G

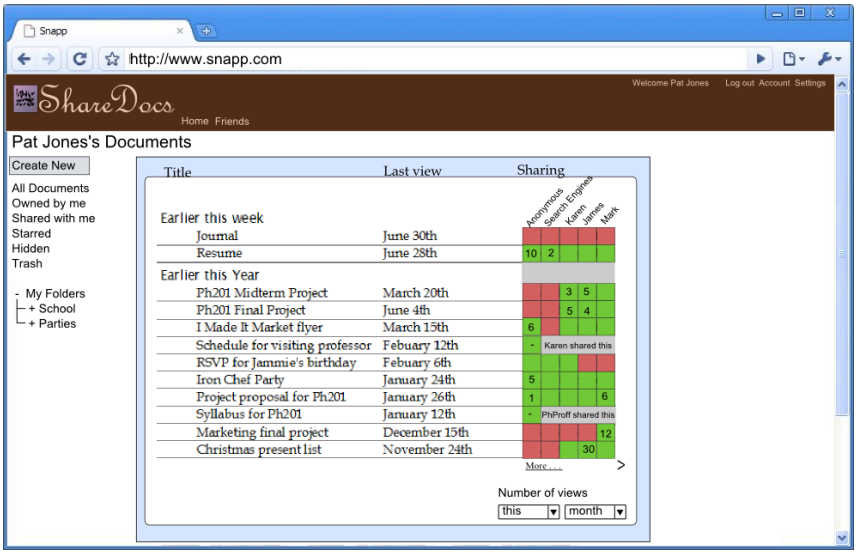

1. Name a person who can see the Ph201 Midterm Project.

2. Name one document James can see.

3. What is the most frequently viewed document?

4. Name a document that was viewed this week?

5. Would this website have helped Alice? Why?

6. Would this website have helped Joe? Why?

7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Appendix B

Eye tracker study (study 2)

B.1 Printed instructions and emails

In the eye tracker study (study 2) participants were given instructions and emails by the researcher on printed sheets of paper. The remainder of this appendix section is all the instructions and emails used in the study. Each instruction or email was printed on its own sheet of paper, but in the interests of saving space, we show only the textual content of the pages. Each box of text was printed on a single page, without the black border.

The pages which give instructions and the pages with emails that initiate tasks were given to all participants. The pages with emails used to prompt the participant, were given to the participant only if the participant did not complete all parts of the task.

Instructions

In this study you will be asked to role play a person called Pat Jones. Every time you have to make a decision or judgment call, I want you to think about how Pat Jones would handle the situation and handle it that way.

During this study you should think about the photo albums you are working with as your own (well, Pat Jones's). If you see something that you would change in your own album then go ahead and change it or just say it out loud so I know what you would have changed if you had time.

Today I will give you several information pages and emails written on pieces of paper. Some of the emails will contain simple and straightforward tasks and some will be less directed to get a better sense of how you approach and complete photograph management tasks in general. When you are ready to respond to an email just say out loud what you would email back. Once you have responded I will hand you another piece of paper with the next email.

We are interested in how you approach and solve the issues presented to you. Remember, we are testing the software and how it supports how you work with photographs. We are not testing you.

Say "Done" when you are finished reading this page.

Instructions: Pat Jones

Your name is Pat Jones. You are an administrative assistant at a large web hosting and data storage company called Global Storage. Your company uses a popular online photo sharing site called Gallery to store and share their photographs. You, your family and most of your friends also use Gallery to store and share photographs.

You use Gallery because it gives each person lots of space, it makes it easy to share with only certain groups of people and it lets people, like your Mom, give others the ability to administer their albums for them without having to give out the password. This makes it easier to help your friends and family when they have problems.

Global Storage has a company wide album on Gallery where company related photographs are posted. As an administrative assistant at Global Storage, one of your jobs is to take photographs of events and post them in the company album. The last administrative assistant wasn't very good at this and left errors all through the albums which you clean up as you find them. Your boss and coworkers often ask you to do photo management tasks to keep the company photo album in order and looking good.

All the Global Storage photographs are in the album called "Global Storage" though some employees keep photographs in their personal albums.

Say "Done" when you are finished reading this page.

To: Pat Jones <pat@globalstorage.com>
From: Angela Wilson <angela@globalstorage.com>
Subject: Sideways photograph

Hello Pat,

I was looking through the *Around the office* album in the "Global Storage Shared Albums" and I noticed that Gerald's photograph is sideways.

Could you please fix that.

Thanks,
Angela

To: Pat Jones <pat@globalstorage.com>
From: Angela Wilson <angela@globalstorage.com>
Subject: Sideways photograph

Hello Pat,

Gerald's photograph still appears to be sideways. You can find if if you go into the "Global Storage Album" and then go to "Around the Office". Gerald's photograph is in the upper right hand corner.

Thanks,
Angela

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I noticed that there is an album entitled *Around the office* inside of the *Global Storage Shared Album* album. The photographs you have there are really great! If I email a link to someone at another company, will they be able to see the photos in that album? Its ok if they can't I just want to know before I send an email.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Emailing photographs

Hi Pat,

Are you sure my friends who aren't in the company will be able to see the photos? I remember doing this before and it didn't work . . .

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Emailing photographs

Hi Pat,

Are you sure? I was about to send off the email when Angela dropped by and she swears she saw you looking at the wrong album when you emailed me.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Kevin Brown <kevin@globalstorage.com>
Subject: Remove photos of me

Hi Pat,

I heard this horrible rumor that you put all our photographs on the Internet where anybody could see them and now the boss is emailing the photos to his friends? I know you take great photographs but I look horrible in photos and I really don't want that on the Internet. Could you please delete the photo of me in the People album?

Thanks,
Kevin

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Remove photos of me

Hi Pat,

Kevin really wants his photo taken down. I had a bit of a talk with him about it because I think it is important to have these photos up. The compromise was that you would take the photograph down and I would have our professional photographer take a photograph of Kevin and put it up later.

So please remove Kevin's photograph from the People album.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I checked with human resources and our lawyer, it is fine to allow employees to add photos to an online album. So go ahead and give Global Storage employees (coworkers) the ability to add photos to the "Around the office" album.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I've been going through the company photo albums all afternoon. Great new photos by the way. I noticed that there is an album called *Around the office* which seems to be great set of photos of day-to-day events in the office. I've noticed that other people sometimes take photographs around the office but they don't seem to be able to add them to this album.

I'd love it if you made it so other people in the office could add to the *Around the office* album. That way we can have all these great pictures in one place.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I know you said you fixed it so other Global Storage employees could add to the "Around the office" album but I just tried and it didn't work. Could you fix it?

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Steve Johnson <steve@globalstorage.com>
Subject: Product fair titling issues

Hello Pat,

The boss has finally decided to pay attention to your photographs; he has been sending me pictures of myself on and off all day. The last one he sent was from the Project Fair and I noticed that you had mistitled my poster. Actually, it looks like you may have switched my title with someone else's, so theirs is wrong too. You should be able to get the correct titles by reading the posters behind each person.

Thanks,
Steve

To: Pat Jones <pat@globalstorage.com>
From: Steve Johnson <steve@globalstorage.com>
Subject: RE: Product fair titling issues

You can find the titles of the posters by looking at the photographs. You can easily read each title behind the person if you just open the photograph instead of looking at the thumbnail.

Sorry, I can't remember my exact title right now.

Thanks,
Steve

To: Pat Jones <pat@globalstorage.com>
From: Steve Johnson <steve@globalstorage.com>
Subject: Product fair titling issues

Hello Pat,

Um, I noticed that you fixed one of the poster titles but not the other one. Could you go fix the other title please?

Thanks,
Steve

To: Pat Jones <pat@globalstorage.com>
From: Ralf Jackson <ralf@globalstorage.com>
Subject: Sort these photos

Hi Pat,

I'm putting together a presentation and I want to use a bunch of photographs of signs that I've been randomly taking over the last couple of years. Could you look through my "Random Photos" album inside the "Ralf Jackson's Album" and move all the photographs of signs to the empty Funny Signs album I made?

Thanks,
Ralf

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Ski trip photos

Hey Pat,

The professional photographer from the company ski trip finally sent me photographs and I want to send them out as part of my weekly email to employees and family members. Please create a new album in the Global Storage albums and put the photographs in it.

I'll send out the email about the photos as soon as you tell me they are up. Don't change any of the titles, in my newsletter I'm going to ask everyone to open the album and create their own titles for the photographs. I don't know if that just works or not.

After you upload the photos can you make sure that none are sideways? Also, make sure there aren't any photos of alcohol or anyone drinking. Susie in marketing may use these later and for some reason she thinks pictures of people drinking are a good thing to send out in a family friendly newsletter, but I don't.

This is going to be great,
Gerald
(The Boss)

The ski trip photographs are on your Desktop in a folder labeled Ski Trip.
To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Put these on your photo site

Hey Pat,

I was just reviewing the ski pictures and I noticed a photo of what looks like alcohol. Please remove it. I don't want any alcohol pictures in this album.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Put these on your photo site

Hey Pat,

I was looking through the ski photographs when I noticed one that was sideways. Please go make sure they are all straight. I don't like untidy photo albums.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Publicity photos

Hi Pat,

This last week we had a public show case of our new product line. I created an album entitled "New Products" of all the great photographs I collected from the event. But I'm not ready to go public with it yet and really don't want anyone but coworkers seeing it. Could you go through and clean things up a bit? All the photos need to have titles. You can pick whatever title you think is appropriate. I already went through and organized them so everything is in the correct order. I had some trouble because Susan in marketing couldn't see or edit the photographs but I fixed that one.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Publicity photos

Hi Pat,

I just looked through the New Products album and I found a photograph that was sideways. Please make sure they are all oriented correctly.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Publicity photos

Hi Pat,

I looked through the New Products album and noticed that some of the photographs still have names like IMG123. Could you please give them English sounding titles. The titles don't have to be complex they can be things like "Examining new product."

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: RE: Ski trip photos

Hey Pat,

I'm really glad I am getting so much engagement on the whole ski trip photo titling idea. But it looks like some employees may have gotten a bit over zealous with their changes.

Susan from marketing is about to put together the monthly news letter and is going to include the ski trip photos and the titles employees have added. As I said before, I'm a bit concerned about what Susan considers to be family friendly photos, not to mention her lack of technical skills. I'm concerned she may try and use a sideways photo or pick something inappropriate. Please go through the photos and make sure they are all in good condition for a newsletter and that everything is ready for Susan.

Thanks,
Gerald
(The Boss)

>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
> To: Pat Jones <pat@globalstorage.com>
> From: Gerald Fredricks <gerald@globalstorage.com>
> Subject: Ski trip photos
>
> Hey Pat,
>
> The professional photographer from the company ski trip finally sent me photographs and I want to > send them out as part of my weekly email to employees. Please create a new album in the Global
> Storage album and put the photographs in it.
>
> I'll send out the email about the photos as soon as you tell me they are up. Don't change any of the
> titles, in my newsletter I'm going to ask everyone to open the album and create their own titles for the
> photographs.
>
> After you upload the photos can you make sure that none are sideways? Also, make sure there aren't
> any photos of alcohol or anyone drinking. Susie in marketing may use these later and for some reason
> she thinks pictures of people drinking are a good thing to send out in a family friendly newsletter, but
> I don't.
>
> This is going to be great,
> Gerald
> (The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: RE: Ski trip photos

Hey Pat,

Some of the Ski Trip photos appear to be sideways. Please fix this.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: RE: Ski trip photos

Hey Pat,

There are still some photographs with titles which clearly mention alcohol. Please change these to some other appropriate title.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: RE: Ski trip photos

Hey Pat,

Since the ski trip photographs will be in the newsletter I want to be sure that they are visible to friends and family. One of the admins claims to have fixed it so that the photos are visible to each employees friends and family. Can you tell me if these photos are visible to your friends and family?

Thanks,
Gerald
(The Boss)

Information: Adventures

Despite having a normal desk job you really like to go out and do fun things on the weekends. When it comes to exciting activities like sky diving you will try anything once. You make sure to post photos of all your adventures so your friends can see. However, your mother is one of those people who panics easily and you know if she ever saw a photograph of you diving out of an airplane you would never hear the end of it. So you make sure not to mention some of your more exciting adventures.

Unlike your work, your friends all put their photos in there own albums.

Say "Done" when you are finished reading this page.

Information

It is now Sunday and you had the weekend off. You are now at your home computer checking email.

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: New photos

Yo Pat,

Here are the better photos from the Building Jumping trip last weekend. Could you put them up on Gallery for me? Just set it up in your album (Pat Jones's Albums) where everyone already knows to look. Also could you title the photos with the people in them? I had the red parachute, George had the green one and of course your's was blue.

When you are finished let me know so I can have all our friends go look at it.

Thanks,
Josh

The photos Josh sent are in a folder labeled *Building Jumping* on your desktop.

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: RE: New photos

Hi Pat,

I'm not going to upload these photos because I don't have the time. Please upload them.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: RE: New photos

Hi Pat,

I see the photos are up but they don't have any titles. Please title the photos with the people in them? I had the red parachute, George had the green one and of course your's was blue.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: Are you ok?

Pat,

Are you all right? Are you ok?

I just sent Aunt Sue a link to Jennifer's Baby pictures and she sent me back this photo of you jumping off a building. A BUILDING! Are you crazy? What were you thinking? Do you realize how dangerous what you are doing is? People die from this!

Uncle David already thinks I'm a poor mother, if he sees these photographs I will NEVER hear the end of it. And he is going to be looking as soon as he gets home because I already sent him a link to Jennifer's Baby pictures. What were you thinking? How could you do this to me?

Please, please make sure no more of our family see these photographs.

Mom

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

I'm so glad you made it to the after dark safari trip last month. Wasn't it cool seeing some of those animals at night?

I added some of the photos from the safari trip to my album but somehow most of them turned out sideways. Could you rotate them? Also, I seemed to have uploaded a photo from one of the Pirates games and can't seem to find how to delete it. Could you delete it for me while you are at it?

When you are done let me know so I can email the link out to all our friends. I can't wait for them to see some of these great shots.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

I just checked and some of the photographs from the Safari trip are still sideways.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

That photograph of the Pittsburgh Pirates is still in the Safari album. Could you please help me delete it?

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

Hey, it just occurred to me. Can our friends even see the Safari album? I'm not sure how to check and I don't want to send it out if they can't see it.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: George Wilson <george@hotmail.com>
Subject: Safari Photos

Hey Pat,

I helped Rich re-build his porch this weekend and I took photographs of the whole thing. I think we did an awesome job but I'll let you judge for yourself.

We ripped out half the porch and put in all new mini foundation pieces for each support. Then we put in a whole new frame and put the surface boards back on and painted.

We thought we were done but then Kerry came out, and well you know how much she likes plants, and Rich and I had all these tools laying around. . . So we went ahead and built her a set of planter boxes for the porch. We started out with some small ones then we built some longer ones too.

Finally, we put in some new stairs that will look better and creak less.

Anyway, the whole reason I'm emailing you is that I can't seem to figure out how to put the photos in order. It looks silly right now with the photos of planter boxes appearing before the photos of us putting in the deck. Could you organize them for me?

Thanks,
George

To: Pat Jones <pat@jones.com>
From: Lisa Williams <josh@hotmail.com>
Subject: New photos

Hi Pat,

Thanks for setting up that great album for our Building Jumping trip. I took some photos too so I went ahead and uploaded them. Could you double check that I didn't mess anything up and that all the photos look ok, your photo sharing system always confuses me.

Thanks,
Lisa

To: Pat Jones <pat@jones.com>
From: Lisa Williams <josh@hotmail.com>
Subject: RE: New photos

I just sent out the link to the building jumping photos but I'm getting complaints because our friends can't see the new photographs. Josh sent me an unflattering email about how I shouldn't be allowed to upload photos. What did I do wrong? Could you please fix it?

Thanks,
Lisa

To: Pat Jones <pat@jones.com>
From: Lisa Williams <josh@hotmail.com>
Subject: RE: New photos

I just sent out the link to the building jumping photos but now Josh is making fun of me because one of the photos is sideways. Is there an easy way to turn it back round? Could you please fix it?

Thanks,
Lisa

Information: Pat's Family

Your parents can barely operate their computer much less manage a photo site. So you let your family post photographs in their own album but you help out by checking each album to make sure it is not visible to everyone on the Internet.

You help your parents manage their photos when they upload new albums. Your mother doesn't understand the photo management software on her computer and tends to make a ton of silly mistakes like once accidentally titling your Dad *Fido*. She is perfectionist and not being able to make her photos look perfect really annoys her so you help her out by fixing up the photographs before she lets her friends and family see anything.

Your mother's name is Samantha and all her photographs can be found in "Samantha Jones's Albums".

Say "Done" when you are finished reading this page.

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: New albums

Hi Pat,
I just uploaded the Christmas photos at Jennifer's to my web album. Aunt Sue has been asking about the Christmas photos for months. I'm so glad I finally found time to do this.

I followed the instructions you gave me last time you showed me how to put photos on your photo site but they were so complex I didn't get through all of them. I'm concerned I might have made a few mistakes. To begin with I think I uploaded some photos from my Mexico vacation into the Christmas album. So could you please go and delete any photos that look out of place. Also, I think I might have mixed up a few titles.

Could you please go look at the albums and fix any mistakes I might have made? Let me know when you are done so I can email the family so they can see the pictures.

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: New albums

Hi Pat,

Thanks for fixing up my new albums. I took a quick glance over them and I think there may still be some errors with the titles. The picture with little Henry holding his pillow at Christmas is still labeled *Susan and new pillow*.

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: New albums

Hi Pat,

Thanks for fixing up my new albums. I took a quick glance over them and I think there may still be some problems. The picture of Susan with her arms out is sideways. Could you please make it straight?

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: Re: New albums

Hi Pat,

Aunt Sue just emailed me and she says she can't see my Christmas photographs. Where did they go? Why can't she see them?

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Separate out some photos

Hi Pat,

I uploaded photos of the snow storm, that Pirates game I went too, and a trip to the Phipps Conservatory during their gargoyles exhibit, into that Misc album you created for me. I even managed to create three new albums for the photographs. The only problem is that I can't seem to get the photos moved from the Misc album to the albums they need to be in.

Thanks,
Jennifer

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Baby photos

Hi Pat,

I just took a bunch of photos of my new baby Angela and I want to share the photos with family, friends and coworkers. Could you create a new album for them in "Jennifer Smith's Albums" and put the new photos in it? When you are done I need you to find the cutest one and make it the album cover.

Thanks,
Jennifer

The photos Jennifer sent are in a folder labeled *Angela* on your desktop.

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Re: New albums

Hi Pat,

Mom was all upset about her photos not being quite right so I had her log in for me and tried to fix them myself. But Mom hated all my changes and wants things back the way they were. Could you go back through her Christmas album and just put everything back the way it was?

Thanks,
Jennifer

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Re: New albums

Hi Pat,

Mom says that all the photos used to be straight and now one is not. She isn't letting me touch the computer anymore, can you please fix it.

Thanks,
Jennifer

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Re: New albums

Pat,

Mom is terribly worried that other people not in our family are looking at her photos. I told her that it was fine but could you please just check.

Thanks,
Jennifer

B.2 Online survey

Gallery Prox Info Display (June 27, 2011)

New Page

1. User ID

Page One

2. Did you find working with Gallery today to be: *

- Very Enjoyable
- Enjoyable
- Neutral
- Unpleasant
- Very Unpleasant

New Page

Gallery uses a set of icons to indicate information about privacy settings. For each icon below describe what you think the icon means.

3. *

4. *

5. *

6. *

7. *

Funny Signs

8. Ralf Jackson asked you to move signs from his "Random Photos" album to another album called "Funny Signs". What was the privacy policy for the Funny Signs album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Funny Signs album. *	○	○	○
Everybody can **view** the Funny Signs album. *	○	○	○
Family can **add** to the Funny Signs album. *	○	○	○
Family can **view** the Funny Signs album. *	○	○	○
Coworkers can **add** to the Funny Signs album. *	○	○	○
Coworkers can **view** the Funny Signs album. *	○	○	○
Friends can **add** to the Funny Signs album. *	○	○	○
Friends can **view** the Funny Signs album. *	○	○	○

9. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Ski Trip

10. Your boss asked you to create an album for the company ski trip. What was the privacy policy for the Ski Trip album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Ski Trip album.	○	○	○
Everybody can **view** the Ski Trip album.	○	○	○
Family can **add** to the Ski Trip album.	○	○	○
Family can **view** the Ski Trip album.	○	○	○
Coworkers can **add** to the Ski Trip album.	○	○	○

Coworkers can **view** the Ski Trip album.	○	○	○
Friends can **add** to the Ski Trip album.	○	○	○
Friends can **view** the Ski Trip album.	○	○	○

11. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

New Products

12. Your boss asked you to review the New Products album for errors. What was the privacy policy for the New Products album when you left it? *

	True	False	Not Sure
Everybody can **add** to the New Products album.	○	○	○
Everybody can **view** the New Products album.	○	○	○
Family can **add** to the New Products album.	○	○	○
Family can **view** the New Products album.	○	○	○
Coworkers can **add** to the New Products album.	○	○	○
Coworkers can **view** the New Products album.	○	○	○
Friends can **add** to the New Products album.	○	○	○
Friends can **view** the New Products album.	○	○	○

13. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Building Jumping

14. Your friend, Josh, sent you some Building Jumping photos and asked you to create an album. What was the privacy policy for the Building Jumping album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Building Jumping album.	○	○	○
Everybody can **view** the Building Jumping album.	○	○	○

Family can **add** to the Building Jumping album.	○	○	○
Family can **view** the Building Jumping album.	○	○	○
Coworkers can **add** to the Building Jumping album.	○	○	○
Coworkers can **view** the Building Jumping album.	○	○	○
Friends can **add** to the Building Jumping album.	○	○	○
Friends can **view** the Building Jumping album.	○	○	○

15. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Safari

16. Your friend Josh Needen asked you to rotate some photos in his Safari album. What was the privacy policy for the Safari album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Safari album.	○	○	○
Everybody can **view** the Safari album.	○	○	○
Family can **add** to the Safari album.	○	○	○
Family can **view** the Safari album.	○	○	○
Coworkers can **add** to the Safari album.	○	○	○
Coworkers can **view** the Safari album.	○	○	○
Friends can **add** to the Safari album.	○	○	○
Friends can **view** the Safari album.	○	○	○

17. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Porch

18. Your friend George Wilson asked you to organize his porch building photos; What was the privacy policy for the Porch Building album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Porch Building album.	○	○	○
Everybody can **view** the Porch Building album.	○	○	○
Family can **add** to the Porch Building album.	○	○	○
Family can **view** the Porch Building album.	○	○	○
Coworkers can **add** to the Porch Building album.	○	○	○
Coworkers can **view** the Porch Building album.	○	○	○
Friends can **add** to the Porch Building album.	○	○	○
Friends can **view** the Porch Building album.	○	○	○

19. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Christmas

20. Your mother asked you to review her Christmas album for errors. What was the privacy policy for the Christmas album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Christmas album.	○	○	○
Everybody can **view** the Christmas album.	○	○	○
Family can **add** to the Christmas album.	○	○	○
Family can **view** the Christmas album.	○	○	○
Coworkers can **add** to the Christmas album.	○	○	○
Coworkers can **view** the Christmas album.	○	○	○
Friends can **add** to the Christmas album.	○	○	○
Friends can **view** the Christmas album.	○	○	○

21. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Baby

22. Your sister asked you to create a new album for her baby photos. What was the privacy policy for the Baby Photo album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Baby Photo album.	○	○	○
Everybody can **view** the Baby Photo album.	○	○	○
Family can **add** to the Baby Photo album.	○	○	○
Family can **view** the Baby Photo album.	○	○	○
Coworkers can **add** to the Baby Photo album.	○	○	○
Coworkers can **view** the Baby Photo album.	○	○	○
Friends can **add** to the Baby Photo album.	○	○	○
Friends can **view** the Baby Photo album.	○	○	○

23. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Misc

24. Your sister asked you to sort some photos from her Misc album to three other albums. What was the privacy policy for the Misc album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Misc album.	○	○	○
Everybody can **view** the Misc album.	○	○	○
Family can **add** to the Misc album.	○	○	○
Family can **view** the Misc album.	○	○	○
Coworkers can **add** to the Misc album.	○	○	○
Coworkers can **view** the Misc album.	○	○	○
Friends can **add** to the Misc album.	○	○	○
Friends can **view** the Misc album.	○	○	○

25. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Confrence

26. In the Global Storage Shared Albums there is an album called Conference. What was the privacy policy for the Conference album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Conference album.	○	○	○
Everybody can **view** the Conference album.	○	○	○
Family can **add** to the Conference album.	○	○	○
Family can **view** the Conference album.	○	○	○
Coworkers can **add** to the Conference album.	○	○	○
Coworkers can **view** the Conference album.	○	○	○
Friends can **add** to the Conference album.	○	○	○
Friends can **view** the Conference album.	○	○	○

27. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

New Desk

28. In your friend George Willson's Albums there is an album called New Desk. What was the privacy policy for the New Desk album when you left it? *

	True	False	Not Sure
Everybody can **add** to the New Desk album.	○	○	○
Everybody can **view** the New Desk album.	○	○	○
Family can **add** to the New Desk album.	○	○	○
Family can **view** the New Desk album.	○	○	○
Coworkers can **add** to the New Desk album.	○	○	○
Coworkers can **view** the New Desk album.	○	○	○
Friends can **add** to the New Desk album.	○	○	○

Friends can **view** the New Desk album.	○	○	○

29. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Sky Diving

30. You have an album called Sky Diving. What was the privacy policy for the Sky Diving album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Sky Diving album.	○	○	○
Everybody can **view** the Sky Diving album.	○	○	○
Family can **add** to the Sky Diving album.	○	○	○
Family can **view** the Sky Diving album.	○	○	○
Coworkers can **add** to the Sky Diving album.	○	○	○
Coworkers can **view** the Sky Diving album.	○	○	○
Friends can **add** to the Sky Diving album.	○	○	○
Friends can **view** the Sky Diving album.	○	○	○

31. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

Angela's Wedding

32. In your mother's (Samantha Jones) albums there is an album called Angela's Wedding. What was the privacy policy for the Angela's Wedding album when you left it? *

	True	False	Not Sure
Everybody can **add** to the Angela's Wedding album.	○	○	○
Everybody can **view** the Angela's Wedding album.	○	○	○
Family can **add** to the Angela's Wedding album.	○	○	○
Family can **view** the Angela's Wedding album.	○	○	○

Coworkers can **add** to the Angela's Wedding album.	○	○	○
Coworkers can **view** the Angela's Wedding album.	○	○	○
Friends can **add** to the Angela's Wedding album.	○	○	○
Friends can **view** the Angela's Wedding album.	○	○	○

33. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
○	○	○	○	○	○

New Page

34. Which of the following photo sharing applications have you used to share photos?

- ☐ Facebook
- ☐ Flickr
- ☐ Kodak
- ☐ Picasa
- ☐ Photie
- ☐ Photobucket
- ☐ Shutterfly
- ☐ SmugMug
- ☐ Webshots
- ☐ Zooomr

35. How often do you upload photos to online photo sharing sites? *

- ○ Daily
- ○ Weekly
- ○ Monthly
- ○ Yearly
- ○ I never upload photographs

36. Which groups do you regularly share photographs with?

- [ ] Co-workers
- [ ] Family
- [ ] Friends
- [ ] Private (visible only to you)
- [ ] Public (visible to anyone on the Internet)

New Page

37. Gender *

- ( ) Male
- ( ) Female

38. What is your age? *

[ ]

39. What is the highest degree you have received? *

- ( ) 12th grade or less
- ( ) Graduated high school or equivalent
- ( ) Some college, no degree
- ( ) Associate degree
- ( ) Bachelor's degree
- ( ) Post-graduate degree

40. What is your occupation? *

- ( ) Administrative Support (e.g., secretary, assistant)
- ( ) Art, Writing, and Journalism (e.g., author, reporter, sculptor)
- ( ) Business, Management, and Financial (e.g., manager, accountant, banker)

- ○ Education (e.g., teacher, professor)
- ○ Legal (e.g., lawyer, law clerk)
- ○ Medical (e.g., doctor, nurse, dentist)
- ○ Science, Engineering, IT professional (e.g., researcher, programmer, IT consultant)
- ○ Service (e.g., retail clerks, server)
- ○ Skilled Labor (e.g., electrician, plumber, carpenter)
- ○ Not currently working/Currently unemployed
- ○ Retired
- ○ Decline to answer
- ○ Student (Please specify area of study) _____
- ○ Other (Please specify) _____

Thank You!

Thank you for taking our survey. Your response is very important to us.

Appendix C

Lab study (study 3)

C.1 Script

C.2 Printed instructions and emails

In the lab study (study 3) participants were given instructions and emails by the researcher on printed sheets of paper. The remainder of this appendix section is all the instructions and emails used in the study. Each instruction or email was printed on its own sheet of paper, but in the interests of saving space, we show only the textual content of the pages. Each box of text was printed on a single page, without the black border.

The pages which give instructions and the pages with emails that initiate tasks were given to all participants. The pages with emails used to prompt the participant, were given to the participant only if the participant did not complete all parts of the task. The first 14 emails were given to the participant in the order they are deplicted here. The remaining emails were presented in a random order.

Instructions

Your name is Pat Jones. You are an administrative assistant at a large company called Global Storage. Global Storage has a company wide photo website, called Gallery, where company related photographs are posted.

Today I will give you emails written on pieces of paper. If you would like to respond to an email just say out loud what you would email back or if you don't want to respond just say "done". Once you have responded I will hand you another piece of paper with the next email.

We are interested in how you approach and solve the issues presented to you. Remember, we are testing the software and how it supports your work with photographs. We are not testing you.

Instructions

Your boss, Gerald, has put you in charge of maintaining the company's online photo website called Gallery. Employees enjoy using this website to share photos amongst themselves and with their family members. Global Storage also uses this website for displaying professional company related photographs.

Many people email you every day asking for you to help them complete photograph management tasks. It is your job to help them but violations of Gerald's photograph policy are not permitted and Gerald has asked you to make any changes necessary to enforce it.

Gerald's photograph policy

1. No photographs containing drugs, alcohol, or anything inappropriate.

2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.

3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.

4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.

5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

To: Pat Jones <pat@globalstorage.com>
From: Charles Taylor <charles@globalstorage.com>
Subject: My new baby

Hi Pat,

Everybody at work has been asking about Dian and my new baby so I took some photographs and posted them to Gallery. Isn't she so cute! Unfortunately, I'm not very good at using Gallery and may have messed a few things up.

The album is called "Charles and Dian's new baby Kerry." The photograph of the card from Dian's mother is sideways and the title has a misspelling that needs to be fixed. Also, I think I accidentally uploaded a photograph of our dog Fido. Could you please delete the photo of Fido?

Thanks,
Charles Taylor

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: People

Hi Pat,

The *People* album has photographs of some of the great people who work here at Global Storage. I think it is wonderful that we have a way to show off some of our employees.

I noticed that Ralf's photograph is in the wrong album. Apparently he put it in "Ralf's Random Photos" album but the prior administrative assistant never moved it to the "People" album. Could you please do so. Also, someone must have thought it would be funny to have a cat as the album cover for the "People" album. Please select some other photograph to be the cover.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Charles' new baby

Hi Pat,

I just saw the photographs of Charles' new baby. I don't know if you noticed but the photograph of the card is sideways. You can tell from the words printed on the card which are sideways.

Please fix it.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Charles' new baby

Hi Pat,

Charles and Dian's new baby is adorable. But the card photograph is titled "Card-ddd" which is not how "Card" is spelled. I'm counting on you to find and fix problems like this in the albums on Gallery.

Please fix it.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Charles' new baby

Hi Pat,

I'm so happy that Charles decided to share the photographs of his new baby with us. However, when I checked the permissions I was disappointed to discover that Everybody on the Internet can see these photographs. I expect you to help employees find and fix problems like this.

Please fix the permissions so they match my policy.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: People

Hi Pat,

Ralph's photograph is still not in the People album. Please move it from the "Ralf's Random Photos" album to the "People" album.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: People

Hi Pat,

The cover of the People album is still a cat. Could you please make the cover be the photograph of Christine.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: What is Gerald's policy?

Hi Pat,

I'm trying to better understand Gerald's policy. Is it ok for me to put the attached photograph from the panel discussion Global Storage hosted on Gallery? If so is there anything I need to make sure to do?

Thanks,
Angela Sebastian

To:	Pat Jones <pat@globalstorage.com>
From:	Angela Sebastian <angela@globalstorage.com>
Subject:	What is Gerald's policy?

Hi Pat,

Sorry to bother you about this again but is it ok for me to put the attached photograph of me trying on wedding dresses on Gallery? If so is there anything I need to make sure to do?

Thanks,
Angela Sebastian

To:	Pat Jones <pat@globalstorage.com>
From:	Angela Sebastian <angela@globalstorage.com>
Subject:	What is Gerald's policy?

Hi Pat,

Last time I bother you about this, I promise. But I also have some photographs from my Bachelorette party. Would the photograph below be ok to post on Gallery?

Thanks,
Angela Sebastian

To:	Pat Jones <pat@globalstorage.com>
From:	Angela Sebastian <angela@globalstorage.com>
Subject:	What is Gerald's policy

Hi Pat,

Does Gerald care what the privacy settings are? Can I just set them up any way I want?

Thanks,
Angela Sebastian

To:	Pat Jones <pat@globalstorage.com>
From:	Angela Sebastian <angela@globalstorage.com>
Subject:	What is Gerald's policy

Hi Pat,

While I was waiting for your email Gerald stopped by and I just asked him what his policy is. I think you may be slightly wrong about what he wants. I've included the policy he told me below.

Thanks,
Angela Sebastian

1. No photographs containing drugs, alcohol, or anything inappropriate.

2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.

3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.

4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.

5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

To: Pat Jones <pat@globalstorage.com>
From: Ralf Jackson <ralf@globalstorage.com>
Subject: Project Competition

Hi Pat,

I was looking at some of our old photographs and I noticed some problems with an album the prior administrative assistant created for me. As you know, Global Storage occasionally hosts college project competitions to help find new talent and to show off how great a company we are.

Could you look through the "Project Competition (2009)" album and fix the errors the last administrative assistant made? All the photographs need to be straight . Also, many of the photographs appear to be duplicates with different titles. Please delete any duplicates.

Thanks,
Ralf

To: Pat Jones <pat@globalstorage.com>
From: Ralf Jackson <ralf@globalstorage.com>
Subject: Funny signs

Hi Pat,

I'm putting together a presentation and I am going to use a bunch of photographs of signs that I've been randomly taking over the last couple of years. I know my random photos album isn't very organized I just like to keep it around so other employees can use some of these random photographs in presentations.

Could you look through the "Ralf's Random Photos" album and move all the photographs of signs to the empty "Funny Signs" album I made?

Thanks,
Ralf

To: Pat Jones <pat@globalstorage.com>
From: Susie Carol <Susie@globalstorage.com>
Subject: New products presentation

Hi Pat,

This last week we had a public show case of our new product line. I created an album entitled "New Products" of all the great photographs I collected from the event.

Could you go through the "New Products" album I just made and clean things up a bit? All the photos need to have titles. You can pick whatever title you think is appropriate. I already went through and organized them so everything is in the correct order.

Thanks,
Susie Carol,
Global Storage Marketing

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Ski trip photos

Hello Pat,

Remember that great ski trip we took together last month? The ski resort photographer finally sent me photographs and they look great. Your friend Daniel looks hilarious fallen over in the snow, I'm sure it is going to take him a while to live that down.

I created an album called "Pat and Ralf's ski trip". Could you please make sure that none of the photos are sideways? Don't worry about changing any of the titles, I already took care of that. Also, can you pick a more exciting cover photograph?

Thanks,
Josh Needam

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Ski Trip

Hi Pat,

I just thought I would check up on how you are doing so I checked the photographs Ralf sent out of your ski trip. I think I need to remind you of my policy about "acceptable" photograph albums.

1. No photographs containing drugs, alcohol, or anything inappropriate.

2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.

3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.

4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.

5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

Thanks,
Gerald
(The Boss)

To:	Pat Jones <pat@globalstorage.com>
From:	Angela Sebastian <angela@globalstorage.com>
Subject:	Conference venue photographs

Hi Pat,

As you know, Global Storage just finished hosting a small conference called BoxTalk and I'm trying to get all the venue photographs posted on Gallery so I can put a link to them on the public website.

You can find the photographs in the "BoxTalk Venue" album. I may have uploaded some of the photographs multiple times so if you see any duplicates feel free to delete them. Also I don't like the current album cover, please select a different photograph and make it the cover. If you see anything else wrong go ahead and fix it.

Thanks,
Angela Sebastian

To:	Pat Jones <pat@globalstorage.com>
From:	Angela Sebastian <angela@globalstorage.com>
Subject:	Conference panel discussion photographs

Hi Pat,

We had a great panel discussion at the BoxTalk conference Global Storage recently hosted. I had Josh take some photographs of the panel discussion which he put on Gallery for me and I would like to put a link to them on the conference forum.

You can find the photographs in the "BoxTalk Panel Discussion" album. Some of the photographs are in the wrong order. The photos of the sandwiches and Jason standing at the podium all need to be at the beginning. The photographs of the panel attendees standing up need to be at the end.

Thanks,
Angela Sebastian

To: Pat Jones <pat@globalstorage.com>
From: Charles Taylor <Charles@globalstorage.com>
Subject: Pirates Game

Hi Pat,

I just realized I still have a great collection of photographs from the trip we took to the Pittsburgh Pirates baseball game that I haven't yet put on the Gallery site. So I thought I had better put them up on the site. Better late than never, right.

I don't actually know who some of these people are so I only titled the people I know. I've put the photographs in an album called "Pittsburgh Pirates". Could you please go and title all the people you recognize?

Thanks,
Charles

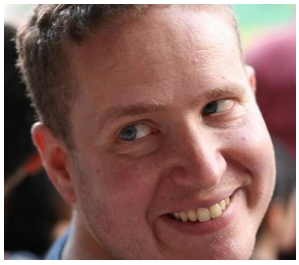
William Barish


Chris Macolm


Cathy Keen

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Cool sculptures in Oregon

Hi Pat,

I recently went on a trip to visit my friends in Oregon. They took me to this great event where contestants make moving sculptures and then race them. They have to both race down the road and successfully peddle them up a sand dune. Some of the sculptures are very inventive.

I uploaded them into an album called "Cool Moving Sculptures" and titled some of them. Could you come up with good titles for the rest? Also could you pick your favorite as the cover?

Thanks,
Josh Needam

To: Pat Jones <pat@globalstorage.com>
From: Susie Carol <susie@globalstorage.com>
Subject: Seattle candlelight parade

Hi Pat,

I don't know if you are aware but every year Seattle has a candlelight parade. All the floats have lights on them and the parade happens after dark. This year Global Storage decided to sponsor a float and I took lots of photographs.

Please help me clean up the "Seattle Candlelight Parade (2011)" album. I took lots of great photographs but I'd rather if this album was all on one page. So please delete your least favorite photographs so that there are no more than 12 photos in this album.

Thanks,
Susie Carol
Global Storage Marketing

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Community service

Hi Pat,

The Global Storage Gives Back official community service day this weekend was a big success. Over half of Global Storage's employees decided to participate by doing everything from helping build houses to cleaning up streets. I volunteered with Habitat For Humanity building a porch on a new house. Susie in Marketing asked me to created an album with all the photographs I took at the event so she can use it to show off how great this company is.

I put all of my photographs in a new album called "Global Storage Gives Back". However, some are sideways, please help me out by turning them around straight. Also, I think I uploaded a bunch of other random photographs into the album by accident. Could you please delete any photographs that don't involve building porches.

Thanks,
Josh

To: Pat Jones <pat@globalstorage.com>
From: Charles Taylor <charles@globalstorage.com>
Subject: Grace's Birthday

Hi Pat,

My daughter Grace just had a birthday and I made sure to photograph the whole event and put the photos in a new album. The cake in particular was very nice looking and I got several shots of that. Ya, I know I made it but that doesn't make it any less awesome.

Please look through the "Grace's Birthday" album and just make sure everything looks ok. I may have gone a bit overboard with photographing the cake, go ahead and pick your favorite(s) and delete the rest.

Thanks,
Charles

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Florence Photographs

Hi Pat,

I just got back from my vacation to Florence, Italy. I took photos while exploring the city and now everyone keeps asking me about all the great sights I saw. So I though I would put together a photo album of Florence.

I put all the photographs in the album "Josh's trip to Florence Italy." Could you please help me out by making up titles for the couple of photographs I couldn't think of good titles for. Also, can you pick your favorite photograph as the album cover? The one I have now is just too generic.

Thanks,
Josh

To: Pat Jones <pat@globalstorage.com>
From: Susie Carol <susie@globalstorage.com>
Subject: Factory Tour

Hi Pat,

In an effort to promote public awareness Global Storage is now offering factory tours at some of our factories. I've taken several photographs from the tour at one of our closer factories and put them on Gallery.

Could you go through the "Factory Tour" album and clean up the titles? All the photographs have titles but some of them have dashes in the middle of the title and I don't want any to have dashes. Also, could you make the photo of people waiting in line be the album cover?

Thanks,
Susie Carol

To: Pat Jones <pat@globalstorage.com>
From: Emma Johnson <emma@globalstorage.com>
Subject: Pumpkin Carving

Hi Pat,

I got a bunch of my girl friends together for some relaxing pumpkin carving fun. We got some great pumpkins that I wanted to show off. I particularly like the one with the witch in the apple (my creation).

Could you please help me fix the order of the photographs in the "Halloween Pumpkin Carving" album? Right now there are photographs of carved pumpkins before the photographs of them being carved. Also, could you make the photo of the pumpkins with the lights out be the album cover?

Thanks,
Emma

Appendix D

Online study (study 4)

D.1 Online survey

Gallery MTurk Questions

Opinion

1. Please indicate if you agree or disagree with the following statements about the **work** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
It was easy to determine if there was an error in the **permissions**. *	○	○	○	○	○
It was easy to determine if there was a **spelling** error in the **title**. *	○	○	○	○	○
It was easy to determine if a photo was **sideways**. *	○	○	○	○	○
It was easy to determine if there was an error in the **tags**. *	○	○	○	○	○

2. Please indicate if you agree or disagree with the following statements about the **home** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I identified and corrected all the **spelling** errors. *	○	○	○	○	○
I identified and corrected all the **rotation** errors. *	○	○	○	○	○
I identified and corrected all the **tag** errors. *	○	○	○	○	○
I identified and corrected all the **permission** errors. *	○	○	○	○	○

3. Please indicate if you agree or disagree with the following statements about the **home** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
It was easy to determine if a photo was **sideways**. *	○	○	○	○	○
It was easy to determine if there was an error in the **permissions**. *	○	○	○	○	○
It was easy to determine if there was an error in the **tags**. *	○	○	○	○	○

It was easy to determine if there was a **spelling** error in the **title**. *	○	○	○	○	○

4. Please indicate if you agree or disagree with the following statements about the **work** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I identified and corrected all the **rotation** errors. *	○	○	○	○	○
I identified and corrected all the **spelling** errors. *	○	○	○	○	○
I identified and corrected all the **permission** errors. *	○	○	○	○	○
I identified and corrected all the **tag** errors. *	○	○	○	○	○

Memory

5. For the **White Water Kayaking album** which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view White Water Kayaking Album	_ can currently view White Water Kayaking Album
Everybody on the internet	☐	☐
Adventure Friends	☐	☐
Animal Shelter	☐	☐
Family	☐	☐
Pat Jones	☐	☐
I don't know	☐	☐

6. For the **Teapots album** which of the following groups would Pat's boss want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Teapots Albums	_ can currently view Teapots Album
Everybody on the internet	☐	☐
Dezig Designers	☐	☐

Innovative Teapots	☐	☐
Purse Central	☐	☐
Starlight Phones	☐	☐
I don't know	☐	☐

7. For the **Bags with Toy album** which of the following groups would Pat's boss want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Bag With Toy Album	_ can currently view Bag With Toy Album
Everybody on the internet	☐	☐
Dezig Designers	☐	☐
Innovative Teapots	☐	☐
Purse Central	☐	☐
Starlight Phones	☐	☐
I don't know	☐	☐

8. For the **Animal Shelter album** which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Animal Shelter Album	_ can currently view Animal Shelter Album
Everybody on the internet	☐	☐
Adventure Friends	☐	☐
Animal Shelter	☐	☐
Family	☐	☐
Pat Jones	☐	☐
I don't know	☐	☐

9. For the **Inspirational Phones album** which of the following groups would Pat's boss want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Inspirational Phones Album	_ can currently view Inspirational Phones Album
Everybody on the internet	☐	☐

Dezig Designers	☐	☐
Innovative Teapots	☐	☐
Purse Central	☐	☐
Starlight Phones	☐	☐
I don't know	☐	☐

10. For the **Family Calendar album** which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Family Calendar Album	_ can currently view Family CalendarAlbum
Everybody on the internet	☐	☐
Adventure Friends	☐	☐
Animal Shelter	☐	☐
Family	☐	☐
Pat Jones	☐	☐
I don't know	☐	☐

Experiences

11. How frequently do you upload and share photographs? *

- ○ A few times a day
- ○ A few times a week
- ○ A few times a month
- ○ A few times a year
- ○ Less than once a year
- ○ Never

12. Which of the following photo sharing sites have you ever used to share photographs? *

- ☐ Flickr
- ☐ Snapfish
- ☐ Photobucket

- [ ] Shutterfly
- [ ] Picasa Web Albums
- [ ] Kodak
- [ ] Phanfare
- [ ] SmugMug
- [ ] Facebook
- [ ] Other

13. Please indicate if you agree or disagree with the following statements *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I am a detail oriented person. *	○	○	○	○	○
I feel confident about my ability to manage the privacy settings on the photo sharing sites I use. *	○	○	○	○	○
I feel confident about my ability to manage tags on the photo sharing sites I use. *	○	○	○	○	○
Most businesses handle the personal information they collect about consumers in a proper and confidential way. *	○	○	○	○	○
I generally notice whether or not a website I am visiting has a privacy policy. *	○	○	○	○	○
I am concerned about threats to my personal privacy online today. *	○	○	○	○	○
I do not care who sees the photos I post online. *	○	○	○	○	○

14. Have you ever had a negative experience after sharing a photograph on a photograph sharing site or a social networking site such as Facebook? *

- ○ Yes
- ○ No

15. Have you ever done any of the following

- [ ] Created a new group and only shared photos with that group.
- [ ] Changed the privacy settings for a specific photo or album.
- [ ] Set privacy settings on a photo sharing site to "Friends Only."
- [ ] Emailed a photo instead of putting it on a sharing site because of privacy concerns.
- [ ] Other

Demographics

16. Gender *

- ○ Male
- ○ Female

17. Age *

18. Select the category that best describes your profession. *

- ○ Accounting / Finance / Banking
- ○ Administration / Clerical / Reception
- ○ Advertisement / PR
- ○ Architecture / Design
- ○ Arts/Leisure / Entertainment
- ○ Beauty / Fashion
- ○ Buying / Purchasing
- ○ Construction
- ○ Consulting
- ○ Customer Service
- ○ Distribution

- Education
- Health Care (Physical & Mental)
- Human resources management
- Management (Senior / Corporate)
- News / Information
- Operations / Logistics
- Planning (Meeting, Events, etc.)
- Production
- Real Estate
- Research
- Restaurant / Food service
- Sales / Marketing
- Science / Technology / Programming
- Social service
- Student
- Other
- N/A - Unemployed / Retired / Homemaker

19. Highest degree obtained *

- 12th grade or less
- Graduated high school or equivalent
- Some college, no degree
- Associate degree
- Bachelor's degree
- Post-graduate degree

20. Household income *

- Less than $25,000

- ○ $25,000 to $34,999
- ○ $35,000 to $49,999
- ○ $50,000 to $74,999
- ○ $75,000 to $99,999
- ○ $100,000 to $124,999
- ○ $125,000 to $149,999
- ○ $150,000 or more

Thank You!

Thank you for completing this study.