

Privacy and social networking sites (SNSs)

Manya Sleeper

4/6/15

Range of privacy concerns for SNSs

- “**Context collapse**” combines separate offline groups (e.g., friends, family, coworkers)
- “**Imagined audience**” may not align with actual audience
- **Temporal changes** in preferences
- **Privacy tools** may be unclear/hard to use
- Minor to severe **consequences**

SNSs are a wide category

- Many different sites: Facebook, Twitter, Google+, Snapchat, Tumblr, etc.
- Varied functionalities, focuses, norms, etc.
- Different types of users
- Range of privacy threats/options

Varied methods to study SNS privacy

- Surveys
- Interviews
- Log data
- Mixed methods (e.g., survey with log data)
- Diary studies/experience sampling
- Experimental lab/field studies

Today's agenda

- Discuss some methods for studying privacy (that work for SNSs)
- Go over some past work that used some of these methods
- Activity!

Method: surveys and interviews

Ask people about behaviors and/or attitudes through surveys and/or interviews

Pros:

- Don't need to collect behavioral data
- Can control questions/sample

Cons:

- Reporting biases
- Sample biases

Reading for today

- Johnson et al. “Facebook and privacy: It’s complicated”
- Method:
 - Used Facebook app to run a targeted survey
 - Asked participants about people from their network/posts they’d shared
- Findings?

Method: log data

Collect data on actual SNS behavior for statistical conclusions

Pros:

- Eliminates reporting biases (actual behavior)
- May allow larger sample pool/easier longitudinal conclusions

Cons:

- Potential ethical considerations (if not opt-in)
- May not provide insight into the “why” of the patterns

Method: mixed methods

Combine multiple methods (often quantitative data and qualitative data)

Pros:

- Provides insights into quantitative results

Cons:

- Potential ethical considerations (if not opt-in)
- Need to be carefully choose corresponding data

Privacy attitudes and behaviors

- Acquisti and Gross. 2006. “Imagined communities: awareness, information sharing, and privacy on the Facebook”
- Used survey and network data to look at privacy attitudes and behaviors
 - Looked at Facebook network data (mined from CMU)
 - Used questions about attitude toward varied issues to measure concern about privacy (e.g., economics, politics, privacy)
- Findings:
 - Some impact of privacy attitudes on membership: non-Facebook users more likely to be concerned about privacy (except undergrads)
 - Limited relationship between privacy concern and provided profile information
 - Many respondents misunderstood their profile visibility

Evolving privacy attitudes and behaviors

- Stutzman et al. 2012. “Silent Listeners: The Evolution of Privacy and Disclosure on Facebook”
- Method:
 - 6-year longitudinal study of privacy and sharing behaviors on Facebook at CMU
 - Database of snapshots of public profiles from CMU
- Facebook has changed since initial paper (and so has society’s reaction to social networks)

Evolving privacy attitudes and behaviors

- Participants shared less data with the general public in general
- But shared more data, in general, with their friends and at the same time with “third-party apps, (indirectly) advertisers, and Facebook”
- Changes in the Facebook interface/defaults led to an increase in public sharing of some types of personal data

Impression/Audience management

- Context collapse (Marwick and boyd):
 - Social networks (e.g., Twitter) flatten “multiple audiences into one” which “makes it impossible to differ self-presentation strategies”
- Group co-presence (Lampinen):
 - Multiple groups (e.g., friends, work, etc.) are present on a social network
 - Identity and privacy need to be maintained through technical and informal mechanisms

Marwick, Alice E and danah boyd. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*. 2011.

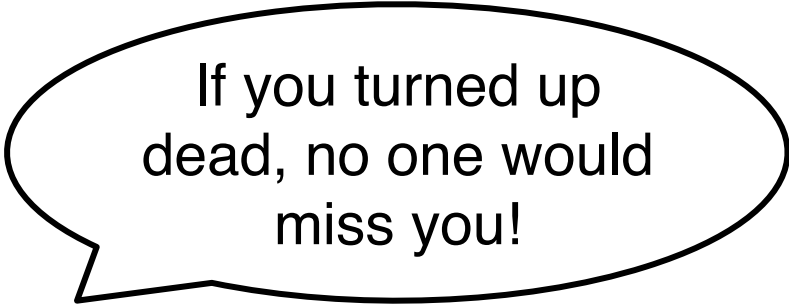
Lampinen, Airi, Sakari Tamminen, and Antti Oulasvirta. "All my people right here, right now: management of group co-presence on a social networking site." *Proceedings of the ACM 2009 international conference on Supporting group work*. ACM, 2009.

“I read my Twitter the next morning and was astonished” A Conversational Perspective on Twitter Regrets

Manya Sleeper, Justin Cranshaw,
Patrick Gage Kelley, Blase Ur,
Alessandro Acquisti, Lorrie Faith Cranor, and
Norman Sadeh




It's easy to **say**
something you
regret.



If you turned up
dead, no one would
miss you!



I hate you!



You look like
you've gained a lot
of weight...

Past
research
analyzed **in-**
person
regret:

- Factors leading to regret
- Types of regret
- Awareness of regret
- Strategies to repair regret

It's also possible
to **tweet**
something you
regret

Thanks for putting me at risk of
getting fired

Maybe, if you would take your
stupid elsewhere...I wouldn't
have to be so blunt

Man, I hate you, you are the
worst person ever, should've
never been born

Twitter allows for:

- Wider audiences
- Lack of face-to-face channel
- Increased persistence



To understand how Twitter regrets
compared to conversational
regrets...

...looked at Twitter users'
regretted messages from
in-person conversations and
on Twitter.

Regretted messages on Twitter and in person

- What **states** lead to regret?
- What **types** of regret occurred?
- How did people become **aware** of regretted messages?
- What **repair strategies** did people use to cope with regretted messages?

Large-scale online survey

- Amazon Mechanical Turk
- 1,221 Twitter users
 - English proficiency
 - Relatively frequent Twitter use
 - Reported a regret



Survey with two conditions


- Conversational and Twitter conditions
- Asked to “recall an occasion when” said or tweeted something and then regretted it
- Described:
 - Regret
 - Circumstances leading to regret
 - How became aware of regret
 - Repair strategies

Data coding and analysis

- Coded open response questions based on in-person conversational regrets literature
- Did not perform statistical comparisons across conditions
 - Different contexts (Twitter/conversation)
 - Qualitative explored themes/trends
 - Performed statistical tests within conditions

States leading to regret

- Negative emotional states common
 - Stress
 - Anger
 - Frustration
- Positive emotions less common



Types of regrets


- Codes from conversational regrets literature (Knapp et al.)
- Most common:
 - Direct criticism
 - Direct attack
 - Implied criticism
 - Expressive
 - Revealed too much
 - Blunder



Types and audience

- Asked participants for **intended** audience
- Twitter participants tended to target multiple people (73% reported)
- Types significantly more likely to be targeted at multiple people:
 - Blunders (82%)
 - Expressive content (84%)
 - Content that revealed too much (80%)

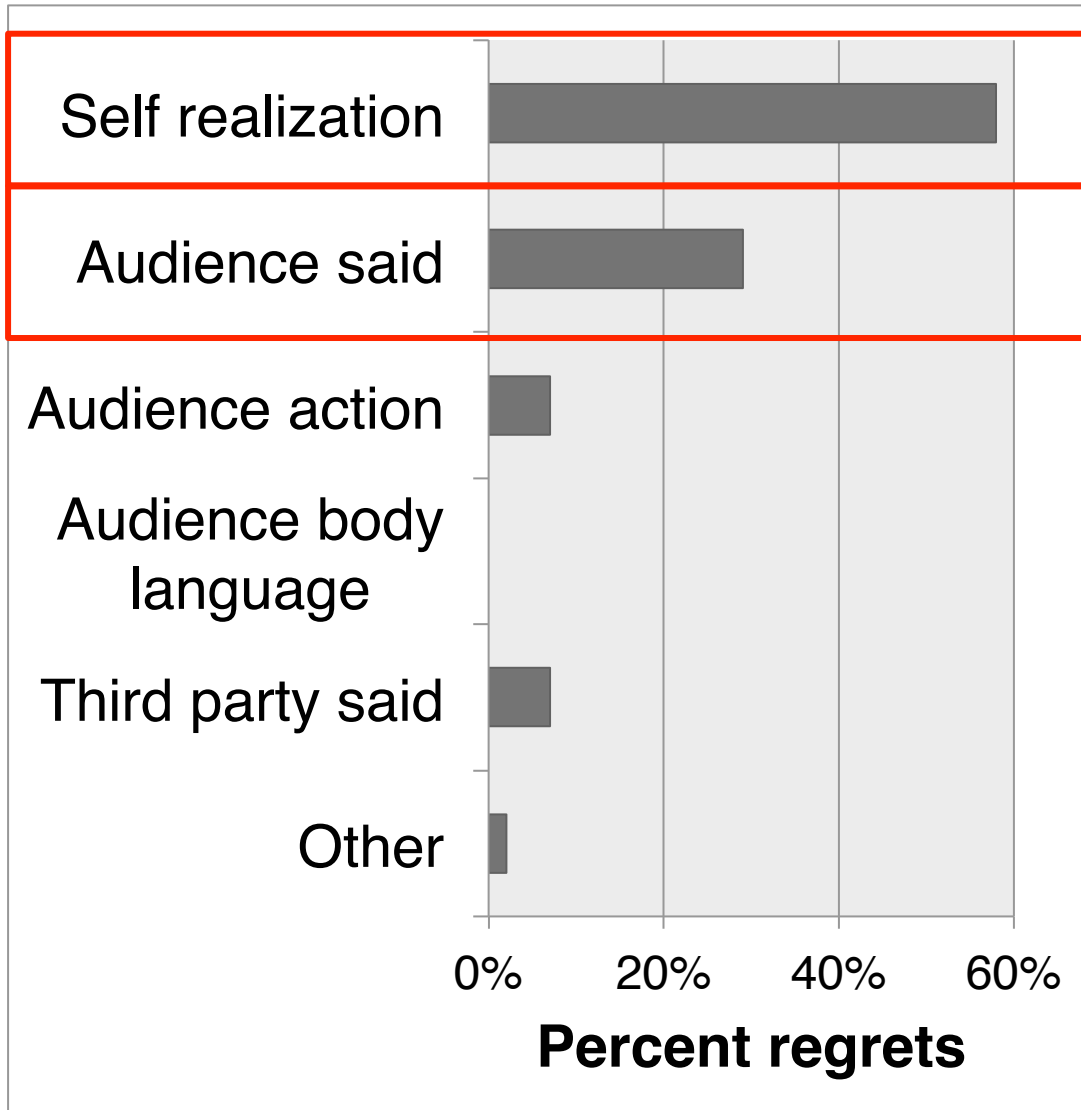
Awareness: Conversation



“When he gave me an angry look and yelled at me”

“When she began to cry I realized how much I hurt her...”

Awareness: Twitter



“As soon as I got a text from the girl I had vented on twitter about. She was none to happy.”


“Re-reading it the

“Once some of my old classmates and friends DM me and told me to stop and that it was very immature of us.”

Repair strategies

Twitter

Conversational



Time to awareness and repair

- Most conversational awareness immediate (63%), repaired within few minutes (52%)
- Twitter participants reported awareness and repair that lagged

Conclusions

- Twitter participants tended to regret critical messages, often targeted at broad audiences
- Twitter participants often became aware of regret hours or days later
- Twitter participants often relied on self-awareness or third-parties to tell them about regrets, absent physical audience cues
- Once aware of regrets, Twitter participants tended to delete the regretted tweet and/or apologize

Method: controlled experiments

Implement a change and test in the lab or field with controls and interventions (between/within subjects)

Pros:


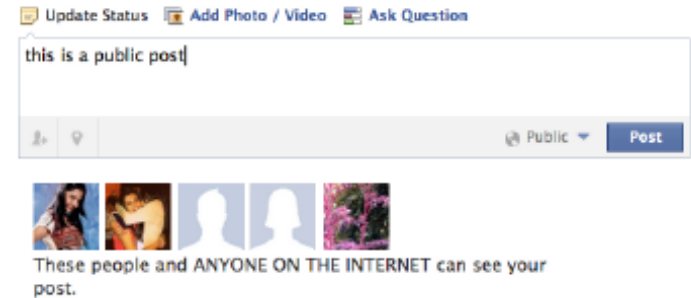
- Able to determine causal links
- Provides behavior under a design change instead of current conditions

Cons:

- Can be difficult to create conditions for statistical significance
- Realistic test conditions can be difficult to implement in the field or lab

Privacy “nudges”

- Designed “nudges” for Facebook to encourage users to consider sharing decisions
- 6-week field trial with Facebook users (n=28)
- Some participants found the “nudges” useful, while others found them annoying



Method: diary studies

Ask respondents to “diary” events from their daily lives as part of a study

Pros:

- Receive data on daily events/hypothetical daily events
- Can be used to ground interviews/survey responses

Cons:

- Act of reporting can change participants’ behaviors/attitudes
- Can be burdensome to participants

The Post that Wasn't: Exploring Self-Censorship on Facebook

Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, Lorrie Faith Cranor



When we look at Facebook we can see what people have posted.

The screenshot shows the Facebook interface from the perspective of a user named 'Manya Sleeper'. The top navigation bar is blue with the 'facebook' logo, icons for home, messages, and notifications, and a search bar labeled 'Search for people, places and things'. Below the navigation bar, the user's profile picture and name 'Manya Sleeper' are visible, along with an 'Edit Profile' link. The left sidebar contains a 'FAVORITES' section with links to 'News Feed', 'Messages' (21), 'Events' (1), and 'Photos'. Below this are sections for 'ADS' (Ads Manager), 'PAGES' (Like Pages, 3), 'APPS' (App Center, Games Feed (20+), Music, Notes, Links, Pokes), and 'GROUPS' (CUPS, Dartmouth '08, Add Group...). A 'MORE' link is at the bottom of the sidebar. The main content area shows the 'News Feed'. At the top, there are buttons for 'Update Status' and 'Add Photos/Video', followed by a status update from 'Manya' asking 'How are you doing, Manya?'. Below this, a post from a user with a red and yellow profile picture asks 'Are we really regressing back to the point where the king can declare people to be outlaws?'. This post includes a link to a news article titled 'Someone Just Leaked Obama's Rules for Assassinating American Citizens - Hit & Run : Reason.com' with a thumbnail image of Barack Obama. The article text states: 'For over a year now journalists, civil liberties advocates, and members of Congress have been asking the Obama administration to release internal memoranda'. The post has 'Like · Comment · Share · 4 minutes ago' and a comment box. Below this, another post from a user with a green and white profile picture says 'I leave my bike unlocked for one night because the lock is frozen and the next morning it's gone; so that's good.' This post has 'Like · Comment · 28 minutes ago near Princeton, NJ'. It shows a thumbs-up reaction from 'Nathan' and a 'bummer' reaction from another user 14 minutes ago. A third user comments 'Big difference from Hanover, huh? Sorry man.' about a minute ago. Both reaction and comment sections have 'Like' buttons and a 'Write a comment...' box.

We don't see
what people
don't post.



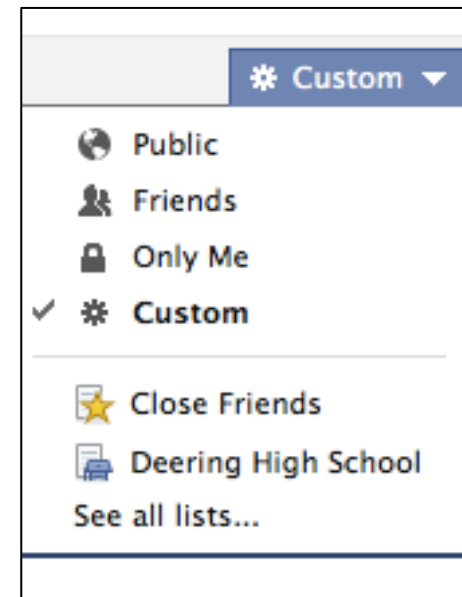
Self-censorship is interesting

- 1. What** are users not sharing?
- 2. Why** do they choose not to share different types of content?

I hate my sister
so much! Oh
wait...no I don't.

Sometimes self-
censorship **can be good**

...but sometimes it
occurs because of
interface limitations



3. What **subset** of currently unshared content could **potentially be shared** if users could exactly target their intended audiences (**optimal selective sharing**)
4. What **attributes** typify the groups with whom users would want to share this content?

Methodology

- 7-day **diary study**
- **SMS** messages
- Whenever thought “of **things** that they would like to post on Facebook but **decide[d] not to post.**”



Methodology

- Details in **nightly surveys**
 - Contained SMS items
 - Asked for additional details, shared content
- Hour-long semi-structured **final interview**

Participants

- 18 participants
- Screened for English, Facebook use, SMS use, and holding back content
- 10 female, 8 male




Data coding

- Coded for:
 - Types of content
 - Reasons for not sharing
 - Types of people would have wanted to share with/block (where relevant)
- Used data from nightly surveys and interview
- Iteratively coded all items

Types of content


Types of content



Drug-related video that a participant decided not to share because her “family in Austin is really religious”

Types of content

Types of content




Decided not to post a “Link to article about young black republicans” to avoid controversy

Wanted to post links to “articles I read on NPR and WeArePowerShift.org – very political stuff” but “I like to keep politics off my Facebook page”

Types of content

Types of content



“My brother-in-law wants to get a tattoo and I was going to comment on how stupid it was; but I decided not to”

Reasons for not sharing


- Presentation of self
- Potentially offensive
- Boring/repetitive
- Avoid argument/discussion
- Inconvenient



Potential for selective sharing

- **Optimal selective sharing:** how much would have shared if could have only targeted particular audiences
- On five-point scale, how likely or unlikely to share if:
 - Could have shared item **only** with people they **wanted to share** it with
 - Could have **prevented** people they **didn't want** to see item from viewing it

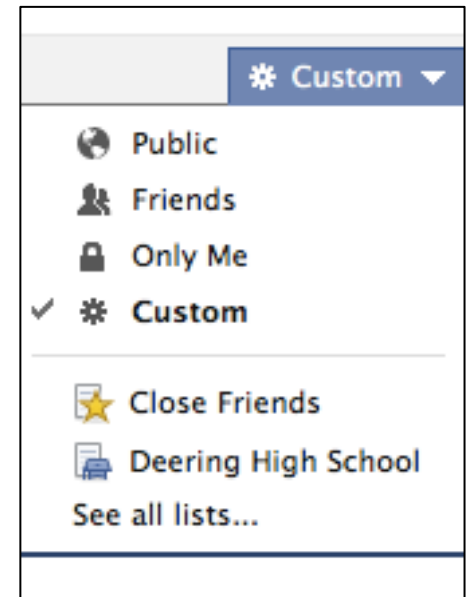
Potential for selective sharing



Approximately half of unshared content would potentially be shared under optimal selective sharing



Types of groups for selective sharing

- To allow for selective sharing would need interface grouping mechanisms
- Asked participants to specify who wanted to share with or block



Types of groups

- **Specific people:** e.g., “my sister”
- **Specific groups:** countable set of people (e.g., 10 close friends)
- **Ambiguous groups:** defined by one or more attributes or relationships (e.g., “hockey friends”)



Groups needed for optimal selective sharing

- Some groups could be captured by Facebook interface
- Need to capture more ambiguous groups
 - Context-specific information
 - Traits potentially unknown to the user

Conclusions

- External content most commonly self-censored followed by personal content
- Presentation of self most common reason
- Half potentially shared under optimal selective sharing
- Some specific individuals and groups potentially could be captured by current interface
- Also wanted to target more ambiguous groups that would require new tools

Activity

How does technical expertise impact choice of privacy strategies for SNSs?

Challenge: Describe at least two studies you could run to study this question. Outline at least two pros and cons of each.