

CSC 495.002 – Lecture 4

Web/Social Networks Privacy: Violations and Regret

Dr. Özgür Kafalı

North Carolina State University
Department of Computer Science

Fall 2017

PREVIOUSLY ON SOCIAL NETWORKS

Sharing and Disclosure

- Common usage scenarios of OSNs
- Common sharing and disclosure patterns of users
 - What content types are shared?
 - Whom are they shared with?
 - How do sharing behaviors change over time?
- Does shared content match intended audience?
- How do users mitigate privacy concerns?

Problem Definition

- Violation: Reality does not meet user expectation about privacy
 - Mismatch between intended and actual audience
 - Unawareness of social links
- Regret: Later become unhappy about negative consequences of sharing behavior
 - Enumerate reasons to share
 - Identify regrettable actions
 - Help users avoid such actions

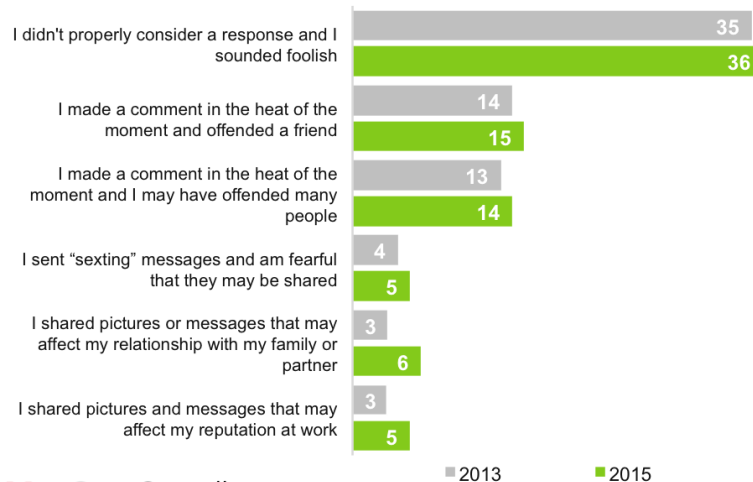
Exercise: Regrettable Actions

- I shared pictures or messages that may affect my relationship with my family or partner
- I shared pictures or messages that may affect my reputation at work
- I made a comment in the heat of the moment and I may have offended many people
- I sent “sexting” messages and am fearful that they may be shared
- I made a comment in the heat of the moment and offended a friend
- I didn’t properly consider a response and I sounded foolish

Regrettable Actions

Which, if any, of the following is your single biggest social media regret? (%)

Base: US adults with social media regrets.



YouGovOmnibus

Julv 13-14 . 2015

http://www.huffingtonpost.com/shane-paul-neil/more-than-half-of-america_b_7872514.html

Exercise: Reasons for Regret

- When I am with friends
- On a laptop/computer
- On my smartphone
- When I have drunk some alcohol
- At home
- Late night when I am tired
- When I am on my own
- When I am busy and respond too quickly
- In the office

Reasons for Regret

Under which, if any, of the following circumstances do you normally make the mistake of sharing information you regret? (%)

Base: All US adults.



http://www.huffingtonpost.com/shane-paul-neil/more-than-half-of-america_b_7872514.html

Violation Types

- Norm violations
 - Norms describe normal (expected) behavior
 - Some norm violations are desirable (to maintain functionality)
- Violations of privacy laws
 - Some norms can be implemented as laws
 - Sanctions applied in case of violations
- Exceptions
 - Depends on user expectations
 - Not all violations are exceptions
 - There might be exceptions even if no violations

Studies

- Look at two studies
 - One formal reasoning method for predicting privacy violations
 - One empirical study about regrets

Detecting and Predicting Privacy Violations in Online Social Networks

Detecting and Predicting Privacy Violations in Online Social Networks

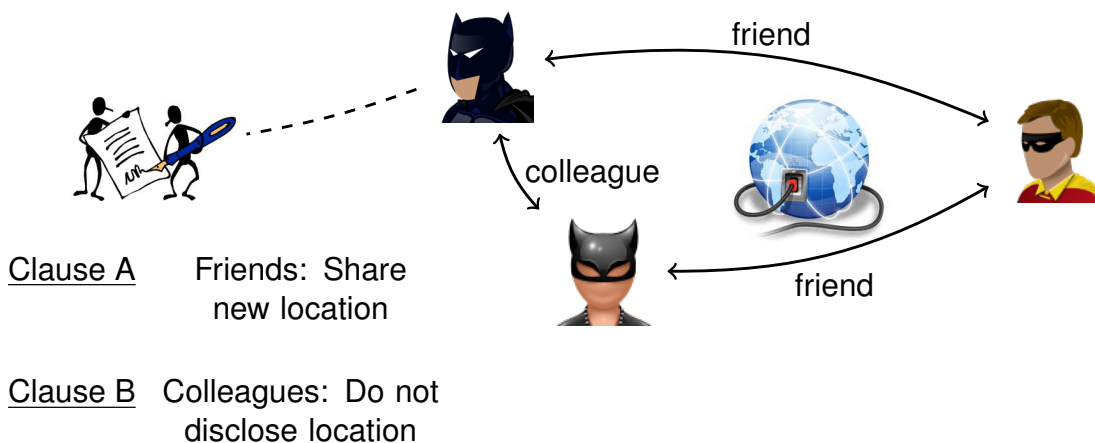
Özgür Kafalı · Akın Günay · Pinar Yolum

Received: date / Accepted: date

Abstract Online social networks have become an essential part of social and work life. They enable users to share, discuss, and create content together with various others. Obviously, not all content is meant to be seen by all. It is extremely important to ensure that content is only shown to those that are approved by the content's owner so that the owner's privacy is preserved. Generally, online social networks are promising to preserve privacy through privacy agreements, but still everyday new privacy leakages are taking place. Ideally, online social networks should be able to manage and maintain their agreements through well-founded methods. However, the dynamic nature of the online social networks is making it difficult to keep private information contained.

We have developed *PROTOSS*, a run time tool for detecting and predicting *PRivacy viOLaTions in Online Social networks*. *PROTOSS* captures relations among users, their privacy agreements with an online social network operator, as well as domain-based semantic information and rules. It uses model checking to detect if relations among the users will result in the violation of privacy agreements. It can further use the semantic information to infer possible violations that have not been specified by the user explicitly. In addition to detection, *PROTOSS* can predict possible future violations by feeding in a hypothetical future world state. Through a running example, we show that *PROTOSS* can detect and predict subtle leakages, similar to the ones reported in real life examples. We study the performance of our system on the scenario as well as on an existing Facebook dataset.

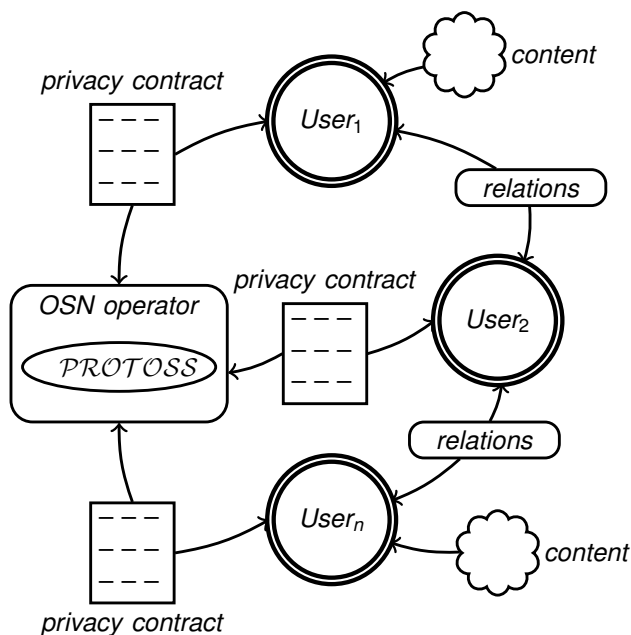
Privacy Contracts



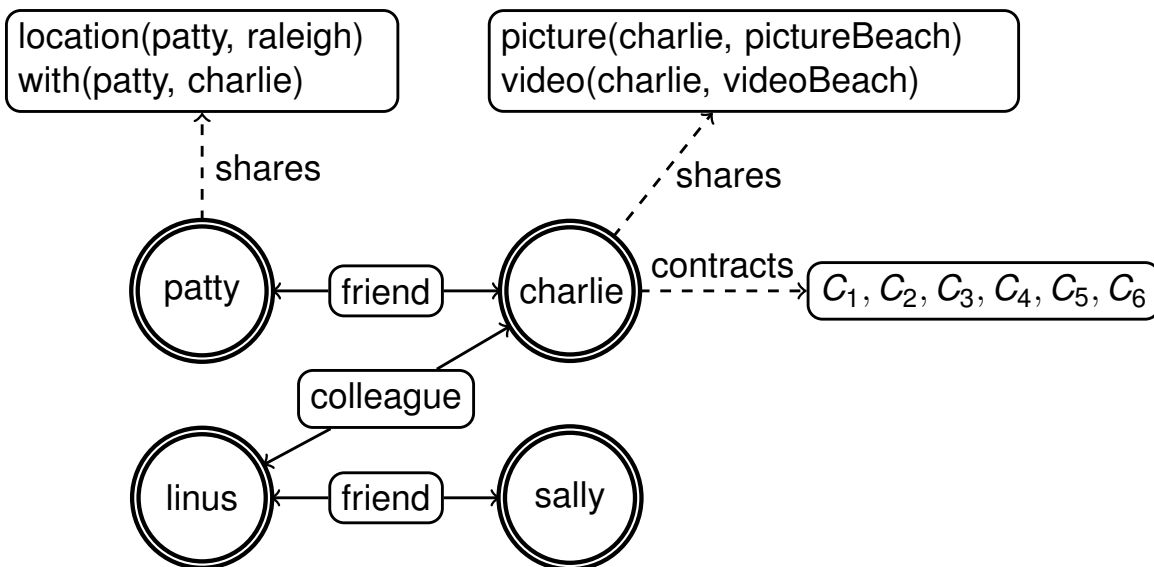
- What happens when Batman checks in at Arkham Asylum, who should know about that?
- What happens when Robin posts a picture together with Batman?

PROTOSS: Contract-based OSN Architecture

- **PR**ivacy **vi**Ola**T**ions in **On**line **S**ocial networks

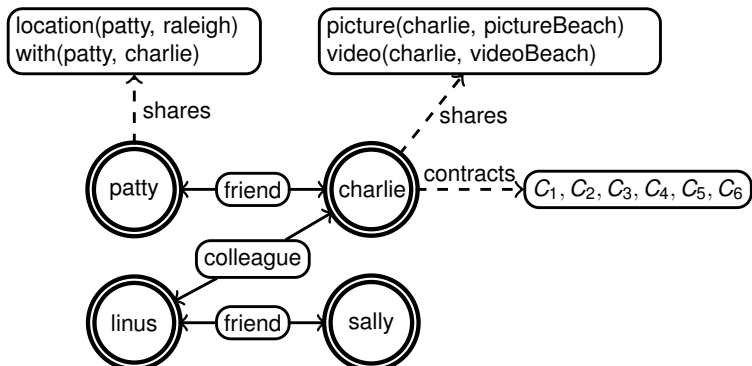


Sharing Example



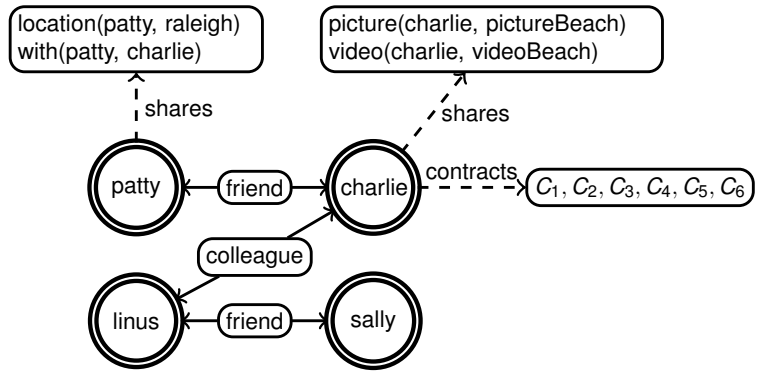
Users

- Charlie
- Patty
- Sally
- Linus



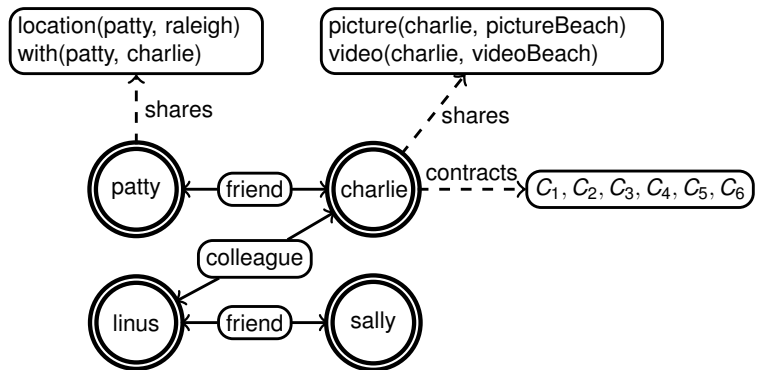
Relations

- $friend(X, Y)$: Users X and Y are friends
- $colleague(X, Y)$: Users X and Y are colleagues
- $friend(patty, charlie)$
- $friend(linus, sally)$
- $colleague(charlie, linus)$



Content

- $location(X, L)$: User X is at location L .
- $with(X, Y)$: User X is with user Y .
- $picture(X, P)$: User X posts a picture P .
- $video(X, V)$: User X posts a video V .



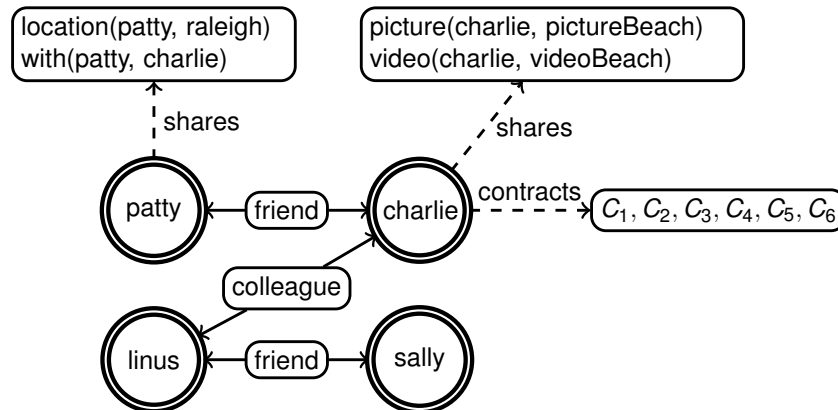
OSN Behavior

- $B_1: \text{visible}(\text{with}(X, Y), Z) \leftarrow \text{friend}(X, Z) \vee \text{friend}(Y, Z)$
- $B_2: \text{visible}(\text{location}(X, L), Y) \leftarrow \text{friend}(X, Y)$
- $B_3: \text{visible}(\text{picture}(X, I), Y) \leftarrow \text{friend}(X, Y)$
- $B_4: \text{visible}(\text{video}(X, V), Y) \leftarrow \text{friend}(X, Y)$

OSN Contracts

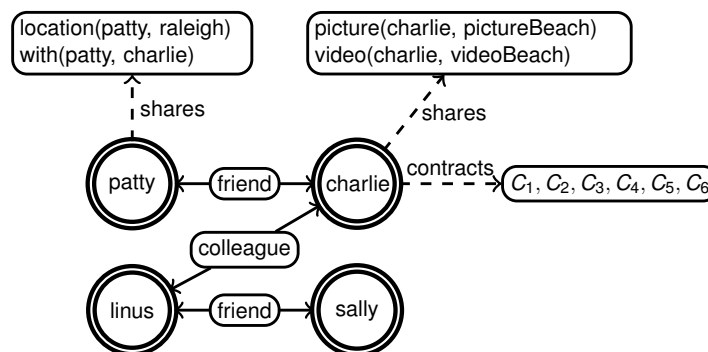
- $C_1(\text{osn}, \text{charlie}, \text{friend}(\text{charlie}, Y), \text{show}(\text{pic}(\text{charlie}, P), Y))$
- $C_2(\text{osn}, \text{charlie}, \text{friend}(\text{charlie}, Y), \text{show}(\text{with}(\text{charlie}, Z), Y))$
- $C_3(\text{osn}, \text{charlie}, \text{friend}(\text{charlie}, Y), \text{show}(\text{loc}(\text{charlie}, L), Y))$
- $C_4(\text{osn}, \text{charlie}, \text{colleague}(\text{charlie}, Y), \neg \text{show}(\text{pic}(\text{charlie}, P), Y))$
- $C_5(\text{osn}, \text{charlie}, \text{colleague}(\text{charlie}, Y), \neg \text{show}(\text{with}(\text{charlie}, Z), Y))$
- $C_6(\text{osn}, \text{charlie}, \text{colleague}(\text{charlie}, Y), \neg \text{show}(\text{loc}(\text{charlie}, L), Y))$

Scenario 1



- According to contract C_4 , pictures of *charlie* should not be revealed to his colleagues
- *linus* should not be able to see *charlie*'s picture *pictureBeach*

Scenario 2



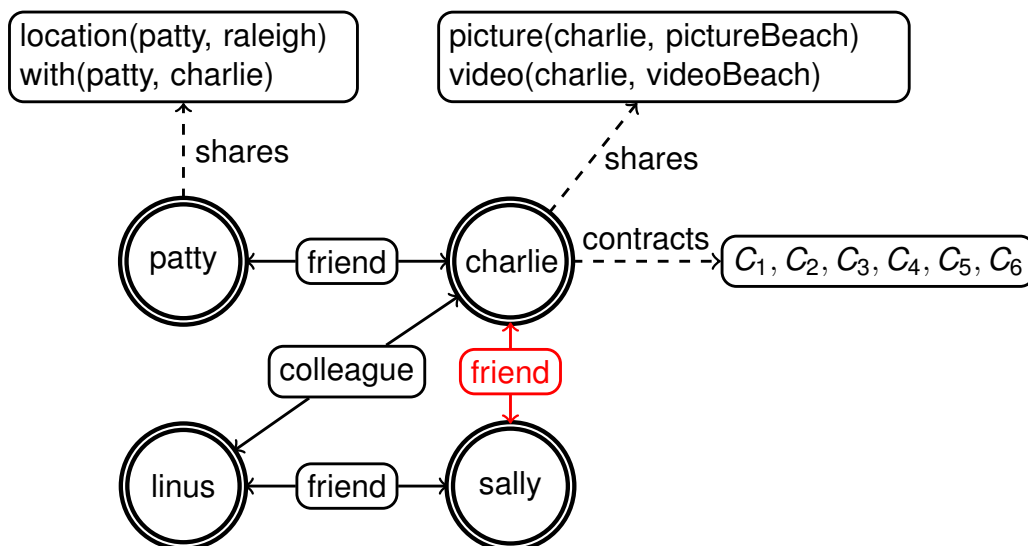
- According to contract C_6 , location of *charlie* should not be revealed to his colleagues
- *linus* should not be able to see *charlie*'s location
- *charlie* does not share his location, but *patty* does (indirectly),
- She shares that she is with *charlie* and she is in Raleigh
- Inference: *charlie* is in Raleigh too

Scenario 3

- As *charlie* stated in his privacy agreement, he does not want his colleagues to view his pictures (contract C_4)
- However, he has not made any statement about his videos (knowingly or unknowingly)
- Is it possible to make further reasoning to infer that videos are by nature similar to pictures?
- If any videos of *charlie* are being seen by colleagues, is it worthwhile to notify him?

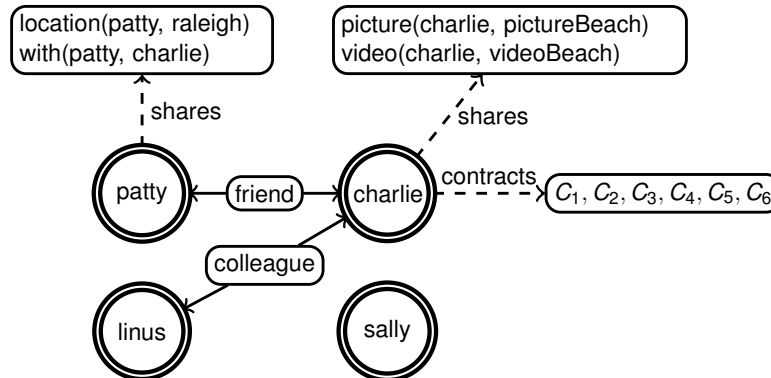
Scenario 4

- Assume that *charlie* meets *sally* in Raleigh and adds her as a friend. Hence, OSN evolves into a new state. The aim is to detect whether *charlie's* picture is visible to *linus*?



Prediction Scenarios

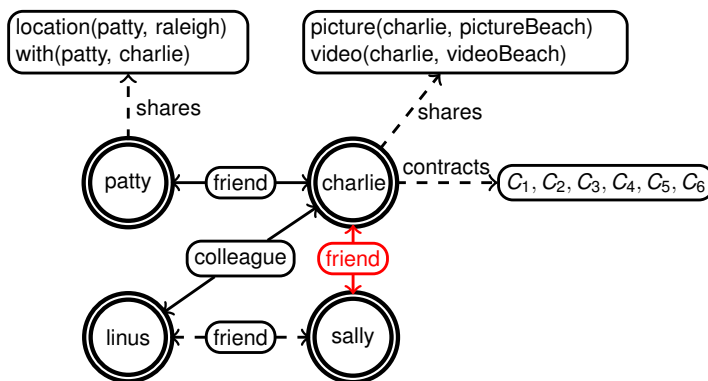
- Go back to the initial state of the OSN, i.e., *charlie* and *sally* are not friends yet
- Look at the OSN from *charlie's* point of view
- *charlie* tries to predict possible future breaches of his privacy depending on the evolution of relations between the users



Scenario 5

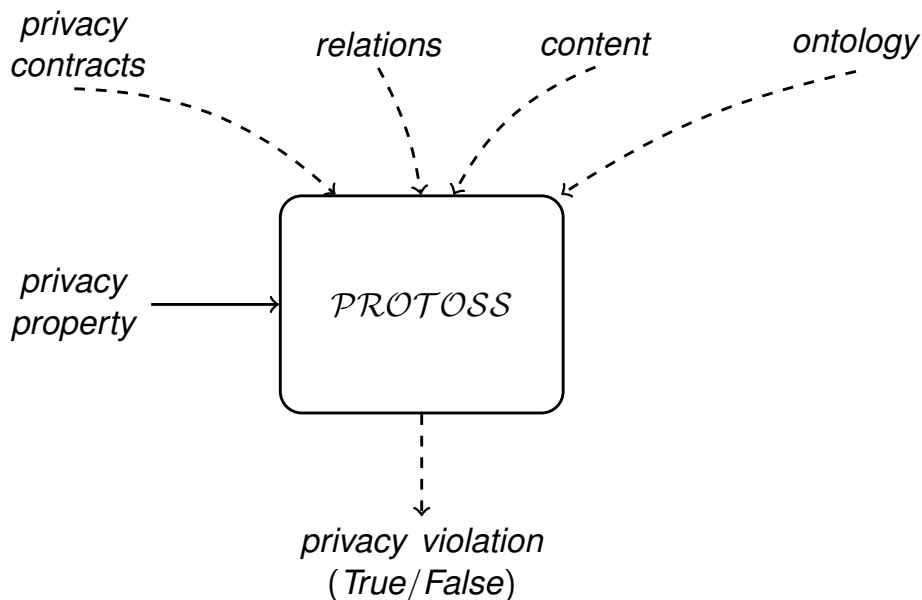
- *charlie* is a cautious user and desires to find out what would it take for *linus* to see his pictures
- That is, what relations in the OSN need to be initiated between the users of the OSN in the future for this information to leak?
- *charlie* chooses not to make any assumptions about the relations of the other users

Scenario 6

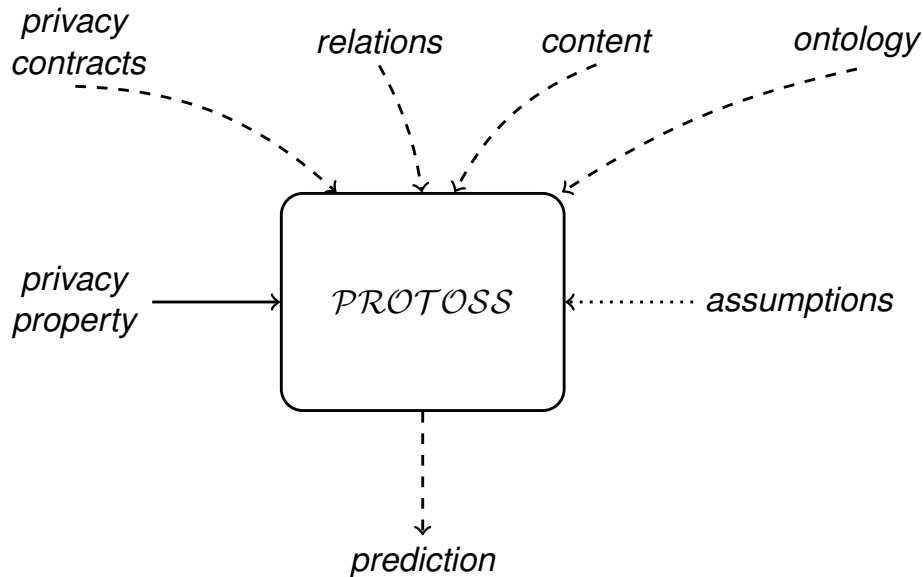


- *charlie* wants to add *sally* as a friend
- He is concerned that this may cause *linus* to see his pictures
- Before adding *sally* as a friend, he wants to find out whether his pictures would be visible to *linus* if he adds *sally* as a friend
- *charlie* assumes *sally* and *linus* are friends
- *charlie* assumes *patty* and *linus* are not friends

Detecting Violations



Predicting Violations



Facebook Dataset

Online Social Networks Research @ The Max Planck Institute for Software Systems

WOSN 2009 Data Sets
 Data from our WOSN 2009 paper is available from the links below. Each of the data sets has been anonymized to protect the privacy of the users themselves. Included is information about the evolving link structure from the networks as well as the communication between users via the wall feature.

Note that we are unable to release any non-anonymized data.
 We are aware that properly anonymizing online social network data is very challenging. Clever schemes have been found to break seemingly well anonymized data sets (e.g. the Netflix data set). For the data we make available, we use a "best effort" anonymization. We do not offer any strong guarantees and we suspect that our anonymization scheme can likely be broken by clever comparisons to other real-world data. We encourage people to help bring problems and fixes to our notice, should they find any.

- **List of links**
 These files contain a list of all of the user-to-user links from the Facebook New Orleans networks. All links are treated as directed, even though they are undirected on Facebook.
 Format: Gzipped ASCII. Each line contains two anonymized user identifiers, meaning the second user appeared in the first user's friend list. Finally, the third column is a UNIX timestamp with the time of link establishment (if it could be determined, otherwise it is "N").
 Data: [Facebook Links](#) (10.4MB)
- **List of wall posts**
 These files contain a list of all of the wall posts from the Facebook New Orleans networks.
 Format: Gzipped ASCII. Each line contains two anonymized user identifiers, meaning the second user posted on the first user's wall. The third column is a UNIX timestamp with the time of the wall post.
 Data: [Facebook Wall Posts](#) (6.8MB)

- Alan Mislove's OSN dataset:
<http://socialnetworks.mpi-sws.mpg.de/data-wosn2009.html>

Dataset Details

User	User	Timestamp
1	18	N
1	20	1217964960
1	23	N
1	24	1227241074

- Each row lists two individuals that are related to each other
- Optionally a date that implies when the relationship between the two individuals were formed
- Does not contain different type of relations or contents
- Assumptions
 - Relations between individuals are friend relations
 - OSN will show the content posted by users (e.g., pictures) to their friends (not anyone else)
 - Users can repost contents initially posted by friends

Methodology

- Research question: Is it possible for a user Y to actually view a content posted by X , even though X and Y are not friends?
- Research question: If so, can we predict it before it happens?
- Take a subset of the dataset such that we begin with one user and add all of her friends and her friends' friends
- Previous work on link prediction has shown that it is very likely for a new friend to be already contained in the friends of friends network

Violation Scenario

- Violation condition: Y ends up viewing a content of X
- There exists a Z that is both friends with X and Y
- Z shares the content of X with Y
- OSN's commitment to X is violated

Performance Results

User	#Users	#Friends	#States	Prediction time
1	27	2129	894.4 K	1.75 s
163	26	1222	396.2 K	0.94 s
1645	29	679	127.1 K	0.89 s
31720	50	2294	557.6 K	1.87 s
48696	16	495	144.1 K	0.50 s

Limitations

- Scalability: Model checking is a computationally expensive approach
- Relaxation of some of the assumptions: Beyond friend of friend

I regretted the minute I pressed share: A Qualitative Study of Regrets on Facebook

“I regretted the minute I pressed share”: A Qualitative Study of Regrets on Facebook

Yang Wang wang@cs.cmu.edu	Saranga Komanduri sarangak@cmu.edu	Pedro Giovanni Leon pedrogl@cmu.edu
Gregory Norcie ganorcie@andrew.cmu.edu	Alessandro Acquisti acquisti@andrew.cmu.edu	Lorrie Faith Cranor lorrie@cmu.edu

Carnegie Mellon University
Pittsburgh, PA

ABSTRACT

We investigate regrets associated with users' posts on a popular social networking site. Our findings are based on a series of interviews, user diaries, and online surveys involving 569 American Facebook users. Their regrets revolved around sensitive topics, content with strong sentiment, lies, and secrets. Our research reveals several possible causes of why users make posts that they later regret: (1) they want to be perceived in favorable ways, (2) they do not think about their reason for posting or the consequences of their posts, (3) they misjudge the culture and norms within their social circles, (4) they are in a “hot” state of high emotion when posting, or under the influence of drugs or alcohol, (5) their postings are seen by an unintended audience, (6) they do not foresee how their posts could be perceived by people within their intended audience, and (7) they misunderstand or misuse the Facebook platform. Some reported incidents had serious repercussions, such as breaking up relationships or job losses. We discuss methodological considerations in studying negative experiences associated with social networking posts, as well as ways of helping users of social networking sites avoid such regrets.

of wine and a mug of beer [14]. These incidents demonstrate the negative impact that a single act can have on an SNS user.

In order to protect users' welfare and create a healthy and sustainable online social environment, it is imperative to understand these regrettable actions and, more importantly, to help users avoid them. In the large body of SNS literature, little empirical research has focused on the negative aspects of SNS usage. We try to address that gap by examining accounts of regrettable incidents on Facebook collected through surveys, interviews, and user diaries.

With more than 600 million users, Facebook has become the world's largest social networking site (according to Alexa, as of August 3, 2010, Facebook has the highest traffic among all SNS sites in the US [6]). While well-evolved norms guide socialization and self-disclosure in the offline world, in the online world it can be more difficult to identify one's audience, control the scope of one's actions, and predict others' reactions to them. As a consequence, Facebook users might not always anticipate the negative consequences of their online activities, and end up engaging in actions that they later regret.

Since they are common experiences that people can recognize and describe, we use regrets as an analytic lens to investigate users'

Goal

- Objective: Understand regrettable actions and help users avoid them
- Research questions:
 - What posts do users regret sharing on Facebook?
 - Why do users make regrettable posts?
 - What are the consequences of these regrettable posts?
 - How do users handle/avoid regrets?

Exercise: Reasons to Post

- Perceived in favorable ways: It's cool, it's funny, ...
- Don't think about consequences: addiction, involuntary action, reflex
- Misjudge culture and norms
- "Hot" state of high emotion
- Misunderstand or misuse Facebook platform

Methodology

- 569 American Facebook users
- Interviews (recruited via Craigslist)
- User diaries
- Online surveys (recruited via Amazon Mechanical Turk)

Survey 1

- Online survey to understand privacy related experiences of users based on previous analysis of users' concerns about Facebook
- People feel certain types of information are more private than others
- People have frustration with Facebook's privacy controls

Interviews

Protocol for Interviews

Interview Guide

<Date>

<Subject #, gender, age range, category of social media use>

<Recording file>

<Setting>

<How long it takes>

Introduction:

Our research group (CUPS) is studying experiences with and impressions of Social Media in general, and Facebook, in particular. We appreciate you taking the time to talk with us today. Everything we talk about will be anonymous and you don't have to answer anything that you don't want to. Having said that, we really want to hear anything that you think will help us understand your experience of Social Media.

Consent - Get their consent (written or oral).

Recorder (if face-to-face or telephone):

Do you mind if we use a recorder? It will be just to make note-taking easier and you can have us turn it off at any point.

1. Brief Intro

Tell us a little bit about yourself:

- Name
- Your social life - strong/weak ties with family, colleagues, friends, etc. How do you stay connected with them?

2. FB attitudes and usage

Tell us about your FB usage. Where, how frequent, what, and why.

- Tell us what you know about Facebook (the tool/site)? How did you hear about it?

If they use it:

Can you log into FB and show us your homepage? (ensure we turn away or use their laptop, ask for the participant's permission if we can video tape the computer screen without recording the participant's face in the video)

- When/how did you start using FB? What did you expect from using FB?
- How frequent do you use FB, when during the day, for how long? Where do you usually log on FB (home/school/workplace/anywhere/mobile)?
- What do you usually do on FB?
- Tell me about the last time you used it? Is this typical?
- How else have you used it? Is that typical?

- Open ended questions
- 19 participants: 8 visited Facebook multiple times a day, 8 once per day, 3 once per week

Diary Study

Diary Study Survey

<Date>

<Subject # >

<Setting>

<How long it takes>

Introduction:

Thank you for participating in our study. For this part of the study, we ask you to answer the following questions on a daily basis. You will be provided with a URL/link to a web form. Please keep the URL private. Your entry to the web form will only be accessible by the researchers.

Your entry will be kept CONFIDENTIAL - only the researchers have access to it.

Questions:

- What activities have you done on Facebook (FB) today?
 - Friend requests. Did you add any new friends? How did you find them on Facebook (e.g., you searched them on FB, or FB recommended them to you, or they sent you requests)? Did you ignore any requests? Why? Who are the people whose friend requests you ignore?
 - Have you posted anything on FB? E.g., status updates or photos. What did you post?
 - Have you changed anything in your profile? What and why?
 - Have you changed anything in your privacy settings? What and why?
- Important incidents (you can write down your experience today or in the past.)
 - Have you posted something on FB and then regretted doing it? Why and what happened?
 - Have you seen any posts from your friends on Facebook that you think they should not have posted? Why and what happened?
 - Did you feel that Facebook invaded your privacy? Why and what happened?
 - Did you feel uncomfortable or embarrassed due to someone learning something about you or your activities on Facebook? Why and what happened?

Exercise: Survey 2

Online Survey II

=====
Page
=====

Please answer this survey only if you have posted something on Facebook and later regretted posting it.

=====
Page
=====

1. Have you ever regretted posting something (status updates, pictures, likes, comments, locations, etc) on Facebook? For example, have you ever posted something that you felt bad about later or wished you hadn't posted?

Yes
 No

2. In the last 12 months, how many times have you regretted posting something on Facebook?

0
 1
 2-5
 6-10
 more than 10

=====
Page
=====

Now, please think about the things that you regretted posting on Facebook (status updates, pictures, likes, comments, locations, etc). Choose the one that you regret the most. For the rest of the survey, think about that post when answering the questions.

=====
Page
=====

3. Describe the post that you regret the most. If it was a status update or comment, what did it say? If you remember the exact words, put them in quotes. If it was a photo or video or something else, describe it.

[free-response text field]

4. Why did you post it?

[free-response text field]

5. Why did you regret posting it?

[free-response text field]

6. What happened after you posted it? For instance, did someone contact you about your post?

[free-response text field]

7. Which of the following reasons best explain why you posted it? (choose all that apply)

I thought it was useful or interesting
 I wanted to congratulate or wish someone happiness
 I wanted to give emotional support
 I wanted to share good news
 I wanted to share bad news
 I thought it was fun or humorous
 I thought it would make me look good
 I wanted to share my feeling or opinion
 I wanted to tell my friends what I was doing

- Based on results (23% reported having regrets), focus on specific aspects of regrets
- Posts: Status updates, pictures, likes, comments, locations

Exercise: What Do People Regret Posting?

- Sensitive content
 - Alcohol and drug use
 - Sexual content
 - Religion and politics
 - Profanity and obscenity
 - Personal and family issues
 - Work and company
- Strong comments, arguments
- Lies and secrets

How to Mitigate Regrets

- Customized privacy settings
 - Friends only: Does it really help when you have 500+ friends?
 - Customs lists
 - Multiple accounts for auditing: Not aware of available options
- Delays
- Self censoring and cleaning
- Appropriate communication channel: Post vs private message

Lessons for Design

- Ultimately, develop models to predict the occurrence and severity of regrettable posts before they are published
- Develop tools capable of identifying posts likely to regret, and integrate into OSNs
- Detect content with strong sentiment and warn the user
- Remind user about friends whom are less interacted with

Limitations

- Privacy paradox: Dichotomy between users' stated privacy preferences and actual behavior
- Participants not necessarily representative of the whole Facebook population
- Does not differentiate between where the regrettable content is
 - Disclosed data: What user posts on own page
 - Entrusted data: What user posts on others' pages
 - Incidental data: What other people post about the user

Connection to Protoss

- How would you use a tool like Protoss to predict some of those potential regrets?
- Predict unintended audiences: People beyond Facebook can be involved

In a Mood? Call Center Agents Can Tell

- News article: <http://www.nytimes.com/2013/10/13/business/in-a-mood-call-center-agents-can-tell.html>
- Links are also on the course website

Things to Look For

- Root cause: What went wrong?
- If it was not intentional, what was the original aim?
- Affected parties
- Implications and similar problems
- Mitigation (using methods we have seen): Prevention, detection, recovery

- Take 10 minutes to look at the incident on your own

- Now discuss with your neighbor
- Also take a look at the summary report: <https://drive.google.com/file/d/0B3m-I0YVAv0EcXIINGN6akI2M2M/view>