

Matt Soave
HCDE 505 Autumn 2013
November 18, 2013

Annotated Bibliography

Below is a selection of reading materials used in conducting research for our project about enhancing social network privacy options.

Binder, Jens, Andrew Howes, and Alistair Sutcliffe. "The Problem of Conflicting Social Spheres: Effects of Network Structure on Experienced Tension in Social Network Sites." *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (2009)*: 965-74. *ACM Digital Library*. Web. 15 Nov. 2013.

This article investigates the tension caused by the "problem of conflicting social spheres," which the authors describe as a disparity between communication's intended audience and its availability to another, unintended audience. They find that increased network diversity led to increased tension, as did a higher number of family connections.

The authors distributed an online survey in which they asked about participants' Facebook to estimate the number of individuals within 16 different social groups, focusing particularly on family, work, and social connections; they used these classifications to quantify network diversity. In order to measure Facebook-borne tension, they asked users to indicate how often they had experienced negative interactions on Facebook. The authors also probed about other behavioral aspects of Facebook use.

The authors found that a more diverse network (*i.e.* more types of groups present) did indeed lead to more tension, especially as caused by the interaction between social spheres that would be kept separate in the offline world. They found that increased "friending" of family members also led to increased tension, which the authors hypothesize as resulting from the naturally intimate nature of family relationships. Another worthwhile finding in this article is that increased levels of tension may discourage network growth as users attempt to forestall future tension.

The degree to which family count correlates with tension is interesting, though perhaps not surprising; since the natures of family membership and friendship are inherently different, it makes sense that users would ideally interact with each in different ways and that lumping them together in "Friends" on Facebook is a violation of offline expectations. Finding ways to discretely separate such connections might prove valuable. Additionally, that the level of tension limits network growth suggests that users would be more willing to form digital connections with their offline connections if they could avoid the resulting tension; improving the ability to manage one's spheres may support users in growing their network to its natural size instead of artificially capping it to avoid tension.

Gage Kelly, Patrick, Robin Brewer, Yael Mayer, and Lorrie Faith Cranor. "An Investigation into Facebook Friend Grouping." *Proceedings of the 13th IFIP TC 13 International Conference on Human-computer Interaction - Volume Part III(2011)*: 216-33. *ACM Digital Library*. Web. 10 Nov. 2013.

Facebook currently uses group-based privacy settings, allowing users to create friend lists that can be used to restrict shared content. This article compares multiple mechanisms for sorting and grouping friends, attempting to identify how each mechanism influences the strategy of creating groups and their resulting composition.

The authors conducted a series of semi-structured interviews, beginning with general privacy behavior questions, followed by a friend sorting task. Each participant was asked to use one method (pre-selected from a set of four) to sort their actual friends: a card sorting method, where participants were presented with an index card representing each of their friends; a grid tagging method, where participants applied markers to several pages full of friends; a file hierarchy method, matching the functionality of a digital folder structure; and the existing Facebook friend list interface.

In their initial questioning, the authors found that only 17% of participants used Facebook's friend lists for controlling privacy, suggesting that users are either unaware of the feature or find it too complicated to use. In their study of mechanism/interface effects on grouping, the authors found that participants either sorted friends one by one ("by friend") or created groups and identified friends that should go in that group ("by group"). They found the card sorting method to be over twice as fast as the other new methods, though all participants placed friends exclusively in single groups, despite Facebook's support for group overlap. Finally, the authors categorized that types of groups that people formed, noting common groups like family, co-workers, college friends, location-based friends, and generic friends. Even after the participants created groups, a subsequent scenario-based exercise found that they didn't use the groups in practice, suggesting that people don't think of sharing in terms of groups. Many users reported that they would either share an item with everyone or not at all because they found the consideration of each user to be too laborious.

The findings suggest that creating sharing groups doesn't match the model by which users want to share. One solution could be to take the focus away from the recipients and instead focus more on the identity or context of the user that's doing the sharing. Additionally, the research's identification of common group types gives insight into the spheres that people already recognize on their list of friends. Finally, the incorporation of a card sort-like function could make categorizing friends more easily, though it may be better applied to a concept different from friend lists.

Johnson, Maritza, Serge Egelman, and Steven M. Bellovin. "Facebook and Privacy: It's Complicated." *Proceedings of the Eighth Symposium on Usable Privacy and Security (2012)*: n. pag. *ACM Digital Library*. Web. 31 Oct. 2013.

This article looked at the feelings Facebook users have about sharing content on the network, such as the main concerns they have and the connections that they want their material to be seen by. It also examined the strategies that users take to avoid oversharing.

The authors deployed a Facebook application as a data gathering and survey method. At a basic level, the application identified the levels of access that users had assigned to various parts of their profile. They asked

participants about their level of concern about various public-facing dangers. To estimate the make-up of friend networks, the authors randomly displayed 9 friends to each participant, asking them to express their level of concern about that person seeing all of their content. Additionally, they randomly displayed 8 pieces of shared content from the participant and asked them how concerned they were about all of their friends seeing the content.

The authors found that people were often concerned with strangers having access to their data, but that 95% of users had blocked either some of the posts or some of their photos from public view, and that 45% had no information visible to the public. Therefore, they concluded that such “outsider threats” were already largely addressed.

They also found concern with “insider threats,” or inappropriately sharing content with some set of their friends; 37% of people were concerned about access either by the 9 friends shown or access to the 8 posts shown, though the authors expect that this number would be even higher if the number of friends and content displayed was higher. People were most concerned with over-sharing with online-only friends, co-workers, and friends of friends. Despite these concerns, only 18% of participants had changed the per-post privacy settings of any of their content; when they did, it was primarily in an exclusive fashion. Another strategy that a few participants undertook was creating multiple accounts to manage social spheres, such as professional vs. social contacts.

Though sharing with strangers was a concern, the concern was significantly lower in participants who had actively managed their privacy settings. This is in line with other research that suggests that users who take more control over their privacy often feel more comfortable and engage more frequently. Beyond strangers, the primary concern of the “insider threat” suggests that people wish to be able to manage the impressions they make upon their connections. The authors suggest that the current implementation of friend lists is insufficient, however, indicating that a more *ad hoc* approach may be more appropriate.

Liu, Yabing, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. "Analyzing Facebook Privacy Settings: User Expectations vs. Reality." *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (2011)*: 61-70. *ACM Digital Library*. Web. 31 Oct. 2013.

In this article, the authors examine the extent to which user content was shared with an audience that is different from what users later identified as the intended audience.

The authors used a Facebook application to survey participants about their general privacy feelings and practices. Additionally, the application asked participants to identify the desired audience for a random set of 10 photos pulled from the participant’s own shared photos. Participants were shown each of their photos, then asked to indicate whether the desired audience is “only me,” “some friends,” “all friends,” “friends of friends,” or “everyone” (public). The authors then compared these values to the actual privacy settings of each photo to determine the degree to which photos were mis-shared.

They found that 36% of photos were shared with the default setting (public), despite participants indicating that only 20% of their photos should be exposed to the public. For 63% of randomly selected photos, the photo was shared with an audience that was different than what the participants later indicated as the desired audience; in 77% of these cases, the photo was shared with an audience larger than desired, and in 51% of these cases of incorrectly shared photos, the photo was shared publicly. This means that 32% of all surveyed photos

were shared publicly despite the user later indicating they should have been restricted.

The authors also looked at photos for which the participants had changed the default setting, finding that even when users actively changed the audience, 61% of photos were still shared with an audience different than intended.

Though the article reveals dramatic disparities between actual audiences and intended audiences for photos, it's worth noting that the study can't differentiate between sharing inconsistencies caused by mistakes or caused by post-sharing regret; that is, some cases may be explained by participants later deciding that they shouldn't have shared the photo as broadly as they originally had, something which wouldn't be solved by interface tweaks. Still, the discrepancies are considerable. One solution could be to periodically prompt users to re-assess the audience of content shared in the past so that the contents' audience reflects the desires of the present-day user and not their past selves. Additionally, a solution to the problem of mis-sharing due to leaving the default setting on is to always prompt the user to make a choice (since users are often notorious for not changing settings).

Morris DiMicco, Joan, and David R. Millen. "Identity Management: Multiple Presentations of Self in Facebook." *Proceedings of the 2007 International ACM Conference on Supporting Group Work (2007)*: 383-86. *ACM Digital Library*. Web. 15 Nov. 2013.

In this article, the authors investigate the use of a single site for managing multiple networks and social contexts. Noting that Facebook and similar sites were initially designed to support a single type of network (e.g. Facebook was designed for intracollegiate networking), the authors were particularly interested in learning how users manage multiple contexts as Facebook expands beyond universities.

The authors conducted a survey by examining the Facebook profiles of colleagues in their own company. They recognized three user segments based on the degree to which the users incorporated friends of different social contexts (college vs. corporate). They recognized that the users who primarily had college-based profiles had not tailored their profile to a professional environment, often revealing personal details that would be inappropriate for corporate context, especially their listed interests and their status updates. Profiles that were more based on a mix of connections were more work-friendly; their actions and provided info suggested that they were more actively presenting a professional version of themselves on Facebook.

The authors also conducted interviews to further understand the landscape. They found that users had varying degrees of concern over their profiles: while some individuals left their profiles unchanged from their college days (explaining that Facebook was meant to be for fun outside of work), others expressed that they cleansed their profiles and changed their habits after recognizing that their online presentation matters.

In terms of our project, this research suggests that some users already recognize and understand the importance of actively managing an online persona so that others perceive you in the intended way; our design should empower users with this management more so than they already are, if possible. Additionally, a design that helped unconcerned users recognize the risks of poor self-presentation might encourage users to change their habits and act more sensibly online.

Staddon, Jessica, David Huffaker, Larkin Brown, and Aaron Sedley. "Are Privacy Concerns a Turn-off?: Engagement and Privacy in Social Networks." *Proceedings of the Eighth Symposium on Usable Privacy and Security (2012)*: n. pag. *ACM Digital Library*. Web. 11 Nov. 2013.

This article examines the relationship between Facebook engagement (defined as posting, commenting, and liking) and multiple privacy metrics. Additionally, the article identifies and categorizes the main reasons that people are concerned about online privacy.

The authors deployed a survey asking participants to self-report their frequency of engagement (how often they visit Facebook, post a status update, post a photo, post a comment, or like a post), as well as their general concern, understanding, and feeling of control over their privacy. They also asked participants to identify the main concerns that they face in using the Internet.

The authors' research found that people who reported more concern or apprehension about privacy on Facebook tended to be less engaged overall on the network. Similarly, they found that a higher self-reported understanding and sense of control of privacy settings were positively correlated with engagement. Unsurprisingly, increased confidence about privacy was exhibited in users who used Facebook's privacy features more frequently. In general, Facebook users who felt more comfortable with their privacy confidence and capabilities were more likely to share and participate in discussions. These findings were true even after controlling for other variables, suggesting that privacy concerns and engagement are closely related. Apprehension about sharing was also strongly correlated with a smaller network size, perhaps indicating that concerned users are more likely to keep their networks limited and closely-knit. Finally, the authors found that identify theft and unintended access to personal data were the biggest perceived privacy threats, suggesting that users were concerned about their information being seen by a larger audience than intended.

While confirming mere correlation and not causation, the research may indicate that increasing awareness and capabilities of privacy controls would lead to more engagement in the form of more communication and more content creation. Such results would be in Facebook's best interest in the form of increased traffic and sharing. For our project, it would be beneficial to design privacy controls that are (or even simply feel) more empowering or that are much more clear and informative about the sharing scope than what currently exist.

Tee, Kimberly, A.J. Bernheim Brush, and Kori M. Inkpen. "Exploring Communication and Sharing between Extended Families." *International Journal of Human-Computer Studies* 67.2 (2009): 128-38. *ACM Digital Library*. Web. 26 Oct. 2013.

In this article, the authors examine the ways in which extended families share life updates, photos, and calendar content, as well as their feeling about their current sharing practices. They attempted to understand the landscape of sharing content with extended family members, such as the complications that might arise and the issues people face.

In conducting their research, the authors used semi-structured interviews and diary studies to identify sharing patterns. They asked participants to map out their extended family, indicating the people with whom they wanted to share and the level of relationship for each member. They asked the participants about their current sharing habits and whether they were satisfied with the amount of sharing and the methods available.

The authors found that while people were generally excited about sharing information with their extended family, some wished that the quality of the interactions were increased; people often received uninteresting, such as chain emails or impersonal updates. People also recognized that there were some things they needed to be careful about in sharing with their family. For example, some participants had had issues with accidentally hurting someone's feelings or oversharing some information. They also identified several barriers to family sharing, such as older individuals who are wary about sharing online due to lack of confidence.

It's clear from this article that many people are very interested in sharing content with their family, especially as social networks have increased the capabilities of interaction even more since the article was written. Currently, family members on Facebook exist primarily the same set of friends as everyone else, which can lead to problems of oversharing or content that wasn't posted with family members as the primary intended audience; one potential improvement could be to segregate family members into a separate identity or section so that the content shared with family members is always targeted towards them.

Wang, Yang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. "'I Regretted the Minute I Pressed Share': A Qualitative Study of Regrets on Facebook." *Proceedings of the Seventh Symposium on Usable Privacy and Security (2011)*: n. pag. *ACM Digital Library*. Web. 1 Nov. 2013.

In this article, the authors looked at the concept of regret on Facebook. Specifically, they focused on what type of content people regretted sharing, why they had originally shared the regretful content, what later made them regret the post, and the strategies they use to avoid sharing potentially regretful content.

The authors conducted a series of surveys and a diary study, both aimed at addressing the questions described above. They asked users to identify specific instances where they regretted sharing information, as well as the context around why they regretted making the post.

The participants' responses were coded into several discrete categories that the authors identified: sensitive issues, such as drug use, sexual content, religious or political beliefs, profanity, personal issues, and work-related complaints; posts filled with strong sentiment, such as negative, offensive, or argumentative comments; and posts that result in the poster or their friends being caught in a lie or having secrets exposed.

The authors identified that users shared such posts in efforts to be perceived as funny or cool, because they weren't thinking about the consequences, or because they were in an emotional or inebriated state. Posts became regretful when users realized that they were seen by unintended audience or they had consequences that the poster did not anticipate. Users also mis-shared information due to lack of understanding of social network privacy settings. In order to avoid sharing potentially regretful information, users used several methods. Among these were the creation of a set of mental rules to guide sharing, to delay posting in case they change their mind, to self-censor, or to read content without sharing any of their own.

Overall, regrets seem to occur primarily when information is shared with an audience larger than intended. Some people handle this by deciding to either share with everyone or not at all; this is a big missed opportunity for the sharing of some content that would ideally go to a limited audience. Others had mental models based on context or the impressions they were trying to manage, such as separating their professional worlds vs their personal worlds, or by differentiating based on social rank (excluding some content from parents, relatives, or teachers). This suggests that users mentally break their world into different identities; supporting this concept could yield a beneficial design.