

An Investigation into Facebook Friend Grouping

Patrick Gage Kelley, Robin Brewer,¹ Yael Mayer,²
Lorrie Faith Cranor, and Norman Sadeh

Carnegie Mellon University, University of Maryland,¹ Harvey Mudd College²

Abstract. With increasingly large friend networks, Facebook users may be losing sight of exactly with whom they are sharing content they post to Facebook. When Facebook released a new privacy interface in summer 2010 they simplified privacy controls; however, group-based permissions remain at the core of fine-grained privacy control. In order to use these fine-grained controls, users must be able to accurately and usefully specify friend groups. In a series of 46 semi-structured interviews, we investigated how participants group their online friends using four different grouping methods. Our results show that these different mechanisms alter the strategies and groups that users create, that groups created a priori need further refinement before they can adequately address privacy decisions, and that users are adapting their online behavior to avoid the need to specify groups in the current Facebook interface. We conclude with several recommendations that would allow users improved group-based access control.

Keywords: grouping, online social networks, privacy, access control

1 Introduction

Media organizations present Facebook privacy as a pervasive problem. With Internet shopping behavior disclosed, employees fired for behavior in tagged photographs, locations being shared, and the frequently changing settings described as a “bewildering tangle of options,” the situation seems bleak [3, 16, 13, 7].

A May 2010 *New York Times* infographic detailed more than 170 privacy options that Facebook users could then manipulate. To restrict access to any group more selective than all of one’s Facebook friends, the “Customize” option must be chosen, at which point pre-specified friend lists or individual friends can be explicitly granted or blocked from that type of content. When Facebook released a new privacy interface in summer 2010 they simplified privacy controls; however, the new controls still rely on the Customize option for selective access control.

While the friend list interface and privacy settings continue to change, as of April 2011 Facebook users continue to have the option to control their privacy using friend lists. In order to use this option, users must first create one or more groups of their friends using the Facebook friend list interface. They can then control access to specific posts or to general categories of content by designating specific friend lists to which access should be granted or denied. We seek here

to examine how users categorize their Facebook friends into lists and whether they tend to do so in a way that supports using friend lists for access control. In addition, we seek to understand the impact of grouping mechanism on the types of lists users create. We conducted 46 semi-structured interviews with Facebook users, investigating how participants group their friends using four different friend-grouping mechanisms.

We found that when presented with the task of grouping their friends, nearly all users relied on similar attributes such as place, time, and context to form initial groups. However, we also found that the different grouping mechanisms we tested affected the strategies users employed. More importantly we found that these mechanisms impacted qualities that one would expect to remain consistent, such as size of groups, number of groups, and overlap. This implies that while it remains possible there could be an internal “ground truth,” defined as a user’s set of ideal groups, the grouping mechanism can be used to—and necessarily does—influence the resultant groups. We also found that one-time grouping is not sufficient for privacy control, even in the small number of possible scenarios we focused on. And finally, we found that Facebook users tend to refrain from posting certain sensitive information on the site rather than posting it and using Facebook privacy controls to restrict access. These findings lead to a number of recommendations for designing group-based privacy controls for social networks.

2 Related Work

Social psychologists and cognitive anthropologists have studied how people categorize objects, concepts, and people for decades. Mervis and Rosch provide a summary of studies regarding the categorization of objects, specifically looking at what categories people define, the fuzziness of categorical boundaries, and the human process of classification [15]. Stangor et al. provide an explanation of how people categorize others, specifically people they do not know, based on simple physical attributes (sex, race) [20]. Categorization has also been looked at within the HCI community, where Malone’s early work explored how people categorize the (paper) files at their desks [12]. Malone recognized, and then studied, the pillars, whose apparently-messy desks do have a physical organizational structure. Kwasnik, commenting on her own investigation of categorization of files adds, “The research indicates that we choose categories, at least for concrete objects, that have the most usefulness for the least cognitive effort” [9]. While these foundational studies provide a holistic picture of how humans group, we seek to understand the grouping task specifically in the context of social networks.

Friend networks are composed of people that users meet in distinct contexts and at different life stages. Facebook and most other social network sites combine all of these people into a single group of “friends,” even though in real life users may have different sharing expectations for different friends and groups [2]. Additionally, sharing across the entire network may not be expected, and one study demonstrated that people having misconceptions about the size and extent of the network they are sharing their information with [1].

In a 232-participant survey, Binder et al. found that as networks continue to grow, broadcasting information across different social groups increases “tension” [4]. In a 20-participant online observation study, Lampinen et al. investigated how people currently managed “group co-presence” on Facebook. Their results showed that participants managed the differences in their groups by dividing the platform into separate spaces such as closed Facebook groups and by “using suitable channels of communication,” i.e. directed messages, private pages, and off-Facebook space. They also relied on self-censorship and trusting their friends to be responsible and not share inappropriate data [10].

The Facebook privacy controls allow users to use their friend lists to customize their sharing preferences according to their groups of friends. However, creating friend groups is a secondary task completed by few Facebook users. In a small thesis study, Smart found that none of his 10 participants used friend lists to control their privacy settings [19]. Facebook founder Mark Zuckerberg stated in October 2010 that approximately 5% of Facebook’s users have created even a single friend list [5].

While, most work in interfaces for privacy has not looked directly at online contact management, many of the same general principles, such as those catalogued by Lederer et al. apply [11], for example they specifically call out excessive configuration over action, a common issue in contact management. In Davis and Gutwin’s research on awareness servers, they found that some types and fidelities of information will be more public than others, yet for more sensitive information the specific recipient will greatly impact the disclosure decision [6].

In a recent experiment, Jones and O’Neill created an automatic method to group Facebook users’ friends. They downloaded 15 participants’ entire friend networks—that is, their friends and their friends’ connections—and used a clustering algorithm to group them. Participants also grouped their friends separately using the card sorting software xSort¹ “as if they were grouping them for controlling their privacy on Facebook.” The xSort software displays the names, but not the photos, of each friend. The participant-created groups were used as ground truth to test against the automatically-generated groups. Jones’ results showed that the automatically-generated groups were at best around 70% accurate in comparison to the groups made by the participants [8].

In the final phase of this study, Jones and O’Neill asked participants to select a privacy-sensitive item from their profile and indicate their “willingness to share” this item with each of 100 contacts. These contacts were a stratified sample of the groups they had made earlier. The researchers found that using the human-generated groups 77.8%-90.8% of contacts were correctly granted or denied access to a privacy-sensitive item. This shows even groups created by participants with privacy in mind may not be adequate for access control [8]. Our research goes beyond the study by Jones and O’Neill to explore how grouping mechanisms impact the groups users create and their suitability for access control.

¹ <http://www.xsortapp.com/>

3 Methodology

To understand how Facebook users group their online contacts, we performed semi-structured interviews with 46 participants in our lab. We view this as an exploratory study, consisting of a pre-survey, an open-ended grouping task, an interview exploring a series of common social network sharing behaviors, and a short post-survey. While we gave each participant one of four different mechanisms to group his or her contacts, we did not set out to test whether one was superior to the others; our goal was to illuminate the effect each interface had on grouping behavior, exploring participants' overall opinions and attitudes toward group-based access controls in social networks.

We recruited participants through a university research pool website, flyers around the city, and local Craigslist postings. Participants completed pre-surveys online before coming to our lab. Upon arrival at our lab, they signed consent forms allowing us to audio and video record (through video cameras or screen capture). For each interview, one member of our research team recorded the time and observations throughout the grouping task, including: size and labels of groups, grouping strategies, participant questions, hesitation, group rearranging, and additional notes. Quotes used in the following sections, from both the grouping task and subsequent privacy scenarios were transcribed from recordings after the interview. The post-survey was completed on paper, immediately after the main interview. Participants were paid \$15 for successful completion of the entire task.

3.1 Grouping Mechanisms

Our participants were each assigned to one of four grouping mechanisms:

- **Card sorting (CS)**. Participants were provided with a set of physical paper cards, each with the name and profile picture of one of their Facebook friends. The participants sorted the cards into piles during the grouping process as shown on the left in Figure 1.
- **Grid tagging (GT)**. Participants were provided the names and profile pictures of each of their friends arranged across several sheets of paper in a grid. Participants “grouped” their friends by tagging the pictures with different colored markers as shown on the right in Figure 1. The participants indicated what group they were creating with each new color used. Participants tagged pictures with multiple colors to place people in multiple groups.
- **File hierarchy (FH)**. Participants were provided a folder on a computer containing profile picture files of each of their friends. The picture files were named with the corresponding friend’s name as shown in Figure 2. To group their friends, participants were instructed to create labeled folders and put the appropriate pictures in the folders. By default, friends were displayed as large icons so pictures were visible. Participants could choose whether to use a Mac or Windows computer for this task.



Fig. 1. Left: A participant using the card sorting mechanism points to a group of one friend during the privacy scenarios section of the interview. Right: A participant using the grid tagging mechanism reviews her friends, after having tagged them with colored markers during the privacy scenarios section of the interview.



Fig. 2. Left: Screen capture of possible view options and friend photos in the file hierarchy mechanism. Right: Screen capture of Facebook’s edit friends page, with the friends tab selected.

- **Facebook friend lists interface (FB).** Participants were asked to log in to their Facebook account and categorize their friends by creating standard Facebook friend lists. We directed participants to the “Account – Edit Friends – Friends” screen (see Figure 2, right) that, while unfamiliar to most users, allows users to see if each of their friends had been placed into a friend list. (By default the edit friends view also includes Facebook “Pages” that a user might also have “liked” and can also be placed into friend lists, which is why we directed users to click the Friends tab on the left).

We did not explicitly tell participants that placing participants in multiple groups was something they could do, so as to not bias them towards doing so. However, we provided a way for this to occur in each mechanism, and did affirmatively answer questions regarding grouping when asked.

Through several early pilot tests we explored other interfaces and variations of those above. The xSort tool mentioned above, used by Jones [8], was eliminated because it cannot display pictures and it has no provision for allowing a

friend to be placed in multiple groups. While we could have implemented our own digital card sorting application that met these requirements, we would then have had a condition fundamentally similar to physical card sorting.

We also tested several variations of our two physical methods, including using post-it notes, allowing for a drawable landscape underneath the card piles (for labels, subgroups, and group relationships), and starting with the cards spread on a table instead of beginning in a stack. Ultimately, these variations introduced more problems than they solved and were thus discarded. For example, beginning with the cards spread across a large surface caused a cognitively troubling and unfamiliar search problem for one pilot participant.

For the final study we chose the four mechanisms described above, representing a diverse set of alternatives, each with their own benefits and challenges. Card sorting was the initial motivation for the study, and remained the most intuitive to sort a set of objects. Grid tagging allowed for a folksonomy tagging structure of friends, where the groups were not central, but the contacts were. The file hierarchy was added to most easily allow for subgroups, encouraging nested folders and simple duplication. Finally, the Facebook friend list interface was used as the status quo, a common baseline and the interface participants have most likely previously used.

3.2 Survey & Interview

Potential participants were directed to an online survey as well as a custom Facebook app, that accessed a list of their Facebook friends. We used the survey and friend data to prepare the grouping task before the participant arrived at our lab. Our pre-survey included basic demographics questions, two questions about general Facebook usage, five questions about trust and privacy concerns on Facebook, four questions about current privacy controls, and six questions about Facebook’s privacy settings changes. Participants completed this survey on their own computer when scheduling an appointment to come to our lab.

At their appointment, participants began with the open-ended grouping task. Next we walked participants through a number of specific scenarios. First, we discussed sharing certain pieces of information, including cell phone number, email address, political and religious views, relationship status, and sexual orientation. We also walked participants through a number of thought experiments to consider who they would share content with if they were posting pictures, sharing their location,² inviting friends to a party, or posting a status update. We also explored some privacy-sensitive scenarios, including asking about vacation photos, gossiping about a superior, untagging a photograph, and removing a post from a participant’s wall.

At the conclusion of the interview task, we had participants fill out a short paper survey that asked about their use and frequency of updating friend lists, their use of Facebook’s “limited profile” feature, if they accept friend requests to “just be nice,” and if they enjoyed the grouping task.

² This study was conducted before the release of Facebook Places, so in the study it was posed as a theoretical feature that offered real-time location sharing.

4 Demographics and Survey Responses

64 participants took our online pre-survey. Of those, 46 participants came to our lab and completed the interview and grouping tasks. For the remainder of the paper we will focus on these 46 participants. For the purposes of the paper we will refer to each participant who completed the study by condition/number (for example, Participant GT2 is the second participant in the grid tagging condition). For a summary of basic demographics see Figure 3. Our study included 27 male participants (59%) and 19 female participants. The average age was 25.6 years, younger than the national population, common among Facebook’s population. 13 participants (28%) self reported that they studied or were employed in a technology-related field.




Fig. 3. Basic participant information including: field of study/employment, gender, and age.

Our participants also closely align with usage patterns of active Facebook users, with all but 2% of participants reporting they log in at least several times per week, and 69% reporting they use Facebook multiple times per day. Facebook reports that 50% of its active users log in on any given day. When asked how frequently users “update your status or post content on Facebook,” 63% of users reported that they contribute information (not just consume) at least once per week. Facebook reports an overall average of 130 friends across all users. Our users have an average of 280 friends (median 222.5 friends).

35% of participants agreed with the statement “I trust Facebook with my personal information” and 43% disagreed (22% neutral). As a possible compensation for users’ distrust in Facebook, 41% of our participants agreed that “I don’t worry about Facebook privacy because I have a strong set of privacy rules” (39% neutral, 19% disagree). A striking 87% of users reported agreement with “Once I put information on Facebook, it’s not truly ‘private’ anymore.” 32% of participants agreed that “Facebook changes privacy controls too frequently.” However, 71% of participants believed “Facebook should ask for user input before making changes,” and 96% agreed that “Facebook should announce any planned changes in advance.” Nonetheless, 83% of participants agreed that they did not read Facebook’s new privacy policy in detail.

From our post-survey at the end of the study we found that, while many might assume the task of sorting all of one’s Facebook friends to be a daunting one, 76% of participants agreed or strongly agreed that they enjoyed the




Fig. 4. Pre-survey responses from our 46 interview participants. Each of the above statements was asked with a 5-Likert scale from 1(Strongly agree) to 5(Strongly disagree).

study, with the rest expressing neutrality. Many participants with several hundred friends stated they enjoyed the study, suggesting people may be willing to group their friends, if the correct incentives are in place.

Of the 46 participants, 17 participants (37%) reported using friend lists in their Facebook accounts. However, only 8 of these people stated that they used these lists to control their privacy settings. In addition, only 9 participants (20%) stated that they try to do monthly maintenance of their friend lists. While these numbers seem to suggest our participants are biased to be more concerned about privacy, these numbers are self reported (we did not check for list usage nor maintenance) we also note that while Facebook normally gives figures distinguishing between active and all users, the 5% figure for friend list usage was not clearly associated with either group, and may have been taken from total usage to make friend lists seem less popular (as this quote is taken from an event with a product intended to replace friend lists as antiquated and useless).

5 Common grouping strategies

Many of our participants when tasked with the intentionally open-ended prompt “using (the cards/the markers/interface), categorize your friends into different groups and tell us what you’re thinking as you do so,” began instinctively and immediately sorting through and categorizing their friends, while others spent

Participant overview

#	Gender	Age	Num. friends		Style	Overlap	secs. per friend	#	Gender	Age	Num. groups		Style	Overlap	secs. per friend
			Age	friends							Age	groups			
Card sorting															
CS1	Female	21	593	8	friend	no	2.10	FB1	Male	22	227	2	group	no	2.29
CS2	Male	24	126	7	friend	no	4.20	FB2	Male	27	169	3	group	overlap	4.37
CS3	Female	21	436	9	friend	no	2.07	FB3	Female	20	229	6	group	overlap	n/a
CS4	Male	32	97	5	friend	no	2.62	FB4	Male	20	428	6	group	overlap	n/a
CS5	Male	22	215	8	friend	no	n/a	FB5	Male	29	152	4	group	overlap	4.13
CS6	Male	23	658	5	friend	no	1.85	FB6	Female	21	407	7	group	no	n/a
CS7	Male	26	99	3	friend	no	n/a	FB7	Male	25	214	5	group	overlap	4.54
CS8	Male	30	212	9	friend	no	2.59	FB8	Female	21	375	2	group	no	0.50
CS9	Male	26	156	6	friend	no	3.67	FB9	Male	21	259	4	group	no	2.08
CS10	Male	28	54	5	friend	no	1.96	FB10	Male	22	395	2	group	no	0.95
CS11	Female	27	6	3	friend	no	1.50	FB11	Male	22	366	5	group	overlap	n/a
CS12	Female	37	274	14	friend	no	n/a	FB12	Male	19	889	11	friend	overlap	6.05
Grid tagging															
GT1	Male	24	108	5	group	no	6.12	FH1	Male	23	145	5	group	no	5.59
GT2	Female	25	218	7	group	no	5.31	FH2	Female	27	149	5	friend	no	5.01
GT3	Male	22	671	30	friend	overlap	5.27	FH3	Female	24	379	6	friend	no	n/a
GT4	Female	19	705	7	friend	no	2.90	FH4	Male	26	269	5	group	no	4.08
GT5	Male	25	391	9	group	overlap	6.47	FH5	Male	55	12	3	friend	no	12.42
GT6	Female	26	62	7	group	no	9.71	FH6	Female	24	206	3	group	no	4.08
GT7	Male	30	164	4	friend	no	3.16	FH7	Female	22	155	13	friend	overlap	5.30
GT8	Female	37	255	8	friend	overlap	n/a	FH8	Male	25	422	8	group	no	4.76
GT9	Male	29	21	4	group	no	9.90	FH9	Female	40	69	4	friend	no	n/a
GT10	Female	24	174	5	group	overlap	4.94	FH10	Female	18	599	2	group	no	1.30
GT11	Female	25	245	11	friend	no	4.16	FH11	Male	23	442	7	friend	no	n/a

Table 1. Overview of our 46 survey participants, including the mechanism they were assigned, the number of friends they had, the number of groups they made, the strategy they used to first categorize their friends, and whether or not their original groups included overlap. Timing information is given in seconds per friend, and was not available for all participants.

time looking at their friends before deciding how to proceed. FB10 began the task stating: “There are all sorts of different ways to categorize the people you know.” GT3 said “Oh gosh, like I kind of have to look at all of them before I start a group.” Some wanted further details. GT10 asked, “Is this open-ended?” FH2 asked, “By different characteristics, or by female, male?” FB1, at a loss for what to do, asked if we had any recommendations for groups he could make.

We observed two common strategies for performing the original grouping task. The first, which we refer to as “by friend,” consists of a participant examining a single friend and placing him/her into a group, creating new groups as needed, and continuing until all friends had been grouped. The second, which we refer to as “by group,” consisted of participants creating a group, and then searching for friends or systematically looking through their list/stack of friends and placing them in that group if applicable. They continued this for each group they could come up with, until they decided they had finished. The strategy each participant used is displayed in Table 1. These strategies are similar to those used by Malone’s file organizing participants [12].

Participants grouped their friends largely based on when or where they met them, or some specific context. Each of the groups was given a label during the grouping task (the interviewer prompted the participant for a group label or description if it was not provided). Each of these labels was coded into general groups, summarized in Table 2. The most used type of label was a description of friends from school/college, which 41 participants (89%) had. Of these, 35 participants (76%) specifically referred to a group of university friends. Family-related labels were also widely used with 29 participants (63%) creating one or more family groups. 14 (30%) participants created a group that specifically referred to a location other than a school or university, such as “Met in New York.” Groups of people the participants could not identify, didn’t know, or described as random were created by 14 participants, with an additional 10 participants creating groups called “other” or “miscellaneous.”

Jones and O’Neill do not give a full breakdown of the labeled clusters in their paper, however they also reported commonly-created groups of close or best friends, groups of weak ties such as acquaintances, “people I hardly know,” or again “random” people. They continue to describe groups based on organizational boundaries (work, school) and specific times and places in their participants lives, all of which we also saw [8]. When comparing our clusters to those typically found in the social sciences research, for example [14], we see fewer family groups, which are almost mandatory, but we hypothesize that due to not everyone having family members present on Facebook, we do not see the same prevalence in our results. We do see that many of the common expected clusters are present, including work, religious affiliations, hobbies, school, childhood friends, and groups of friends met through other friends or relatives.

Our 46 participants created 287 groups (6.5 on average), ranging from four participants who made only 2 groups, up to GT3 who made 30 groups. The groups ranged in size from several single-member groups, to FB12’s group of 394 high school friends. As expected, the more friends a participant had, the longer it took for them to group their friends.

Group Labels

Cluster	Total uses	% of users	Example labels
General friends	85	76%	
Location based	23	30%	Friends from China, France, Met in New York
Generic friends	16	35%	General, Good friends, Friends
Close friends	16	33%	Very good friends, Close set of friends, Best friends
Friends of friends	14	24%	Friends of my boyfriends, Friends' friends
College	79	76%	
General College	38	76%	College, Undergrad friends, University friends
Club or group	24	33%	Dancing club, Crew team, Band people, Fraternity, Freshman hall, Gaming/clan, Swimmers
Other education	40	70%	
High school	27	59%	HS friends, High school, Schoolmates
Grade school	9	11%	Middle school and below, Elementary, Junior high
Family	35	63%	Family, Cousins, Siblings, Relatives
Work	15	33%	Work, Work mates, People I've worked with, Office
Church	5	9%	Church, Bellefield church friends
Don't know	17	30%	People I hardly know, Never met in person, Unidentifiable, Random people who friended me

Table 2. Our participants created and named a total of 287 groups. Afterwards we classified the group names into the above clusters, many of which contained subclusters.

6 Mechanisms impact groups

We observed that the mechanism that a user was given to group their friends impacted the resultant groups. While there were some expected differences (such as time taken to group), we also observed differences in the composition of the groups, specifically with friends who belonged to multiple groups, and also in the strategies used to create the groups.

6.1 Mechanisms

We briefly describe the trends we saw in each of the four mechanisms.

Card sorting. All participants in the card-sorting group used an identical strategy of beginning at the top of the pile of cards and working their way through it in the order it was given to them (random). For each friend, they decided on a group to place them into, placed the card on the table in the area they had designated for that group, and continued. They created new groups only when they got to a card for a person who didn't fit into the groups that had already been created. Each card-sorting participant used the by-friend approach.

All twelve card-sorting participants placed each of their friends into exactly one group. One-third of our card sorting participants created single-person groups when there was not a better fitting group. Many participants seemed to hesitate with certain people, in some cases because they admitted they could

not figure out who that person was, in other cases because they had a person who crossed multiple places, times, and contexts in their life. CS1, when asked if there was overlap between her high school and college groups said: “Yeah, a number of them. I put them in the high school group, because that is where I first met them.” None of the participants in this condition placed similar people between two groups, or “duplicated” the card, although there were blank cards and pens within sight.

Grid tagging. The grid-tagging condition did not lead to a common grouping strategy among participants. Some participants followed an approach similar to that used in the card-sorting condition, where each person was assigned a group before moving on to the next. Others attempted to list all of the groups they believed they would need before placing any friends, and others completed a page at a time, scanning the 20 friends per page before moving on, an artificial limit that our mechanism imposed.

Many participants did not place some of their friends into any group. 4 of our 11 participants in the grid-tagging condition had multiple friends who received more than one tag. While overlap did occur, many participants questioned if it was permissible. GT10 specifically asked, “Can you put someone into more than one category?”

File hierarchy. While the main intent of the file-hierarchy condition was to provide participants with an easy and familiar method to create subgroups and organize friends into a hierarchical structure, we found participants did not take advantage of this. We did not explicitly say or demonstrate that subgroups were possible, as with all the other conditions. In some of the follow-up privacy scenarios participants did create subgroups, but even these events were extremely infrequent. More commonly participants would navigate into the groups, select or point at the users that were relevant to the given privacy scenario, and then navigate away. We did have some participants in the other conditions also try to use “subgroups” by creating groups that they later combined in response to certain scenarios.

Facebook friend list interface. The Facebook friend list interface was flexible enough that our participants used multiple strategies. Some participants created each of their ideal lists first before placing participants, others created lists as they went. Participants grouped both by group or by friend.

The edit friends page, where we had users begin, was not necessarily their preferred interface on Facebook. FB8 used the standard interface instead of the edit friends page, requiring her to recall and then search for friends, instead of paging through her entire network on the edit friends page (50 friends per page).

FB11 asked if duplicating members was allowed, and then added he had “probably accidentally done that,” as if it were a mistake (or possibly was prohibited in the study). Sometimes users were aware that there were overlapping groups. The number of friends placed in multiple groups may be particularly high for the Facebook interface because when adding a friend to a list, that friend’s

current memberships are not shown, and a user may not remember where a given friend has been placed.

Although the Facebook interface was generally faster, some users spent considerable time with it. For example, FB12, with 889 friends, spent time reviewing many of his contacts' profile pages to determine who they were. The Facebook interface was the only mechanism we tested that allowed participants to click through to view full profile pages, a frequently requested feature across the other conditions as many participants were unable to identify friends by just photograph and name.

6.2 Mechanism comparisons

In general, the participants using the card-sorting method were much faster than the grid-tagging and file-hierarchy methods (average times in seconds per friend, card-sorting: 2.52, grid-tagging: 5.79, file-hierarchy: 5.32). The Facebook interface seems to be approximately as fast as card-sorting (average 3.11 seconds per friend), but many participants were not thorough and left many friends unplaced. Additionally, verifying all Facebook friends have all been placed is considerably more difficult using this mechanism than with the other three. Participants took on average 16.8 minutes to group their friends, yet, ranged in time from 9 seconds for CS11 to group her 6 friends into 3 groups, to the 89 minutes FB12 took to group his 889 friends. While these times give us an estimate of how long grouping takes, our participants were also told to think aloud, which can dramatically affect times. For statistical comparisons between mechanisms, non-think-aloud, larger-scale tests would need to be conducted.

Just under half of the participants in the grid tagging and the Facebook interfaces placed at least one friend in multiple groups, while only one file hierarchy participant duplicated a friend, and no one in the card sorting method did. Multiple placement, if used in future interfaces, may create rule conflicts, but it may remove decisions such as whether a friend really belongs in a high school or college group. To create an interface that aligns well with people's internal feelings about their friends, allowing multiple groups seems natural. However, even for the methods where group overlapping occurred, participants often asked the researchers if they could place friends into more than one group before they did it. So, although users may think of their friends as occurring in multiple groups, the interface must make it obvious that this is possible and encouraged.

7 One-time grouping is insufficient for privacy

Immediately after the completion of the grouping task, we walked our participants through a series of Facebook information-sharing and related scenarios. We found that the groups that they had just created were wholly ineffective for meeting participants' privacy expectations in these scenarios. Many participants avoided using the groups they had created until prompted by the interviewer. Often when participants tried to use these groups they found them insufficient.

Even though participants had just spent an average of 17 minutes grouping their friends, nearly all of our participants attempted to work through the privacy scenarios without the use of the groups on the table or screen in front of them,

opting instead to describe how an item of information would be shared with everyone or not placed on Facebook at all. Interviewer prompting was required to have the participants think in terms of the groups, and additional encouragement was needed to have the participants select certain groups they would or would not share items with.

For specific scenarios, such as hosting a party or event, many participants felt they would need to invite people on a case-by-case basis, and often did not consider the groups they created as a helpful starting point. CS4, when considering who to invite to a party said, “I would want a selective way of doing that, not just a group message.” After CS4 eventually suggested a group that might work and the interviewer asked if everyone from that group could come, he added “No, no, not all of them.”

On the other hand, in some scenarios participants stated they would provide the same access to multiple groups. FH2 responded to one question by allowing full access to four separate groups: “These four folders are OK, I have met them, talked with them.”

Additionally, for certain types of content sharing, such as religious and political views, participants needed to create completely new groups or tags to correctly identify a set of friends that they would share these more specific types of information with.

We set up the grouping task, as mentioned earlier, without explicitly mentioning privacy or sharing, just as Facebooks friend lists currently are: under a vague Edit Friends settings page, without any explicit reason for their existence. However even doing this, we saw what we believe are the types of groups people will make even if prompted for privacy (Jones prompted users for privacy and the groups users created share many similarities to the groups in our study [8]). Until we have more information on what an ideal privacy/sharing solution is on Facebook, and a real-world study testing dissemination through groups over time, it is difficult to say which grouping-mechanism is the best, and therefore we attempted to gain an understanding of the difference between mechanisms, while also investigating how people already deal with privacy management on Facebook (without skewing the grouping task).

When participants did decide to indulge the interviewer and think about our specific privacy questions in terms of groups, the groups quickly proved insufficient. When a participant thoroughly reviewed a group, the group was frequently cleaved in two, or often required a few special cases to be removed before it would be suitable to share a specific photo gallery or vacation information with. When asked who GT4 would not have wanted to share her vacation information with she said, “Oh, my cousins.” When we asked if this was all or some of her cousins, she stated, “just a few” and proceeded to highlight which cousins should not have access to this information. A family group alone, which many users commonly create, could never have captured this specific event.

Additionally, participants often had difficulty relocating a particular participant, occasionally not being sure which group they placed them into a few minutes earlier. FH9 described not being able to locate a friend she had just placed: “This Teresa, oh wait, she might be in classmates. Or did I put her in acquaintances, maybe I put her in acquaintances....That’s weird. I don’t see her name.... I can’t find Teresa, this is so weird, it is like she disappeared.”

8 Privacy-related observations

Many participants appear to have developed strategies for protecting their privacy on Facebook without using groups..

For example, many of our participants seem to believe that regardless of the Facebook privacy settings, things they are putting on Facebook may be seen by anyone, and will thus only post information they consider public. People believed that they were able to use their discretion and not post any compromising pieces of information or anything that could be misconstrued as being compromising. When asked about possibly sensitive topics such as political and religious views, some participants stated that it didn't matter who saw these, and thus they were suitable for Facebook, while others said they would never put these topics on Facebook. This furthers the argument that users are unable or unwilling to utilize the full extent of Facebook's privacy settings.

Similar to previous studies [17], our participants varied in what they were comfortable posting. FH2 said of her vacation photographs: "Those pictures, they are just pictures, there is nothing about privacy." CS12 said of her cell phone number: "I let all my friends have it... if I am friends with them on Facebook I feel like they can talk to me anytime." This is further evidence that controls must be flexible enough to allow for users to specify desires that match their own personal preferences.

GT11 created two groups she explicitly had specific sharing requirements for, a group called "ex-boyfriends" and another created for a past stalker. Many other participants created groups of friends with whom they were not close, with labels such as "don't know very well" or "people I've grown apart from." In total, 11 of our participants made a group that we coded as "distant friends."

Even less closely related than distant friends, groups of people our participants could not identify, didn't know, or were described as "random people" were created by 14 of our subjects. As we discussed earlier, profile photographs significantly benefited users (names alone were often not enough). However, Facebook profile links were also desired (and used in the Facebook friend list condition), and without them some users admitted to being unable to identify a small percentage of their friends. GT10 said, "I don't know who this woman is—she just friended me, I replied. I can't remember who she was." FH2 stated of one friend: "I don't even know this person." GT4 said of her "random" group: "The random people I don't even remember where and who they are." CS1 had a group she described as: "These are people I cannot remember for the life of me."

Many participants simply believed they had too many friends. FB10 expressed this best, just a few minutes into the original grouping saying, "I should have cleaned out my friends list before I came here." GT3's comments were similar: "Wow I wish I didn't have so many friends." In general, participants had friends they didn't recognize, didn't know, and didn't want to expend the effort into grouping. Based on the post-survey 19 (41%) of our participants said they sometimes, often, or very often accept friend requests, "just to be nice."

In addition, participants were asked if they had any friends who sent too many application invites. If they answered yes, they were asked if they would

consider placing these friends in a group to block them from sending application invites. What we found was that most would not do this and would continue to block the applications themselves. They didn't want to censor their friends in any way, as one user put it. This shows that users find applications themselves to be annoying, rather than the people who use the applications.

We also asked participants whether they had untagged themselves from photos. Most often, participants would tell us they did so because the picture was unflattering or perhaps inappropriate. This begs the question, what do people consider inappropriate enough that they would restrict others from seeing it? And more importantly, are there still people they would like to grant permission to see that content through privacy settings?

9 Recommendations

Our findings suggest some recommendations for designers of contact-grouping mechanisms in online social networks. These recommendations are not to be taken as a hard and fast list of requirements, but rather to provide a direction for improving interfaces to address real-world access-control needs. Our analysis has focused on privacy and security decisions and the actual content and sharing decisions that our participants make on Facebook.

Provide ad-hoc grouping. One of our key findings above was that a priori user-created groups do not meet users' needs for use as access-control groups. While these groups provide a starting point for users to rationalize about their sharing decisions and understand the space of their online friend network, without allowing ad-hoc modifications they fail. We recommend that contact-grouping interfaces allow for usable, in-situ, temporary group modifications. This would allow a user preparing to post content they are aware is privacy-sensitive to select some number of pre-created groups, briefly review the summed memberships, and remove or add misaligned contacts.

Tagging and filtering. In addition to ad-hoc grouping controls, it seems that because many groups are based on specific contexts, contacts overlapping between groups could be common. For example, Family, Church, and Hometown might all share certain members. To allow for this model without requiring the user to necessarily construct groups, the interface could allow searching based on information Facebook already has, such as friends with a similar current city, or contacts with the same high school or university listed. These filters could then be constructed into permanent groups (with user modifications made).

Additionally, we recommend that users be able to tag their friends with other freeform terms that could be further filtered or searched across. These tags could be similar in nature to Facebook's own friend relationships once requested upon friending; however, they should be private and not require external approval. Ideally these could also not only be added through a privacy/friends management interface, but also directly through a drop-down while viewing a friend's profile.

Assistance through automation. Automated approaches for helping users create and modify groups could be enormously useful. In addition, approaches that would suggest access-control settings or remind users that they might want to restrict access to sensitive information could be helpful. These are areas where additional research is needed.

Holistic view. Previous work has suggested that access-control interfaces should contain a “holistic view” that allows a high-level, single-screen look at the entire policy [18]. Currently, the closest Facebook comes to that is the previously mentioned edit friends screen, which is still paginated and therefore does not actually provide a single complete view (nor does it scale particularly well). This is one of the greatest strengths of the physical card sorting mechanism: every participant was able to know that they had placed each of their friends, while understanding the relative sizes and relations of each group to the others. At the end of the grouping task, the table became their holistic grouping policy, which they could easily point to, gesture at, and modify throughout the privacy scenarios.

Iterative management tools. While our grouping task occurred for participants in a single session in our laboratory, the makeup of online social networks, and members’ relations, change over time. While we have little understanding how users would like to and should modify their groups over time, we do know that the current interfaces do not assist this process. Indeed, nearly all (80%) of our subjects reported that they have either never or only once updated their friend lists.

Grouping is not the primary task. While we believe all of these recommendations will lead toward better interface design, we want to emphasize that the act of friend grouping is not a primary task. Friend groups should be defined to assist users in managing their communications, content sharing, and general privacy on social networks. If the groups do not accomplish this task, then the privacy they seem to create is simply an illusion.

References

1. A. Acquisti and R. Gross. Imagined communities: Awareness, information sharing, and privacy on the facebook. *Privacy Enhancing Technology Symposium*, 2006.
2. P. Adams. The real life social network v2. <http://www.slideshare.net/padday/the-real-life-social-network-v2>.
3. S. Berteau. Facebook’s Misrepresentation of Beacon’s Threat to Privacy: Tracking users who opt out or are not logged in., November 29, 2007. <http://community.ca.com/blogs/securityadvisor/archive/2007/11/29/facebook-s-misrepresentation-of-beacon-s-threat-to-privacy-tracking-users-who-opt-out-or-are-not-logged-in.aspx>.
4. J. Binder, A. Howes, and A. Sutcliffe. Conflicting social spheres and experienced tension in social networking sites. *CHI*, 2009.

5. A. Carr. Facebook's New Groups, Dashboards, and Downloads Explained, October 6, 2010. <http://www.fastcompany.com/1693443/facebook-big-announcements-dashboards-personal-information-downloads-friend-group-lists>.
6. S. Davis and C. Gutwin. Using relationship to control disclosure in awareness servers. In *Proceedings of Graphics Interface 2005*, GI '05, pages 145–152, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2005. Canadian Human-Computer Communications Society.
7. G. Gates. Facebook Privacy: A Bewildering Tangle of Options, September 18, 2010. <http://www.nytimes.com/interactive/2010/05/12/business/facebook-privacy.html>.
8. S. Jones and E. O'Neil. Feasibility of structural network clustering for group-based privacy control in social networks. *SOUPS*, 2010.
9. B. Kwasnik. How a personal document's intended use or purpose affects its classification in an office. In *Proceedings of the 12th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '89, pages 207–210, New York, NY, USA, 1989. ACM.
10. A. Lampinen, S. Tamminen, and A. Oulasvirta. All my people right here, right now: Management of group co-presence on a social networking site. *International Conference on Supporting Group Work*, 2009.
11. S. Lederer, I. Hong, K. Dey, and A. Landay. Personal privacy through understanding and action: five pitfalls for designers. *Personal Ubiquitous Comput.*, 8:440–454, November 2004.
12. T. W. Malone. How do people organize their desks?: Implications for the design of office information systems. *ACM Trans. Inf. Syst.*, 1:99–112, January 1983.
13. R. Marsden. Facebook unveils its latest trick but the implications are worrying, September 18, 2010. <http://www.independent.co.uk/life-style/gadgets-and-tech/news/facebook-unveils-its-latest-trick-ndash-but-the-implications-are-worrying-2082708.html>.
14. C. Mccarty. Structure in personal networks. *Journal of Social Structure*, 3, 2002.
15. C. Mervis and E. Rosch. Categorization of natural objects. *Annual Review of Psychology*, 32:89–115, 1981.
16. A. Ostrow. FACEBOOK FIRED: 8% of US Companies Have Sacked Social Media Miscreants, August 10, 2009. <http://mashable.com/2009/08/10/social-media-misuse/>.
17. S. Preibusch. W3c workshop on privacy for advanced web apis. *Journal of Interactive Advertising*, 1213 July 2010. <http://www.w3.org/2010/api-privacy-ws/slides/preibusch.pdf>.
18. R. Reeder, L. Bauer, L. F. Cranor, M. Reiter, K. Bacon, K. How, and H. Strong. Expandable grids for visualizing and authoring computer security policies. *CHI*, 2008.
19. D. Smart. How do people limit their behaviour in response to their perception of potential social tension? a study of facebook. *Manchester Business School Bsc Interactive Systems Design*, 2009. <http://intranet.cs.man.ac.uk/Intranet.subweb/library/3yrep/2009/7024300.pdf>.
20. C. Stangor, L. Lynch, C. Duan, and B. Glass. Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology*, 62(2):207 – 218, 1992.