Everyday online sharing

Thesis Proposal

Manya Sleeper
Computation, Organizations, and Society
Carnegie Mellon University

July 29, 2015

Abstract. People make a range of everyday decisions about how and whether to share content with different people, across different platforms and services. These sharing decisions can encompass complex preferences and a variety of access-control dimensions. I plan to explore how varied conceptions of privacy, as well as the sharing "ecosystem" created by the platforms and services with which, and people with whom, one shares can shape these sharing decisions. I also plan to look at how the affordances that shape these decisions can fall short, resulting in regret or self-censorship. Drawing on prior and proposed work, I plan to propose and examine improved modalities for sharing on SNSs.

1 Introduction

People make a range of everyday decisions about how and whether to share content with different people, using varied devices and services. These sharing decisions can range from a choice of whether to post a potentially controversial status update on Facebook while waiting for the bus, to deciding how best to share photos with family members and friends from college, to a group decision around how to share documents with collaborators for group editing, to deciding how to share an event invite so that only the intended audience is able to view and share it.

Each of these sharing decisions encompass a range of dimensions, including available access-control channels, settings, the affordances of different mediums, people or relationships involved in the process, context, types of content, need for verification or security, and underlying preferences. Sharing preferences are also seldom binary (i.e., seldom only hinge on willingness to allow people to view or not view a piece of content at a single point of time). Instead, these preferences can range across a spectrum from wanting to allow some mix of full access, partial access, access based on context or purpose, push/pull-based access, or access on request (reactive access) [25, 31, 33, 34, 51].

In some cases people are able to share content in ways that matches their sharing preferences. However, in other situations on-the-ground decisions may fall short of actual preferences at the time of sharing or in the future [6, 51]. This gap between actual preferences and sharing decisions can occur for a variety of reasons. In some cases technology may fall short. For example settings, tools, or other access-control mechanisms may not be available to meet users' desired preferences, or these tools may not be usable enough for regular use. In other cases, users may have underlying goals that come into conflict and negatively impact their sharing preferences. For example a user may wish to garner attention but may also wish to manage their self-presentation or identity [45, 50]. In still other cases users may feel that they are acting in a way that meets their sharing preferences at the time of sharing but, due to changes in life state, knowledge, mood, etc. may later come to regret their sharing behaviors [4,46,52]

When users' sharing behaviors don't meet their (conscious or unconscious) preferences, a range of suboptimal outcomes can occur, including:

- Coping behaviors to address technological shortfalls [28,50]
- Regret at the time of posting or in the future [46,52]
- Threats to identity or presentation-of-self
- Undersharing or self-censorship [15,19,45,50]
- Inefficiency [51]

The proposed dissertation will explore the dynamics of everyday online sharing decisions with a focus on the attributes and modalities associated with content that drive non-binary sharing decisions and will examine how to instrumentalize these attributes and modalities to help people make better (less subpotimal) sharing decisions on social networking sites (SNSs).

2 Thesis statement

This thesis will explore current everyday online sharing decisions, with a focus on attributes that drive these decisions and current shortfalls, and will apply knowledge of these behaviors to suggest and test an improved method for social-network-site-based content sharing.

This thesis will focus on three primary components:

Understand shortfalls of current access-control mechanisms This thesis will draw on user studies to understand when current access-control mechanisms fall short for different platforms and types of content. It will explore the types of content, people, and sharing decisions for which current access-control mechanisms do not currently meet user needs. I will also draw on past work, performed with co-authors, to explore the impact of current access-control mechanisms falling short, including coping mechanisms, self-censorship, and regret. I will use this knowledge to inform potential mechanisms and designs for improving SNS-based sharing mechanisms.

Understand users' everyday online sharing decisions Through a user study I will explore the full range of peoples' everyday online sharing decisions (across platforms, groups of people, types of content, etc.). I will establish an understanding of the range of sharing decisions people face throughout the day as well as factors that lead people to choose different platforms to share different types of content with different groups of people (e.g., affordances, ease of use, settings/access control, trust in the platform, etc.).

Explore the potential of attribute or modality-based improvements to SNS-based sharing decisions Based on user studies I will select several attributes or modalities around which people tend to try to base their content sharing decisions (e.g., interest-based sharing, ability to provide levels of access, etc.). I will use survey, lab, and field studies to evaluate the potential of these attribute or modality-based improvements for SNS-based sharing decisions.

3 Related work

In everyday situations people share different types of content with others using varied platforms and services. Prior work has addressed the importance and challenges of online access control, the impact of the shortfalls of current online access-control mechanisms, and potential improved access-control systems both for general use and specific to SNSs.

3.1 Online access-control is necessary but difficult

Everyday content management, including file management, sharing communications through formats like status updates, and sharing and distributing media like photos and videos, is increasingly moving online. Thus, users with a range of expertise levels rely on online platforms to manage and share content with others with increasing frequency. However, prior work has found that people, even in everyday scenarios, can have complex privacy preferences. Providing usable access-control options for these preferences is challenging; however, when access-control needs are unmet, suboptimal access-control decisions can have negative results.

3.1.1 Online sharing presents a challenging environment

Many people share content online during everyday interactions using a range of services (e.g., email, SNSs, photo-sharing sites, etc.) [18, 34, 42]. However, everyday online sharing presents a challenging environment for designing and providing usable access-control settings.

During everyday interactions users often want to share content, sometimes with one or more people, often to varying degrees [18, 31, 34]. However, desired access control policies can be complex, include exceptions, and be difficult or impossible to implement using existing tools and settings [7, 8].

Additionally, the desire to selectively share is typically in pursuit of an end goal (e.g., publish photos, organize an event, edit a document, etc.) rather than an ends in itself. This desire may be further complicated by users who have benefit-oriented goals (e.g., getting attention for a photo) that are contradictory to access-control or privacy-focused preferences [23, 50].

Online platforms, like SNSs, also have potential to cause "context collapse," which necessitates that users communicate with different audiences in a single context when using a platform [32]. The platforms create an environment of "group co-presence" by combining different social groups on one application that would not otherwise be combined in the users' offline or professional lives [27, 28, 56]. This can create additional challenges as people either create content for an "imagined audience" that may not match the actual audience who can access the content they share, or when users try to create access-control policies to distinguish between the varied groups who have access to the content [1,32].

Users must also cope with managing their privacy and choosing access-control preferences when sharing decisions can depend on other peoples' preferences and actions, as well as the actions of companies and platform providers. This issue can be especially apparent on platforms that include photo sharing and tagging functionality. Users depend on others not to post potentially compromising photos, and rely on a number of social norms and offline strategies to address unwanted photos that are posted online or tagged [10,27].

Sharing in everyday environments, or when using non-professional platforms or tools is further complicated because users' expertise levels, available devices and platform, demographics, and technical backgrounds may vary more than they might in professional environments. Prior work has found that people use different platforms, including Facebook and email, for different purposes [18], and users with different demographic or childhood technical backgrounds, tech savviness levels, knowledge or experience with systems, and access to systems may use both access-control settings as well as general communication channels differently [14,21, 22,31,49]. Additionally, features of the content being shared, such as data size, as well as system affordances including perceived convenience, reliability, privacy/security, and features such as alerts and the ability to comment may impact choice of communication channel for content sharing [9,14,43,55]. These differences may be more controlled in work environments in which devices, applications, formal access-control policies and hierarchies, are provided. However, in more informal, everyday settings, people often use a wide range of devices or tools, resulting in variety of strategies and ad hoc policies. [34].

Thus, users cope with a challenging environment for online sharing. In this context, users also have complex preferences complicating the need for usable, transparent access-control and sharing options.

3.1.2 Varied mental models of privacy

Privacy preferences and perceptions, and corresponding attitudes toward sharing are contextually dependent. Such preferences can vary based on a number of factors, including how the data will be used, social norms, and knowledge [5,29,48]. Nissenbaum's theory of contextual integrity outlines some of this complexity, describing how privacy depends on "norms of information flow" that can vary based on "context not only of place but of politics, convention, and cultural expectation" [41].

Some of this complexity can be illustrated by "mental models," which are peoples' mental representations of privacy-related concepts. People have a wide range of models for thinking about privacy and security-related concepts, ranging from security-related "folk models" related to viruses or hackers [53] to, as Camp describes, five categories of models grounded in real-world analogies: physical, medical, warfare, market, and criminal [13]. Mental models can vary widely, and this variance can be correlated with factors such as expertise [3,12].

3.2 Current access-control decisions can be suboptimal

In this challenging environment, faced with complex preferences for sharing and access-control, sharing decisions can have suboptimal outcomes. In the absence of usable sharing or access-control options, people often fall back on coping strategies to try to achieve access control, which can result in suboptimal or inefficient use of services. Alternatively, users sometimes share content that they come to regret at the time of sharing or in the future [46, 52].

3.2.1 Coping strategies

Available online access-control mechanisms can often be unwieldy, time-consuming, difficult to understand, untrustworthy, or not provide options that meet users' needs.

When they lack usable access-control options people often fall back on informal or ad hoc methods to try to achieve access control. This can be seen in SNS-access-control management. To achieve desired selective sharing, SNS users rely on a variety of informal coping strategies, including using multiple profiles to post content for different groups of people, trusting their friend group to maintain the groups' privacy norms, creating stated group rules around sharing, mentally targeting a particular audience when sharing, or deciding not to think about posting in a specific manner [27, 28, 32, 47]. Users may also sometimes use different services for different purposes, for example using Facebook and email for different types of social interactions [18].

Another common coping strategy on SNSs is deciding not to post, or to self-censor some subset of content. Users choose not to post certain types of content because it might not be appropriate for all audiences who might view it, rather than relying access-control tools to target specific audiences [15, 23, 27, 28, 32, 45].

In some cases coping strategies are successful at preventing undesired access, but result in varied inefficiencies. When users self-censor, for example, it can result in suboptimal use of the service. In a study of Facebook users' self-censored posts, we found that participants would have posted about half of self-censored posts if they had access to access-control tools that let them do so optimally [45].

In other cases these coping strategies are unsuccessful at preventing undesired audiences from viewing the content. This can occur when users' mental models of access control do not match actual access control. For example, access-control transference when users comment on or re-share SNS posts can be non-transparent, which can result in inadvertent or unclear sharing [54]. In home access control, mental models for security and privacy may also not match reality, leading to ineffective access-control or security strategies [34].

3.2.2 Regret

Suboptimal outcomes can also result in regret, either at the time of sharing or at a later time. Prior work has examined failure events, including regrets [36]. In work on regretful messages during conversation, people tended to regret a variety of types of messages including blunders, attacks or criticism, making stereotypical references, expressive or cathartic messages or otherwise revealing too much, lies, or telling someone to behave in a specific way [26]. People tended to associate highly emotional negative states as well as "having a lot on their mind" with saying regretful things [37] and often tried to repair regrets by apologizing [38].

On Facebook, people tended to regret posting content related to potentially sensitive topics like alcohol/drugs, sex, profanity, or religion/politics, as well as negative or argumentative content [52]. Similarly, on Twitter, participants in our study of Twitter regrets tended to regret critical statements, blunders, and tweets that revealed too much. About half of these participants were able to successfully repair their regret, often deleting the tweet, and/or apologizing. However, compared to offline regrets, participants with regrets on Twitter took longer to realize they should regret statements and to repair the regret [46].

Thus, suboptimal online sharing can result in a range of types of regret. Compared to offline contexts, online scenarios present a challenging environment for preventing and addressing such regrets.

3.3 Varied modalities and attributes could improve online access control

Prior work has also focused on increasing the usability of online access control, both generally, and with a focus on selective sharing for SNSs. Work has focused on modalities around which access-control decisions can be based, for example changing the timing of access-control decisions or making access-control decisions attribute-based. Work has also focused on the attributes around which access control could be based, for example, different types of relationships or measures of relationship closeness. In line with these types of modalities and attributes, prior work has also focused on exploring, improving, and automating grouping tools for selectively sharing content.

3.3.1 Tag or attribute-based access control

Prior work has focused on improving the usability of online access control by allowing users to control access to content by defining access-control rules using user-created tags or other attributes.

Klemperer et al. found that tags created naturally by users for photos were viable for access control, and that users found tag-based hypothetical policies usable [25]. In online content-sharing systems, there are a range of potential attributes that can be drawn on for such systems, ranging from user-defined tags to system-defined metadata that users commonly draw on for search and recall like file location, type or format, time of last usage, keywords, or events associated with the content [11].

Several systems have been proposed that offer access-control decisions using tag or attribute-based policies. For example, Au Yeung et al. created a prototype system for creating access-control policies for Flickr using descriptive tags and linked data for photos [57]. Hart et al. created a tag-based system to provide access control for Wordpress, an online blogging system. They found that users were able to create policies more quickly, and with equal accuracy, using the tag-based tool as compared to traditional tools [20]. More broadly, Mazurek et al. proposed a distributed, attribute-based file access-control system that was able to express policies for personas drawn from user studies with low overhead [35].

3.3.2 Grouping tools

Prior work has focused on exploring the groups of people with whom users share content and creating improved tools and interfaces for providing users with automatically-created groups of people with whom to share content.

Several SNSs, including Facebook and Google+, provide manual or partially-automated grouping mechanisms, like Facebook's friend lists or Circles on Google+. Researchers examined the types of groups that emerge on these sites, how well these mechanisms capture the types of groups that emerge from friends present on these sites, and how these mechanisms are typically used.

In a field study of the Circles created by early Google+ users, Kairam et al. found that participants tended to create Circles that reflected either "life facets" like work or school or strong or weak tie strength [23]. Work also found that users used these Google+ Circles for a range of purposes, beyond privacy-based selective sharing, including directing content to appropriate audiences, for appropriateness, relevance, and to try to maximize audience [23, 54].

Prior work has also extensively examined both the use of Facebook's friend lists feature, as well as how well it matches the types of groups people naturally create from their friends in different scenarios. Several studies have found that, when asked to group their Facebook friends, people typically group them into groups that correspond to life-stage or contextual relationships. Kelley et al. found that, when asked to group their Facebook friends using various lab-based methods, groups participants created tended to correspond to school, family, specific locations, and people the participants couldn't identify [24]. De Wolf et al. had similar results for a study of young adults in Belgium, finding that participants tended to categorize their friends according to interest-based categories, geographic-community-based categories, people who knew each other, mutual friends, types of contacts, or personality traits [16]. In a study of categories of SNS-friends in Singapore, Zhang et al. also found that participants tended to describe school, work, interest, and family-based groups, but found some variation by gender, ethnicity, and age [58]. Alternatively, Wiese et al. found that rather than specific groups, self-reported closeness was the strongest predictor of willingness to share various types of information [56].

Group-creation tools face an additional challenge in helping users remember and visualize all the groups, both for use, and to spot inconsistencies and problems. Netter et al. presented a system designed to cluster Facebook friends by "social role" by clustering the friends based on similar permissions [40]. They draw on Reeder's expandable grid interface to usably present the varied permissions [44].

Prior work has also examined creating machine-learning-driven automated grouping tools that use various attributes to algorithmically predict groups for sharing or privacy settings. Although prior work has found that permanent access-control list membership can be difficult to predict using traditional algorithms [17, 39], several systems have used automated or partially-automated interfaces to provide more usable options than traditional interfaces.

For example, Amershi et al. created an interactive automated grouping tool called ReGroup for sharing content on Facebook using seventeen demographic, life stage, interest, and social features. It iteratively learned, and presented, groups to users and suggested additional members and characteristics for filtering the groups. Participants found it more effective for creating large and varied groups than the traditional grouping models [2].

3.3.3 Timing of access control

Another method for improving the effectiveness of access-control tools is by changing when users can set or adjust access-control policies.

Mazurek et al. looked at the usability and utility of allowing users to make reactive, rather than a priori access-control decisions through an experience sampling study. They found that reactive policies can facilitate policies that are contextually-dependent but difficult to define using traditional models. Additionally, many participants preferred reactive, or partially-reactive, systems to traditional systems [34]. Similarly, Bauer et al. looked at a smartphone-based door-unlocking system and found that the ability to reactively provide permission rather than distribute keys a priori helped create policies that better matched user preferences [7].

Work has also looked at the possibility of sharing impermanent content. Ayalon and Toch found that sharing preferences for Facebook content faded over time and changed based on life events, suggesting the potential for impermanent sharing mechanisms on Facebook [4]. However, Bauer et al. found that while some participants wanted the visibility of posts to change over time, they tended not to be able to accurately predict which ones they would prefer would become more private, indicating difficulties in creating a priori fading mechanisms [6].

4 Understanding shortfalls of current access-control mechanisms

To understand user needs for improving online-access-control mechanisms, it is necessary to understand when current mechanisms fall short. In this section I discuss three studies I performed with co-authors, or am in the process of performing, to better understand how and when current online access control falls short, the outcomes of these shortfalls, and possible implications of these shortfalls for designing improved online sharing mechanisms.

4.1 Regrets from Twitter sharing

Oversharing content, or sharing with the wrong audience can lead to regret. Expanding on prior work on regrets on Facebook [52], we performed a survey-based user study to look at the dynamics of regretted content-sharing on Twitter. We focused on how regretted sharing on Twitter compared to conversational, offline regret, drawing on methods used in prior work on conversational regrets [26, 37, 38]. Our goal was to explore how online environments, and, specifically, the affordances provided by Twitter, impacted states that lead to regret, types of regret, methods for determining that one should regret content, and strategies for repairing a regretted message.

In the summer of 2012, we ran a 1,221-participant Mechanical Turk survey of US Twitter users to compare regretted sharing on Twitter to regretted messages in conversation. We asked each participant to fill out a survey in which they described one regret, either on Twitter or from a conversation.

We found that, although many regrets were comparable across Twitter and conversations, some types of regrets were more typical of Twitter, related to the broadcast environment. For example, Twitter regrets tended to be more cathartic/expressive or reveal too much information. Twitter regrets also tended to be targeted at broader audiences than conversational regrets. We also found that, because Twitter participants lacked in-person feedback to their messages, they took a longer time to become aware that they should regret the tweets, and, therefore, took longer to try to take repair actions.

This study is complete and was published at CHI 2013 [46].

4.2 Self-censorship as a coping mechanism

Regret is a possible outcome when one shares sub-optimally. However, another possible suboptimal outcome is sharing with too few people, not sharing on certain topics or in certain styles, or self-censoring content one would benefit from sharing or would like to share in the presence of better access-control mechanisms. We used a qualitative diary and interview-based study (n=18) to explore the types of content participants chose to self-censor on Facebook. We also looked at the subset of this unshared content participants might choose to share in the presence of improved access-control tools.

Participants used SMS messaging, over a one-week period, combined with short, nightly online surveys to tell us about all the content they posted on Facebook as well as every time they thought about posting something on Facebook but then decided not to post. At the end of the one-week period we performed hourlong interviews with each participant, in a lab setting. During the interviews we asked participants for additional details about each of the pieces of content they thought about posting but decided not to post. To explore the potential for improving access-control mechanisms to allow for increased sharing we also probed about desired audiences and whether they would have shared pieces of content if they could have more directly targeted desired audiences.

We found that participants chose not to share personal content, external content (e.g., entertainment, political and news items), and conversational items. Reasons for choosing not to share content included self presentation strategies, for example a desire to seem positive, a desire not to start an argument or discussion, a desire not to offend anyone, a desire not to be boring or repetitive, and posting not being convenient (i.e., time, technology, or location making it difficult to post).

When asked, participants indicated that they would have shared about half the unshared items if they had optimal access-control mechanisms and could have exactly targeted their desired audiences. When they described their desired audiences, these groups included specific individuals, specific groups of people, often people they were close to (e.g., friends, family, people who lived close by), and ambiguous groups of people, who were often context-dependent, for example people interested in a post (known or unknown), people at whom a post was targeted, or people who wouldn't be offended by a post.

This study is complete and was published at CSCW 2013 [45].

4.3 Understanding conceptions of privacy

Access control can also fall short when user understanding of privacy or access control, or related concepts, are misaligned with company or designer conceptions of privacy. Throughout interviews during the self-censorship study, we found that participants tended not to trust Facebook privacy mechanisms because they had low levels of trust or understanding of how the access-control tools work.

There are a range of commonly-used concepts related to privacy that appear in many companies' privacy settings, controls, and privacy policies. To help provide insight into general conceptions of these privacy concepts, across a range of contexts, we performed a large-scale online survey to probe high-level understandings of privacy-related concepts in online and offline contexts.

In winter of 2015 we asked 3,178 US Mechanical Turk workers to react to one of 15 commonly-used privacy-related concepts (e.g., "privacy," "personal information," "using your information to improve your experience," etc.) taken from a review of companies' privacy policies and privacy-settings pages. Participants were given the concept in one of five contexts (using email, using a search engine, using Facebook, using a map application, or checking out at the grocery store). Participants provided their initial reactions, any perceived benefits, any perceived risks, and and any ways they thought they could reduce the perceived risks. Participants also responded to a range of potential risks, benefits, and ways to reduce risks drawn from pilot rounds of the survey, presented as Likert -scale questions.

We found that participants' associations with the terms, across the five contexts, tended to revolve around five high-level dimensions: ads/marketing, uses/benefits of the application, privacy/security threats, types of information collected, and negative emotional reactions. However, the manner in which users thought about the dimensions varied by context. For example, respondents tended to perceive grocery-store ads more positively than online ads, and tended to associate different security and privacy threats, and methods to reduce the threats, with the different contexts. These differences can provide insights into the design of privacy-related settings and notifications in different contexts.

The study is complete and currently under submission for CSCW 2016.

4.4 Implications for design

Regret and self censorship occur on SNSs, often because people overshare or would share more if they had access-control tools that allowed them to better target desired audiences. We also found some attributes or modalities that may guide peoples' decisions for sharing. This provides some initial, potential implications for design.

Metrics for suboptimal outcomes: For access control, both under-sharing (self-censorship) and oversharing (with regret) should be considered metrics for suboptimal outcomes. Participants both regretted overshared content but also under shared when provided with mechanisms that did not meet their needs.

Content-related sharing: While grouping tools may create static, pre-defined groups (e.g., pre-defined friends lists), we saw that many attributes around which participants wanted to base sharing decisions were tied to the content. Thus, improved sharing mechanisms should allow for more ad-hoc, dynamic, context-specific sharing.

Time-of-decision support: Study participants needed fast, usable solutions at the time of sharing. In the Twitter regrets study, participants tended to seek support and share regretful content in highly emotional states or with a lot on their minds. In the self-censorship study participants tended to bypass access-control tools that they felt were too unwieldy or untrustworthy. Thus, access control must be easily available at the time of a sharing choice, as well as highly transparent.

5 Understanding everyday online sharing decisions

These results around the dynamics of sharing decisions, and related shortfalls, on SNSs demonstrate complex attributes and modalities that may underly sharing decisions on SNSs. However, these results are limited to a subset of content-sharing decisions, specifically social-network updates.

People use the online content-sharing ecosystem to share a wide range of content with many different different people, to a range of levels (e.g., full access, partial access, reactive access on request, etc.). To better understand the full space of online sharing decision attributes and modalities it is important to look beyond current-state SNSs to better understand factors that influence the full range of everyday on-the-ground sharing decisions. In this section we describe a currently ongoing study that seeks to better understand the multi-dimensional process encapsulated in everyday sharing decisions.

We are using a combination interview and diary study-based approach to focus on capturing and understanding the full range of participants' online sharing decisions. The study takes part in three parts. In an introductory interview we ask participants to describe their typical online sharing behaviors, with a focus on the services they use, types of content they share, and people they share with, as well as their general reasons for choosing each service. We then use this interview data to create personalized diary studies for each participant, which they use to report their online sharing behavior over a one-week period. This diary-study methodology is roughly based on Lindley et al.'s diary study-based work on Internet use [30]. Finally, after the one-week reporting period, we interview each participant about their sharing behaviors with a focus on understanding factors that drove their specific online sharing decisions as well as times when available access control mechanisms fell short.

In both interviews we focus on understanding the factors that drive participants' online sharing decisions, including the affordances of the different services, communication rhythms, knowledge/experience of the participant and their sharing partner(s), access control options, daily dynamics, habits, trust, security and privacy needs, etc. We will use the results of this study to develop a better understanding of how and why people typically share content, affordances and other factors that influence different types of online sharing decisions, and factors that can make everyday content-sharing decisions suboptimal (e.g., less efficient, cause regret, cause someone to share more broadly or narrowly than desired). We will also use this study's results to refine our list of potential attributes and modalities around which mechanism-based access control decisions could be based.

This study is currently in progress. We plan to finish interviews by mid-July and analysis by mid-August. We plan to submit this study to CHI 2016.

6 Exploring potential improvements to SNS-based sharing decisions

I will use the results of these user studies and surveys (Twitter regrets, self-censorship, everyday sharing) to find a set of attributes and modalities around which people commonly base sharing decisions but that aren't currently captured by SNS interfaces. I will use this to create a set of attributes and modalities that could be used to potentially improve SNS-based sharing decisions. For example, based on initial studies we find that having the option to perform non-binary sharing (i.e., push/pull based sharing), create ad-hoc, temporary, interest-based groups, or provide varied levels of access in environments where this is not easily facilitated might improve sharing. I plan to use two additional studies to explore the potential of these attributes and modalities to improve sharing.

6.1 Lab-based exploration of attributes or modalities

I will first explore these attributes by measuring their ability to increase desired expressiveness and their potential to reduce regret. To do so, I will ask participants, in a lab or survey environment, to go over content shared using a SNS (e.g., Facebook). I will ask them about the intended audience for each post as well as the desired audience. I will then present them with the chosen attribute or modality and explore whether it allowed them more closely match intended audiences with desired audiences.

6.2 Instrumentalization of a method to improve SNS-sharing

Based on the initial study I will chose a promising attribute or modality. I will then instrumentalize it by creating a prototype or partial prototype of the form of sharing. I will test the effectiveness of the prototype through a lab and field-study based approach.

First, in a lab setting I will present participants with the prototype, pull up the participants' Facebook feed for the time period, and ask them about their sharing preferences, actual sharing, and how they would share using the prototype. Second, I will present field study participants with the prototype. I will ask these participants to periodically report any content they share (on any platform) that they would have shared on Facebook if the prototype method was available. I will ask these participants about how the material was shared currently, who they shared it with, and who they would have shared it with on Facebook. I will use a follow-up interview to probe the reasons behind these decisions.

7 Thesis outline

This sections describes an outline of the proposed dissertation.

1. Introduction and motivation
2. Background and related work
    a) Online access control is necessary but difficult
        i. Online sharing is a challenging environment
        ii. Varied mental models of security and privacy
    b) Current access control decisions can be suboptimal
        i. Users resort to coping strategies when given suboptimal options
        ii. Suboptimal online sharing can result in regret

    c) A range of modalities and attributes could potentially improve onion access control

        i. Tag or attribute-based access control

        ii. Improved grouping tools

        iii. Timing of access control

3. Regrets on Twitter from suboptimal sharing

4. Self censorship as a coping mechanism

5. Understanding conceptions of privacy

6. Understanding everyday sharing decisions

7. Exploring potential improvements to SNS-based sharing decisions

    a) Lab-based exploration of attributes and modalities

    b) Instrumentalization of a method to improve SNS-based sharing

8. Conclusion

8 Task list and timeline

The following lists the planned task list for this dissertation, including completed tasks and a planned timeline for the remaining work.

1. Regrets on Twitter study. Complete (CHI 2013)

2. Self-censorship study. Complete (CSCW 2013)

3. Understanding user conceptions of privacy-related concepts study. Complete (Submitted to CSCW 2016)

4. User study on everyday sharing. Target completion date: August 31, 2015

    a) Plan and pilot user study of everyday sharing. Complete

    b) Run user study
    Targeted completion date: July 15, 2015

    c) Analyze and write up results
    Targeted completion date: August 31, 2015

5. User study to explore attributes and modalities to improve SNS sharing. Target completion date: March 1, 2016

    a) Plan and pilot user study to explore attributes and modalities to improve sharing
    Targeted completion date: October 1, 2015

b) Run user study
Targeted completion date: November 1, 2015

c) Analyze and write up results of user study
Targeted completion date: November 15, 2015

d) Plan and pilot method or tool to improve sharing
Targeted completion date: January 1, 2016

e) Run user study to test method or tool
Targeted completion date: February 1, 2016

f) Iterate and improve method or tool
Targeted completion date: March 1, 2016

6. Thesis writing. Targeted completion date: April 1, 2016

References

[1] Acquisti, A., and Gross, R. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Privacy enhancing technologies*, Springer (2006), 36–58.

[2] Amershi, S., Fogarty, J., and Weld, D. Regroup: Interactive machine learning for on-demand group creation in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2012), 21–30.

[3] Asgharpour, F., Liu, D., and Camp, L. J. Mental models of security risks. In *Financial Cryptography and Data Security*. Springer, 2007, 367–377.

[4] Ayalon, O., and Toch, E. Retrospective privacy: Managing longitudinal privacy in online social networks. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, ACM (New York, NY, USA, 2013), 4:1–4:13.

[5] Barkhuus, L. The mismeasurement of privacy: Using contextual integrity to reconsider privacy in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2012), 367–376.

[6] Bauer, L., Cranor, L. F., Komanduri, S., Mazurek, M. L., Reiter, M. K., Sleeper, M., and Ur, B. The post anachronism: The temporal dimension of Facebook privacy. In *Proceedings of the 12th ACM Workshop on Workshop on Privacy in the Electronic Society*, WPES '13, ACM (New York, NY, USA, 2013), 1–12.

[7] Bauer, L., Cranor, L. F., Reeder, R. W., Reiter, M. K., and Vaniea, K. A user study of policy creation in a flexible access-control system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, ACM (New York, NY, USA, 2008), 543–552.

[8] Bauer, L., Cranor, L. F., Reeder, R. W., Reiter, M. K., and Vaniea, K. Real life challenges in access-control management. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2009), 899–908.

[9] Bergman, O., Whittaker, S., and Falk, N. Shared files: The retrieval perspective. *Journal of the Association for Information Science and Technology 65*, 10 (2014), 1949–1963.

[10] Besmer, A., and Lipford, H. Tagged photos. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems - CHI EA '09* (Boston, MA, USA, 2009), 4585.

[11] Blanc-Brude, T., and Scapin, D. L. What do people recall about their documents?: Implications for desktop search tools. In *Proceedings of the 12th international conference on Intelligent user interfaces*, ACM (2007), 102–111.

[12] Bravo-Lillo, C., Cranor, L. F., Downs, J., and Komanduri, S. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security & Privacy 9*, 2 (2011), 18–26.

[13] Camp, L. J. Mental Models of Privacy and Security. SSRN Scholarly Paper ID 922735, Social Science Research Network, Rochester, NY, 2006.

[14] Capra, R., Vardell, E., and Brennan, K. File Synchronization and Sharing: User Practices and Challenges.

[15] Das, S., and Kramer, A. Self-censorship on Facebook. In *ICWSM* (2013).

[16] De Wolf, R., and Pierson, J. Whos my audience again? understanding audience management strategies for designing privacy management technologies. *Telematics and Informatics 31*, 4 (Nov. 2014), 607–616.

[17] Eslami, M., Aleyasen, A., Moghaddam, R. Z., and Karahalios, K. Friend grouping algorithms for online social networks: Preference, bias, and implications. In *Social Informatics*. Springer, 2014, 34–49.

[18] Farnham, S. D., and Churchill, E. F. Faceted identity, faceted lives: Social and technical issues with being yourself online. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, ACM (2011), 359–368.

[19] Ferwerda, B., Schedl, M., and Tkalcic, M. To post or not to post: The effects of persuasive cues and group targeting mechanisms on posting behavior.

[20] Hart, M., Castille, C., Johnson, R., and Stent, A. Usable privacy controls for blogs. In *International Conference on Computational Science and Engineering, 2009. CSE '09*, vol. 4 (Aug. 2009), 401–408.

[21] Hsieh, Y. P. Online social networking skills: The social affordances approach to digital inequality. *First Monday 17*, 4 (2012).

[22] Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D. J., and Kamm, C. The character, functions, and styles of instant messaging in the workplace. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, ACM (2002), 11–20.

[23] Kairam, S., Brzozowski, M., Huffaker, D., and Chi, E. Talking in circles: Selective sharing in Google+.

In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2012), 1065–1074.

[24] Kelley, P. G., Brewer, R., Mayer, Y., Cranor, L. F., and Sadeh, N. An investigation into Facebook friend grouping. In *Human-Computer InteractionINTERACT 2011*. Springer, 2011, 216–233.

[25] Klemperer, P., Liang, Y., Mazurek, M., Sleeper, M., Ur, B., Bauer, L., Cranor, L. F., Gupta, N., and Reiter, M. Tag, you can see it!: Using tags for access control in photo sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2012), 377–386.

[26] Knapp, M. L., Stafford, L., and Daly, J. A. Regrettable messages: Things people wish they hadn't said. *Journal of communication 36*, 4 (1986), 40–58.

[27] Lampinen, A., Lehtinen, V., Lehmuskallio, A., and Tamminen, S. We're in it together: Interpersonal management of disclosure in social network services. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2011), 3217–3226.

[28] Lampinen, A., Tamminen, S., and Oulasvirta, A. All my people right here, right now: Management of group co-presence on a social networking site. In *Proceedings of the ACM 2009 international conference on Supporting group work* (2009), 281–290.

[29] Lin, J., Amini, S., Hong, J. I., Sadeh, N., Lindqvist, J., and Zhang, J. Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy Through Crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, ACM (New York, NY, USA, 2012), 501–510.

[30] Lindley, S. E., Meek, S., Sellen, A., and Harper, R. "It's simply integral to what I do": Enquiries into how the web is weaved into everyday life. In *Proceedings of the 21st International Conference on World Wide Web*, WWW '12, ACM (New York, NY, USA, 2012), 1067–1076.

[31] Litt, E., and Hargittai, E. Smile, snap, and share? A nuanced approach to privacy and online photo-sharing. *Poetics 42* (Feb. 2014), 1–21.

[32] Marwick, A. E., and boyd, d. I tweet honestly, i tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* (July 2010).

[33] Mazurek, M. L., Arsenault, J. P., Bresee, J., Gupta, N., Ion, I., Johns, C., Lee, D., Liang, Y., Olsen, J., Salmon, B., and others. Access control for home data sharing: Attitudes, needs and practices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2010), 645–654.

[34] Mazurek, M. L., Klemperer, P. F., Shay, R., Takabi, H., Bauer, L., and Cranor, L. F. Exploring reactive access control. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2011), 2085–2094.

[35] Mazurek, M. L., Liang, Y., Melicher, W., Sleeper, M., Bauer, L., Ganger, G. R., Gupta, N., and Reiter, M. K. Toward strong, usable access control for shared distributed data. In *FAST* (2014), 89–103.

[36] McLAUGHUN, M., Cody, M., and O'HAIR, H. The management of failure events: Some contextual determinants of accounting behavior. *Human Communication Research 9*, 3 (1983), 208–224.

[37] Meyer, J., and Rothenberg, K. Repairing regretted messages: Effects of emotional state, relationship type, and seriousness of offense. *Communication Research Reports 21*, 4 (2004), 348–356.

[38] Meyer, J. R. Regretted messages: Cognitive antecedents and post hoc reflection. *Journal of Language and Social Psychology* (2011).

[39] Mondal, M., Liu, Y., Viswanath, B., Gummadi, K. P., and Mislove, A. Understanding and specifying social access control lists. In *Symposium on Usable Privacy and Security (SOUPS)* (2014).

[40] Netter, M., Weber, M., Diener, M., and Pernul, G. Visualizing social roles-design and evaluation of a bird's-eye view of social network privacy settings.

[41] Nissenbaum, H. Privacy as contextual integrity. *Washington law review 79*, 1 (2004).

[42] Ojala, J., and Malinen, S. Photo sharing in small groups: Identifying design drivers for desired user experiences. In *Proceeding of the 16th International Academic MindTrek Conference*, MindTrek '12, ACM (New York, NY, USA, 2012), 69–76.

[43] Rader, E. Just email it to me!: Why things get lost in shared file repositories. In *GROUP'07 Doctoral Consortium papers*, ACM (2007), 9.

[44] Reeder, R. W., Bauer, L., Cranor, L. F., Reiter, M. K., Bacon, K., How, K., and Strong, H. Expandable grids for visualizing and authoring computer security policies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2008), 1473–1482.

[45] Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., and Cranor, L. F. The post that wasn't: Exploring self-censorship on Facebook. In *Proceedings of the 2013 conference on Computer supported cooperative work*, ACM (2013), 793–802.

[46] Sleeper, M., Cranshaw, J., Kelley, P. G., Ur, B., Acquisti, A., Cranor, L. F., and Sadeh, N. i read my Twitter the next morning and was astonished: A conversational perspective on Twitter regrets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2013), 3277–3286.

[47] Stutzman, F., and Hartzog, W. Boundary regulation in social media. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, ACM (2012), 769–778.

[48] Ur, B., Leon, P. G., Cranor, L. F., Shay, R., and Wang, Y. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*, ACM (2012), 4.

[49] Van den Berg, P. E., Arentze, T. A., and Timmermans, H. J. New ICTs and social interaction: Modelling communication frequency and communication mode choice. *new media & society 14*, 6 (2012), 987–1003.

[50] Vitak, J., and Kim, J. "You can't block people offline": Examining how Facebook's affordances shape the disclosure process. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work &#38; Social Computing*, CSCW '14, ACM (New York, NY, USA, 2014), 461–474.

[51] Voida, S., Edwards, W. K., Newman, M. W., Grinter, R. E., and Ducheneaut, N. Share and share alike:

Exploring the user interface affordances of file sharing. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM (2006), 221–230.

[52] Wang, Y., Norcie, G., Komanduri, S., Acquisti, A., Leon, P. G., and Cranor, L. F. I regretted the minute i pressed share: A qualitative study of regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, ACM (2011), 10.

[53] Wash, R. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, ACM (2010), 11.

[54] Watson, J., Besmer, A., and Lipford, H. R. +your circles: Sharing behavior on Google+. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, ACM (New York, NY, USA, 2012), 12:1–12:9.

[55] Whalen, T., Toms, E., and Blustein, J. File sharing and group information management. *Personal Information Management: PIM 2008* (2008).

[56] Wiese, J., Kelley, P. G., Cranor, L. F., Dabbish, L., Hong, J. I., and Zimmerman, J. Are you close with me? are you nearby?: Investigating social groups, closeness, and willingness to share. In *Proceedings of the 13th International Conference on Ubiquitous Computing*, UbiComp '11, ACM (New York, NY, USA, 2011), 197–206.

[57] Yeung, C.-m. A., Kagal, L., Gibbins, N., and Shadbolt, N. Providing access control to online photo albums based on tags and linked data. In *AAAI Spring Symposium: Social Semantic Web: Where Web 2.0 Meets Web 3.0* (2009), 9–14.

[58] Zhang, X., Gao, Q., Khoo, C. S. G., and Wu, A. Categories of friends on social networking.