Temporal privacy/deletion privacy

Mainack Mondal

CS 60081
Autumn 2021

Roadmap

- Passwords/multi factor authentications
- Usability for security developers
- Online tracking
- **Temporal aspect of privacy**
- Privacy notices/dark patterns

Temporal Privacy: Changing privacy settings

Need to revisit old privacy settings

2009

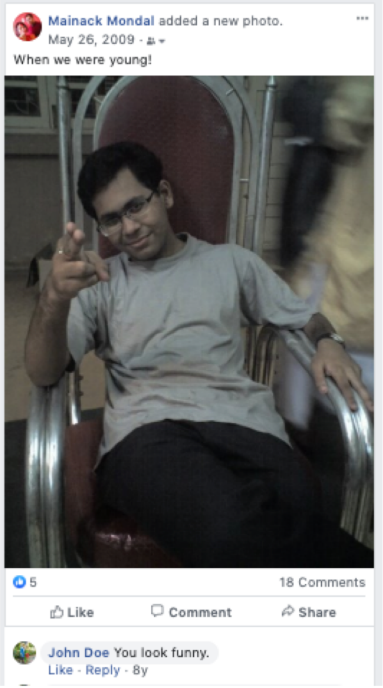

Mainack Mondal added a new photo.
May 26, 2009 ·

When we were young!

5 · 18 Comments

John Doe You look funny.
Like · Reply · 8y

Privacy setting: “**all friends**”

Need to revisit old privacy settings

2009

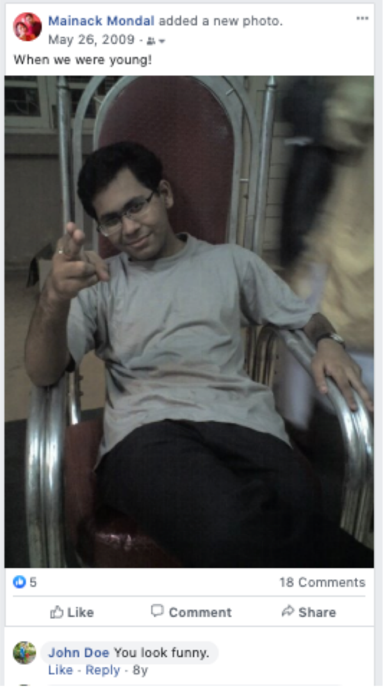

Mainack Mondal added a new photo.

May 26, 2009 ·

When we were young!

5

18 Comments

John Doe You look funny.

Like · Reply · 8y

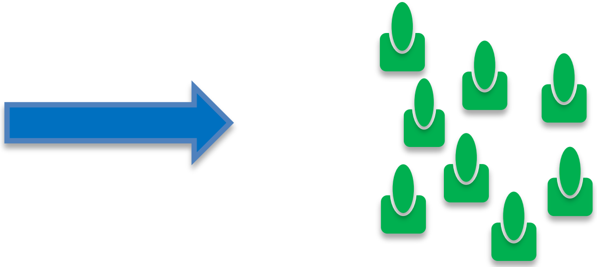

Undergraduate friends

Privacy setting: “**all friends**”

Need to revisit old privacy settings

2019

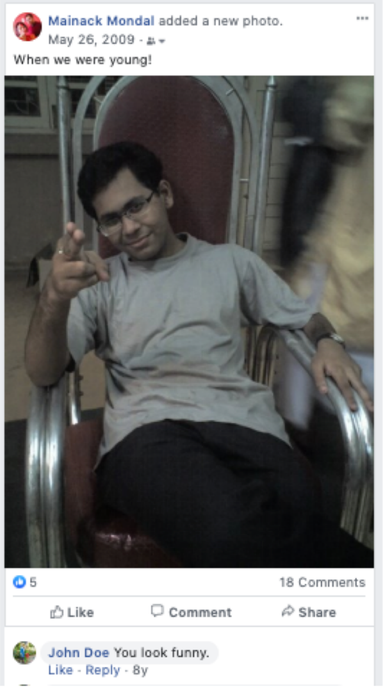

Mainack Mondal added a new photo.
May 26, 2009 ·

When we were young!

5 · 18 Comments

John Doe You look funny.
Like · Reply · 8y

Privacy setting: “**all friends**” !!

Need to revisit old privacy settings

2019

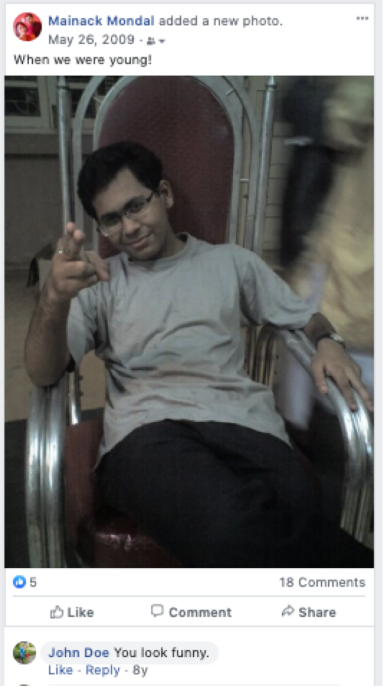

Mainack Mondal added a new photo.
May 26, 2009 ·

When we were young!

5

18 Comments

John Doe You look funny.

Like · Reply · 8y

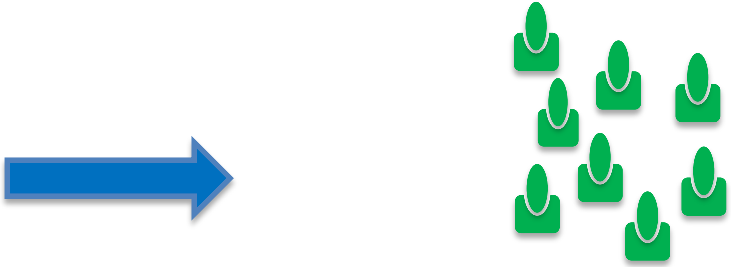

Undergraduate friends

Privacy setting: “**all friends**” !!

Need to revisit old privacy settings

2019

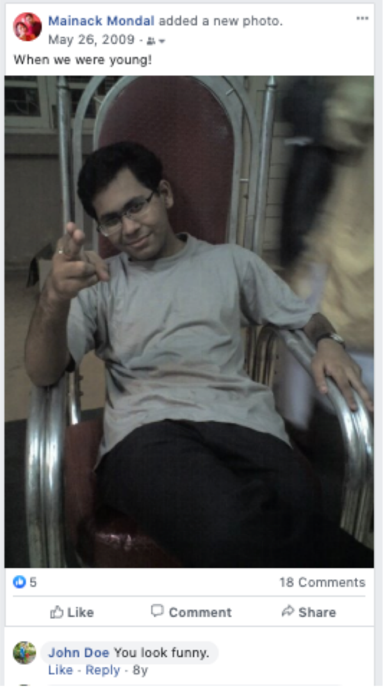

Mainack Mondal added a new photo.
May 26, 2009 ·

When we were young!

5 18 Comments

John Doe You look funny.
Like · Reply · 8y

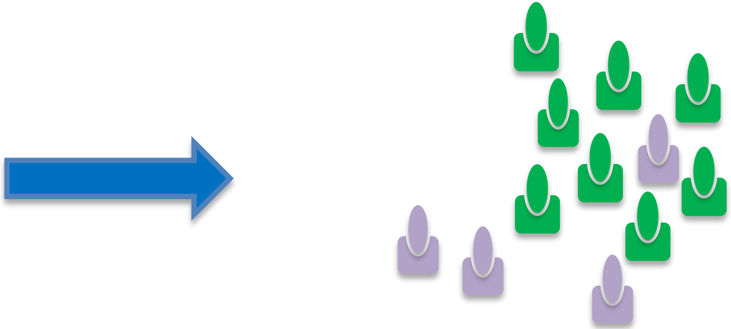

Undergraduate friends

Graduate school friends

Privacy setting: “**all friends**” !!

Need to revisit old privacy settings

2019

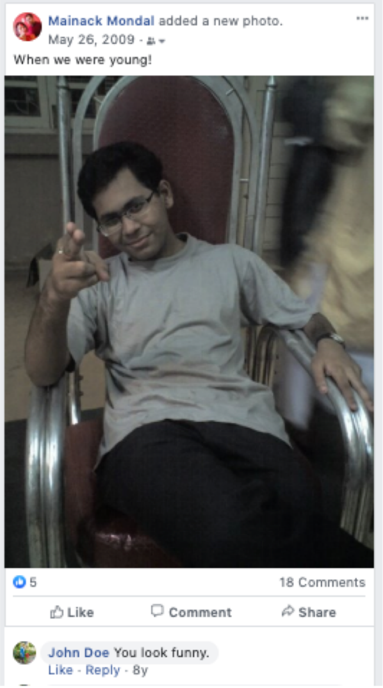

Mainack Mondal added a new photo.
May 26, 2009 ·

When we were young!

5 · 18 Comments

Like · Comment · Share

John Doe You look funny.
Like · Reply · 8y

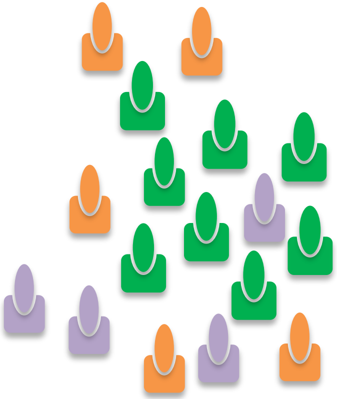

Undergraduate friends

Graduate school friends

Professional colleagues

Privacy setting: “**all friends**” !!

Need to revisit old privacy settings

2019

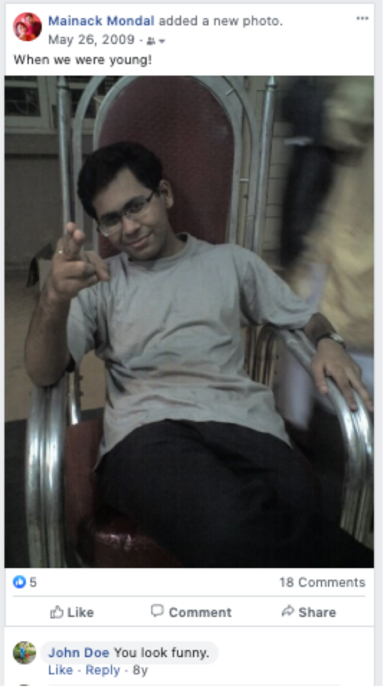

Mainack Mondal added a new photo.
May 26, 2009 ·

When we were young!

5 18 Comments

John Doe You look funny.
Like · Reply · 8y

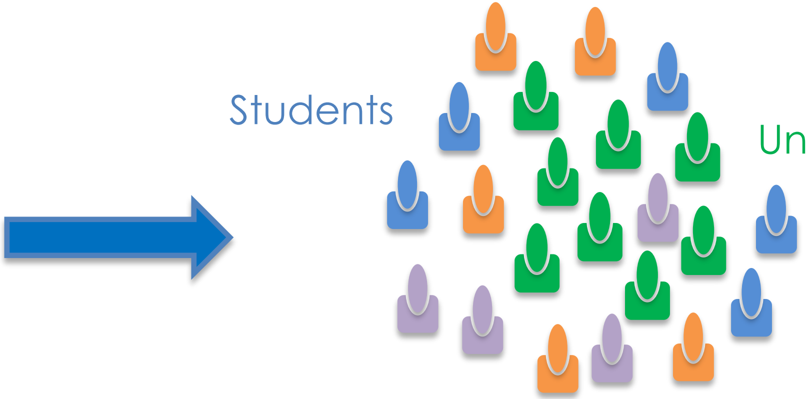

Students

Undergraduate friends

Graduate school friends

Professional colleagues

Privacy setting: “**all friends**” !!

Issue: Users take a “**set-it-and-forget-it**” approach to privacy settings for social media posts

Need: Retrospectively manage privacy

Retrospective privacy management is difficult

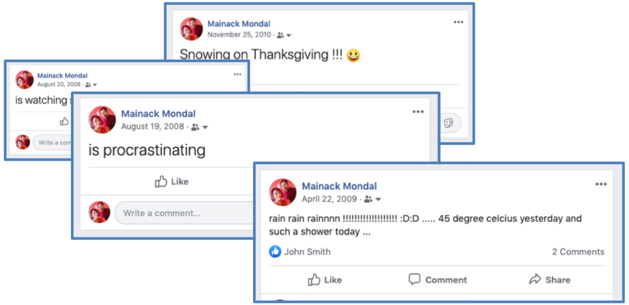

Mainack Mondal
November 25, 2010

Snowing on Thanksgiving !!! 😃

Mainack Mondal
August 20, 2008

is watching

Mainack Mondal
August 19, 2008

is procrastinating

Mainack Mondal
April 22, 2009

rain rain rainnnn !!!!!!!!!!!!!!!!!! :D:D ..... 45 degree celcius yesterday and such a shower today ...

John Smith

2 Comments

Undergraduate friends

Retrospective privacy management is difficult

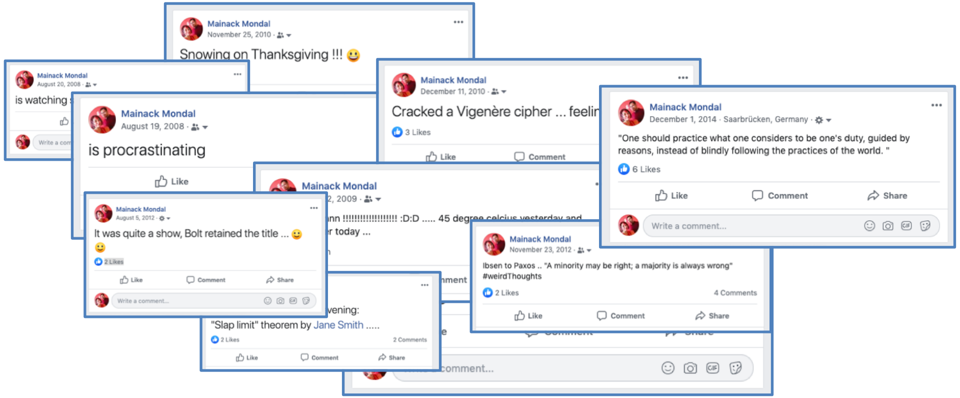

Mainack Mondal
November 25, 2010
Snowing on Thanksgiving !!! 😀

Mainack Mondal
August 20, 2008
is watching

Mainack Mondal
August 19, 2008
is procrastinating

Mainack Mondal
December 11, 2010
Cracked a Vigenère cipher ... feelin
3 Likes

Mainack Mondal
December 1, 2014 · Saarbrücken, Germany
"One should practice what one considers to be one's duty, guided by reasons, instead of blindly following the practices of the world. "
6 Likes

Mainack Mondal
2, 2009
nn !!!!!!!!!!!!!!!!!! :D:D ..... 45 degree celcius yesterday and
r today ...

Mainack Mondal
August 5, 2012
It was quite a show, Bolt retained the title ... 🙂 🙂
2 Likes

vening:
"Slap limit" theorem by Jane Smith .....
2 Likes 2 Comments

Mainack Mondal
November 23, 2012
Ibsen to Paxos .. "A minority may be right; a majority is always wrong" #weirdThoughts
2 Likes 4 Comments

Like
Comment
Share
Write a comment...

Students

Undergraduate friends

Graduate school friends

Professional colleagues

Retrospective privacy management is difficult

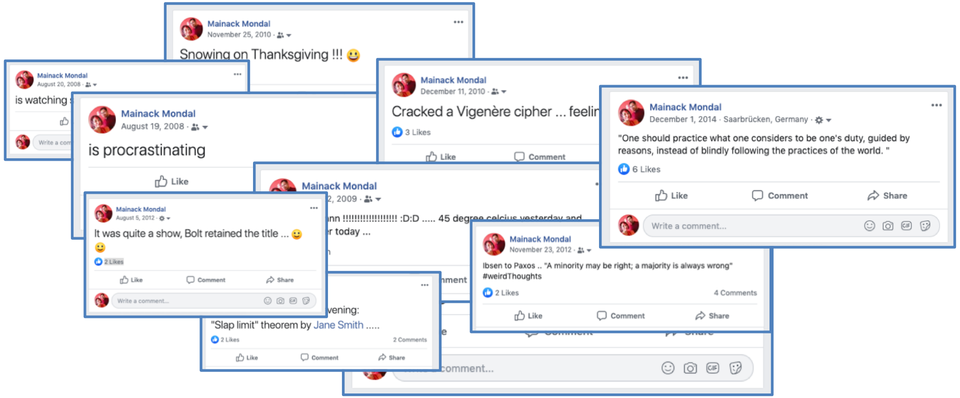

Mainack Mondal
November 25, 2010
Snowing on Thanksgiving !!! 😀

Mainack Mondal
August 20, 2008
is watching

Write a com

Mainack Mondal
August 19, 2008
is procrastinating

Like

Mainack Mondal
December 11, 2010
Cracked a Vigenère cipher ... feelin

3 Likes

Like Comment

Mainack Mondal
December 1, 2014 · Saarbrücken, Germany
"One should practice what one considers to be one's duty, guided by reasons, instead of blindly following the practices of the world. "

6 Likes

Like Comment Share

Write a comment...

Mainack Mondal
2, 2009
nn !!!!!!!!!!!!!!!!!!! :D:D ..... 45 degree celcius yesterday and
r today ...

mment

comment...

Mainack Mondal
August 5, 2012
It was quite a show, Bolt retained the title ... 🙂 🙂

2 Likes

Like Comment Share

Write a comment...

vening:
"Slap limit" theorem by Jane Smith .....

2 Likes 2 Comments

Like Comment Share

Mainack Mondal
November 23, 2012
Ibsen to Paxos .. "A minority may be right; a majority is always wrong" #weirdThoughts

2 Likes 4 Comments

Like Comment Share

Students

Undergraduate friends

Graduate school friends

Professional colleagues

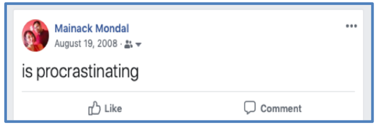

Mainack Mondal
August 19, 2008
is procrastinating

Like Comment

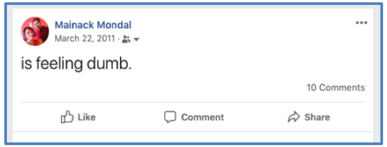

Mainack Mondal
March 22, 2011
is feeling dumb.

10 Comments

Like Comment Share

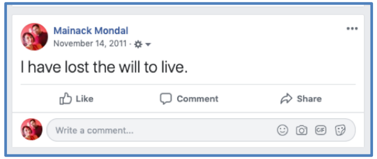

Mainack Mondal
November 14, 2011
I have lost the will to live.

Like Comment Share

Write a comment...

State of the art

No proposal for a predictive model or mechanism

[Bauer et al. 2013]
[Ayalon et al. 2013]

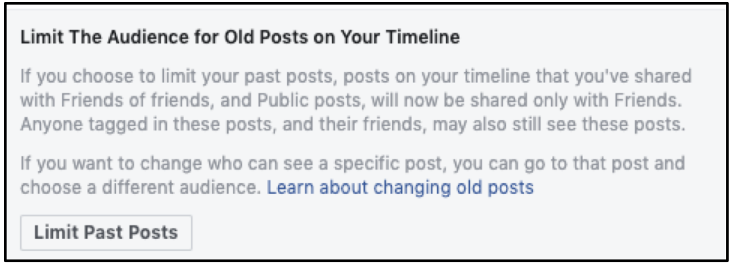

Limit The Audience for Old Posts on Your Timeline

If you choose to limit your past posts, posts on your timeline that you've shared with Friends of friends, and Public posts, will now be shared only with Friends. Anyone tagged in these posts, and their friends, may also still see these posts.

If you want to change who can see a specific post, you can go to that post and choose a different audience. Learn about changing old posts

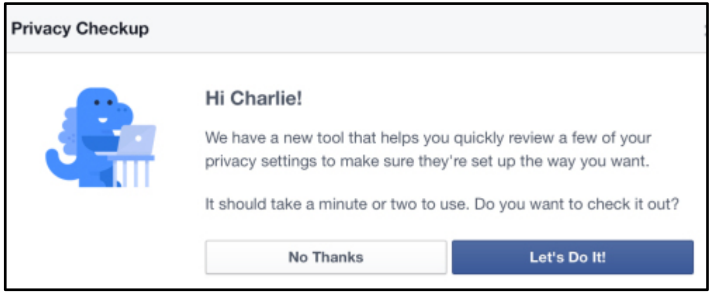

Privacy Checkup

Hi Charlie!

We have a new tool that helps you quickly review a few of your privacy settings to make sure they're set up the way you want.

It should take a minute or two to use. Do you want to check it out?

State of the art

No proposal for a predictive model or mechanism

[Bauer et al. 2013]
[Ayalon et al. 2013]

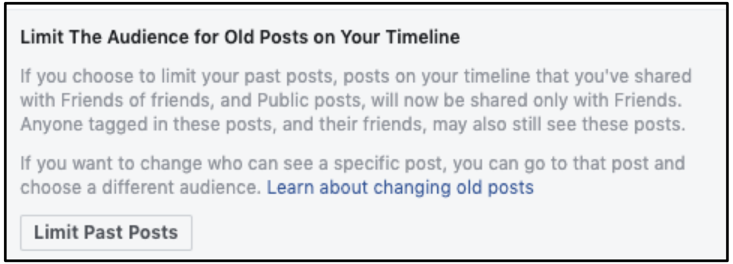

Limit The Audience for Old Posts on Your Timeline

If you choose to limit your past posts, posts on your timeline that you've shared with Friends of friends, and Public posts, will now be shared only with Friends. Anyone tagged in these posts, and their friends, may also still see these posts.

If you want to change who can see a specific post, you can go to that post and choose a different audience. Learn about changing old posts

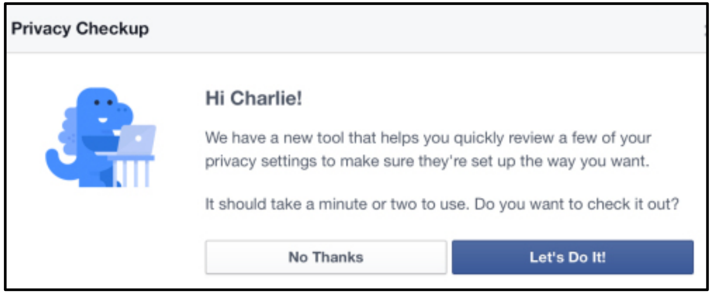

Privacy Checkup

Hi Charlie!

We have a new tool that helps you quickly review a few of your privacy settings to make sure they're set up the way you want.

It should take a minute or two to use. Do you want to check it out?

Focus of our study

- Measure privacy activity and preferences
- Predictive models for retrospective privacy management

Assisting users in temporal privacy management

Our data collection approach

Privacy settings and friend network over time

Preferences for changing privacy settings

Automated classifiers

Assisting users in temporal privacy management

Our data collection approach

Privacy settings and friend network over time

Preferences for changing privacy settings

Automated classifiers

Study overview

amazon
mechanical turk

Privacy-preserving data-collection Infrastructure

78 Facebook users

Two surveys

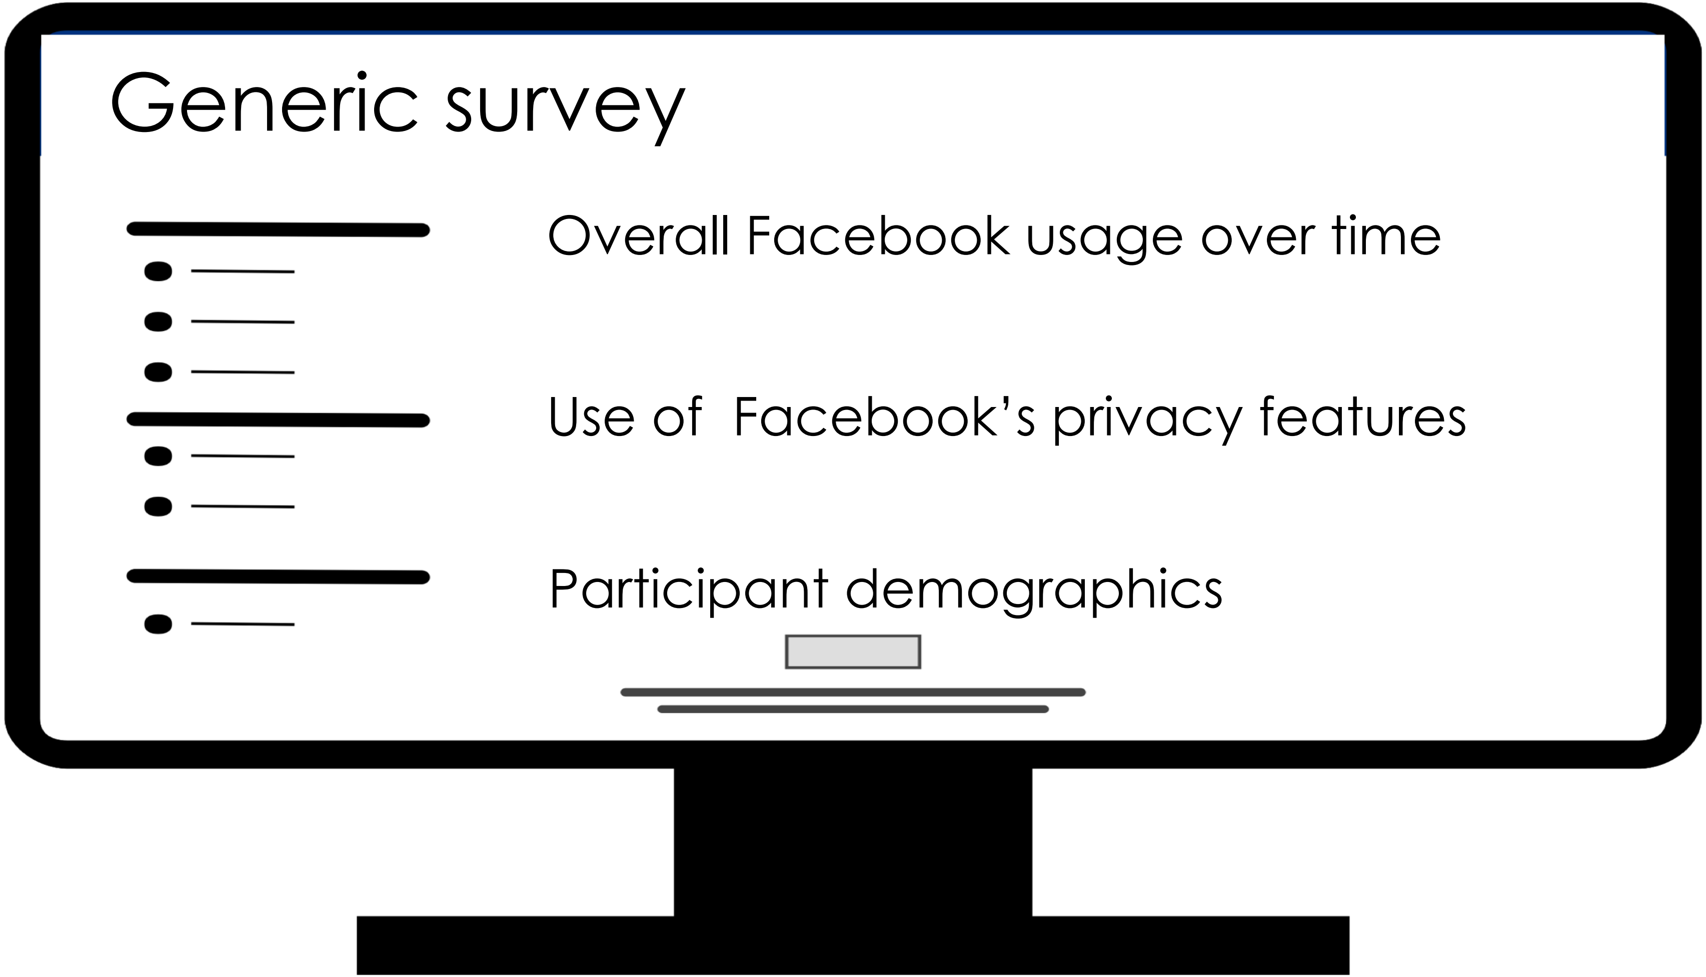

Generic survey

Overall Facebook usage over time

Use of Facebook's privacy features

Participant demographics

Consent process

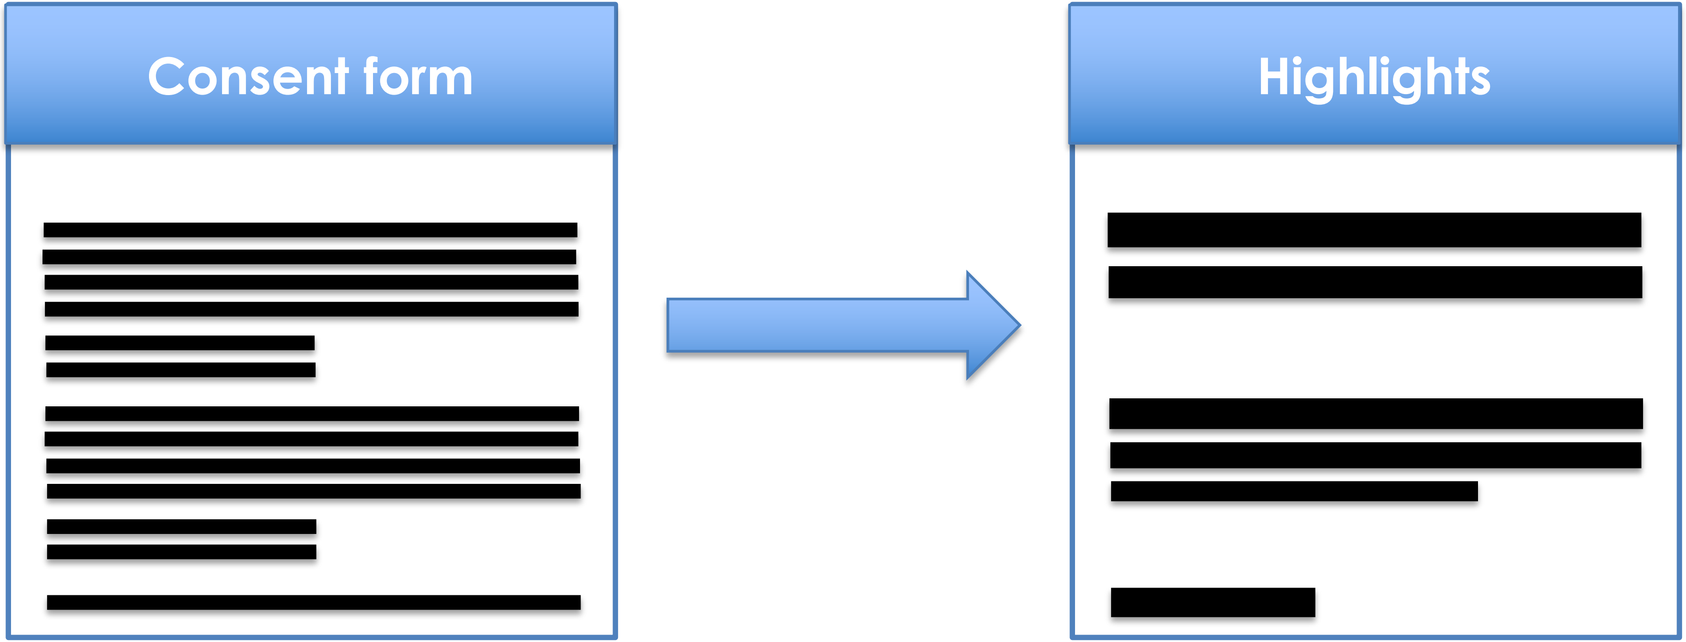

Consent form

Highlights

Data collection process

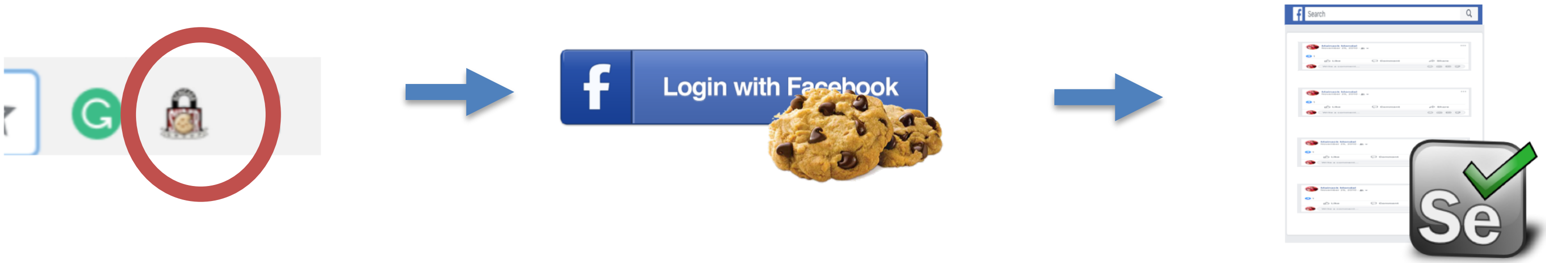

Data collection process

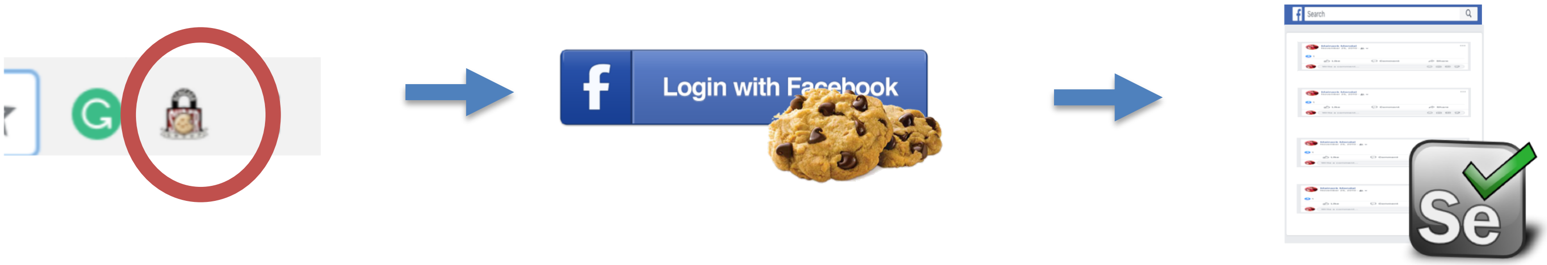

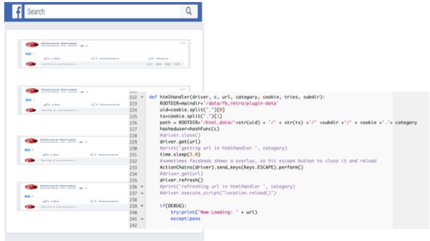

```python
def htmlHandler(driver, c, url, category, cookie, tries, subdir):
    ROOTDIR=maindir+'/data/fb_retro/plugin-data'
    uid=cookie.split(".")[0]
    ts=cookie.split(".")[1]
    path = ROOTDIR+'/html_data/'+str(uid) + '/' + str(ts) +'/' +subdir +'/' + cookie +'.'+ category
    hasheduser=hashFunc(c)
    #driver.close()
    driver.get(url)
    #print('getting url in htmlHandler ', category)
    time.sleep(5.0)
    #sometimes facebook shows a overlay, so hit escape button to close it and reload
    ActionChains(driver).send_keys(Keys.ESCAPE).perform()
    #driver.get(url)
    driver.refresh()
    #print('refreshing url in htmlHandler ', category)
    #driver.execute_script("location.reload()")

    if(DEBUG):
        try:print('Now Loading: ' + url)
        except:pass
```

Programmatic
No humans ever view raw HTML

Hash names and IDs;
No images collected

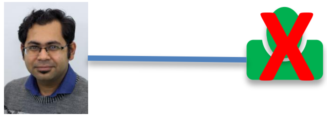

Never access
friends' profiles

Facebook Timeline data

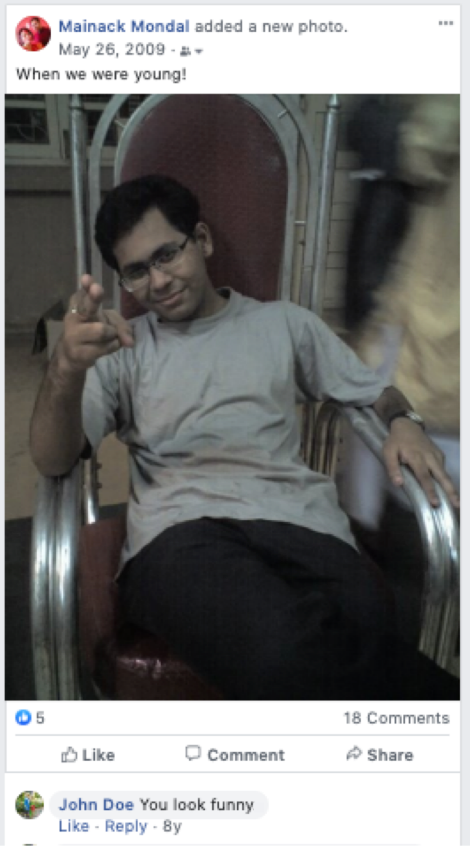

Mainack Mondal added a new photo.
May 26, 2009 ·

When we were young!

5

18 Comments

John Doe You look funny

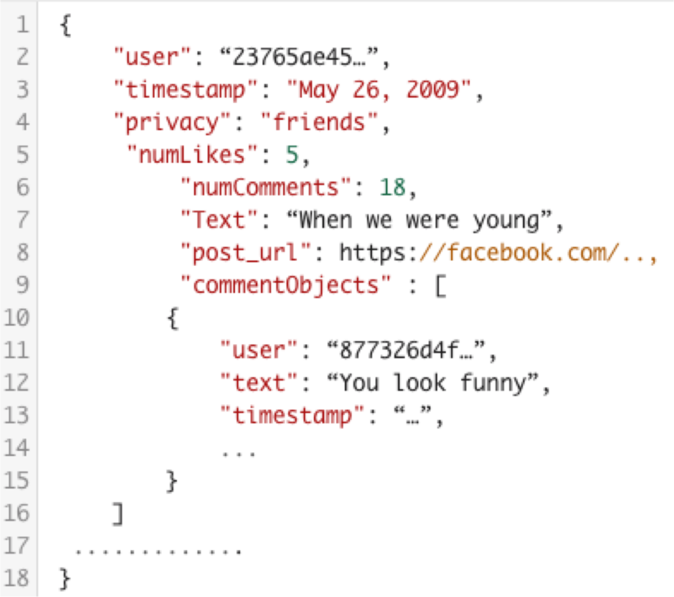

```
{
        "user": “23765ae45…”,
        "timestamp": "May 26, 2009",
        "privacy": "friends",
         "numLikes": 5,
             "numComments": 18,
             "Text": “When we were young”,
             "post_url": https://facebook.com/..,
             "commentObjects" : [
            {
                "user": “877326d4f…”,
                "text": “You look funny”,
                "timestamp": “…”,
                ...
            }
        ]
  .............
}
```

Chose not to store images!

Facebook Activity Log data

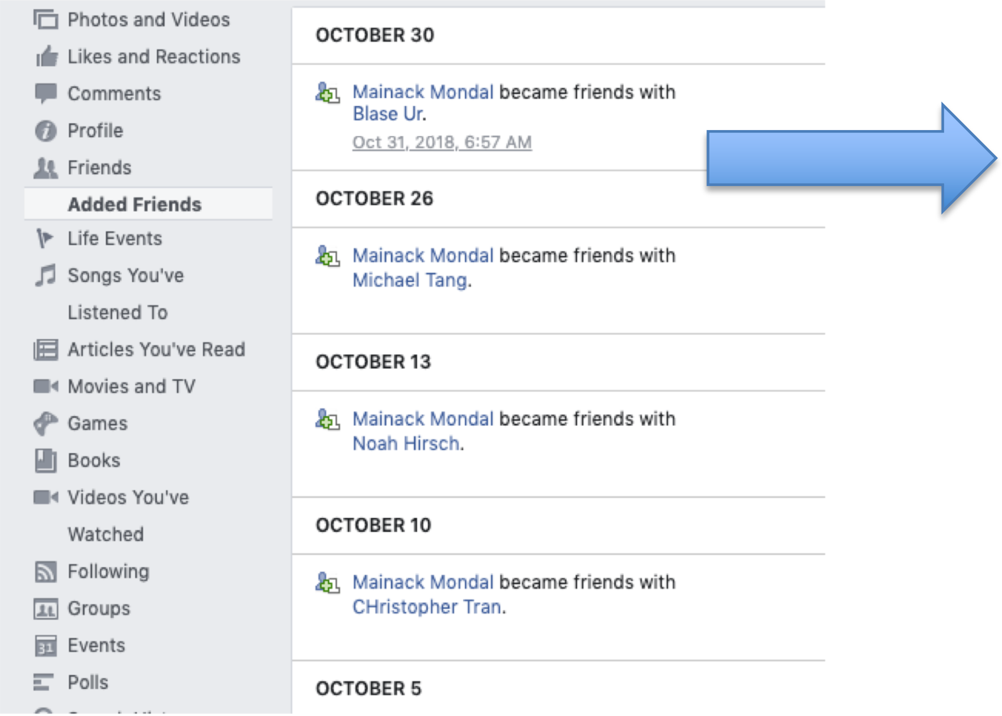

OCTOBER 30

Mainack Mondal became friends with Blase Ur.

Oct 31, 2018, 6:57 AM

OCTOBER 26

Mainack Mondal became friends with Michael Tang.

OCTOBER 13

Mainack Mondal became friends with Noah Hirsch.

OCTOBER 10

Mainack Mondal became friends with CHristopher Tran.

OCTOBER 5

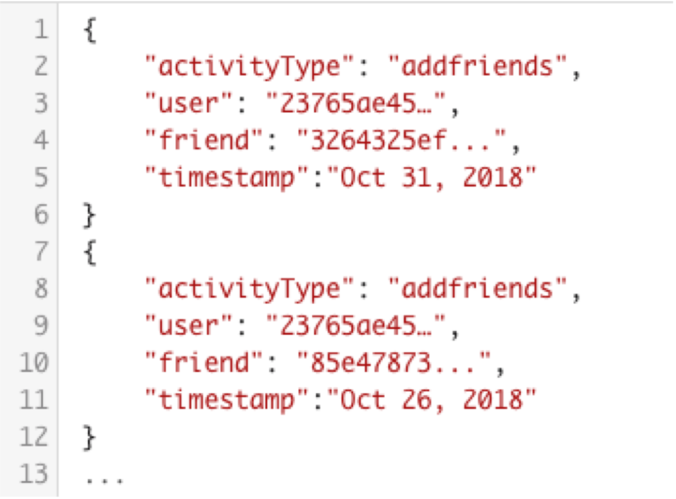

```
{
     "activityType": "addfriends",
     "user": "23765ae45…",
     "friend": "3264325ef...",
     "timestamp":"Oct 31, 2018"
}
{
     "activityType": "addfriends",
     "user": "23765ae45…",
     "friend": "85e47873...",
     "timestamp":"Oct 26, 2018"
}
...
```

ALL Facebook activities by user (friendship, likes, comments,…)

Post-specific survey

1. Desired privacy settings for 5 random posts per user

Post-specific survey

1. Desired privacy settings for 5 random posts per user

Click here to see Post 1. Current privacy setting: Public

Post-specific survey

1. Desired privacy settings for 5 random posts per user

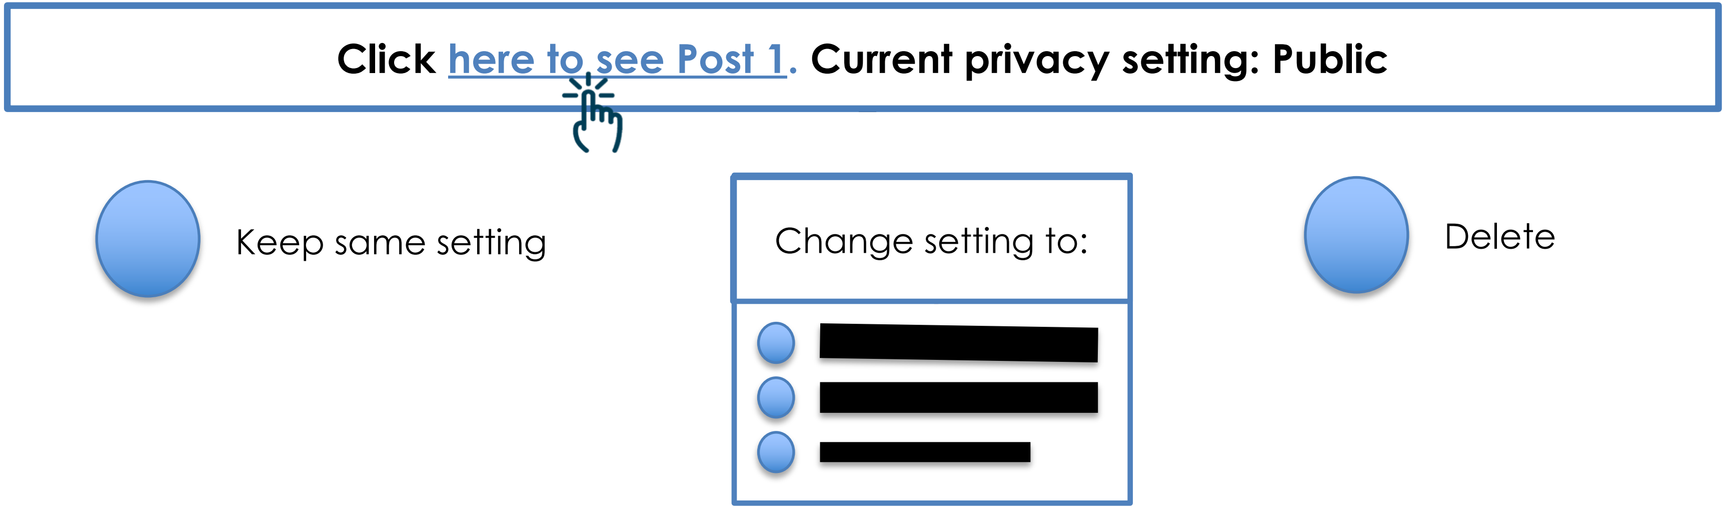

Click here to see Post 1. Current privacy setting: Public

○ Keep same setting

Change setting to:

○ [redacted]

○ [redacted]

○ [redacted]

○ Delete

Post-specific survey

1. Desired privacy settings for 5 random posts per user

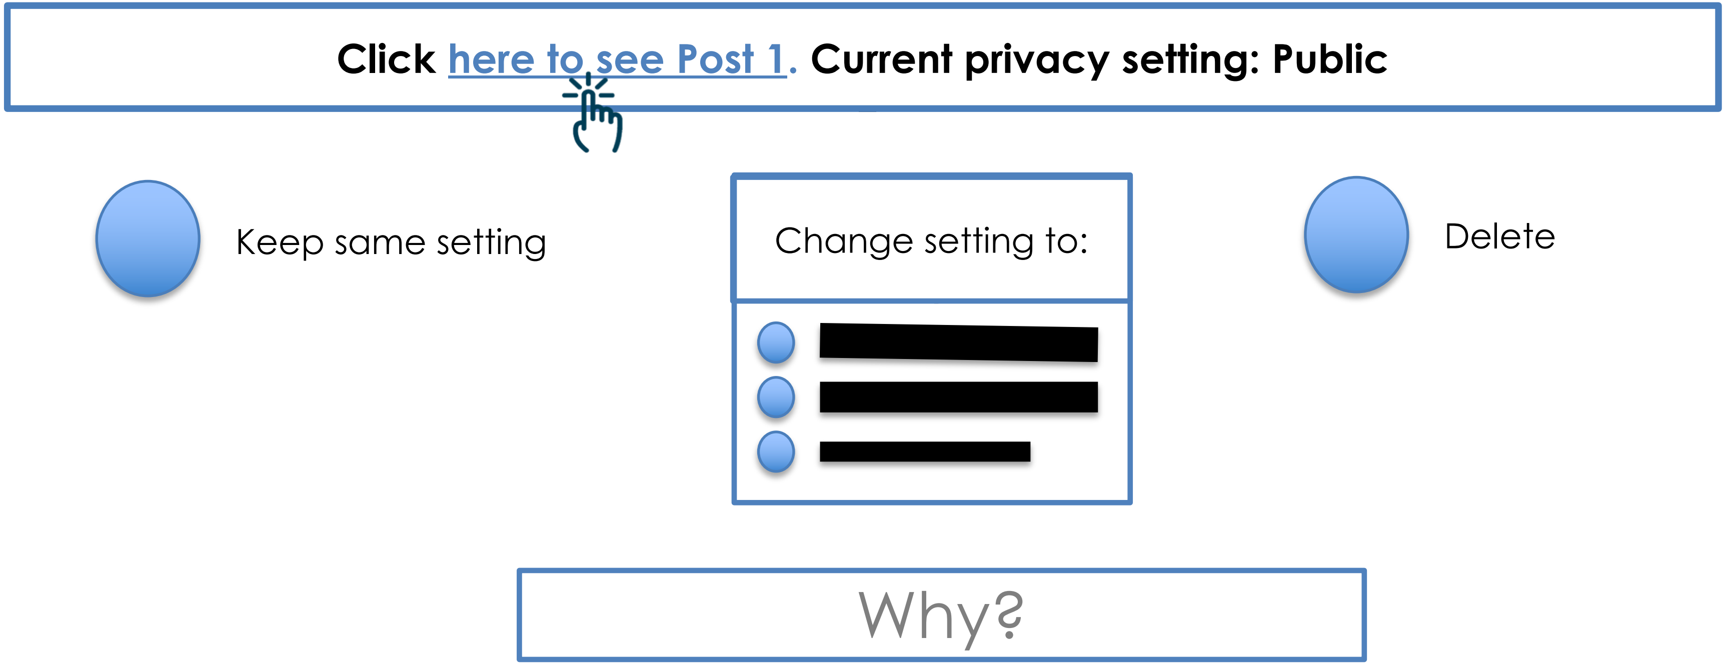

Click here to see Post 1. Current privacy setting: Public

○ Keep same setting

Change setting to:

○ [redacted]

○ [redacted]

○ [redacted]

○ Delete

Why?

Post-specific survey

1. Desired privacy settings for 5 random posts per user
2. Desired privacy settings for 6 specific friends per post

Post-specific survey

1. Desired privacy settings for 5 random posts per user
2. Desired privacy settings for 6 specific friends per post

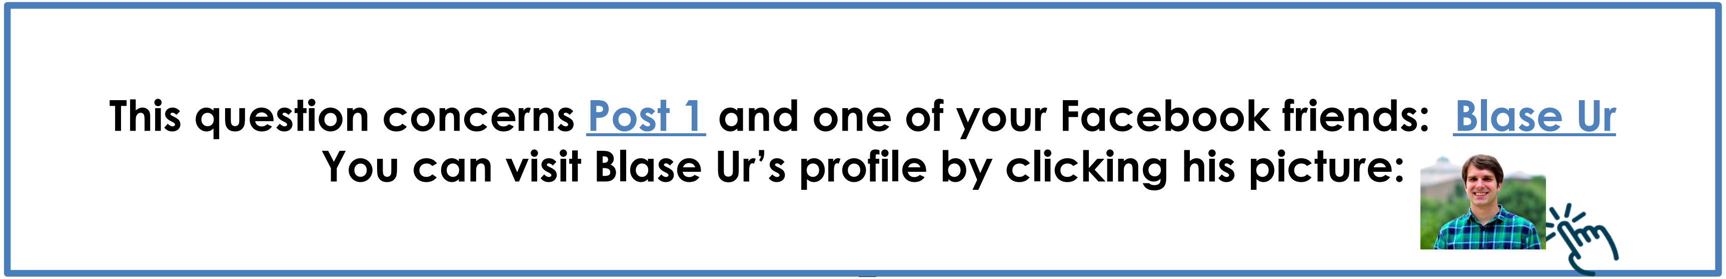

This question concerns Post 1 and one of your Facebook friends: Blase Ur
You can visit Blase Ur's profile by clicking his picture:

Post-specific survey

1. Desired privacy settings for 5 random posts per user
2. Desired privacy settings for 6 specific friends per post

This question concerns Post 1 and one of your Facebook friends: Blase Ur
You can visit Blase Ur's profile by clicking his picture:

Keep sharing post 1 with Blase Ur ○ ○ ○ Stop sharing post 1 with Blase Ur

Post-specific survey

1. Desired privacy settings for 5 random posts per user
2. Desired privacy settings for 6 specific friends per post

This question concerns Post 1 and one of your Facebook friends: Blase Ur
You can visit Blase Ur's profile by clicking his picture:

Keep sharing post 1 with Blase Ur ○ ○ ○ Stop sharing post 1 with Blase Ur

Why?

Demographics

AMT workers from US

69% identified as female

46% reported age 25-34

18% reported CS background

Facebook usage

	Total	Median
Account age (Years)	-	10
#Friends	-	224
#Timeline posts	253,122	1,840
#Activity log entries	1,738,303	20,263

Facebook usage

	Total	Median
Account age (Years)	-	10
#Friends	-	224
#Timeline posts	253,122	1,840
#Activity log entries	1,738,303	20,263

Active users with old accounts and lots of posts

Facebook usage

	Total	Median
Account age (Years)	-	10
#Friends	-	224
#Timeline posts	253,122	1,840
#Activity log entries	1,738,303	20,263

Active users with old accounts and lots of posts

67% reported reduced Facebook usage over time

Assisting users in temporal privacy management

Our data collection approach

Privacy settings and friend network over time

Preferences for changing privacy settings

Automated classifiers

Privacy settings over time

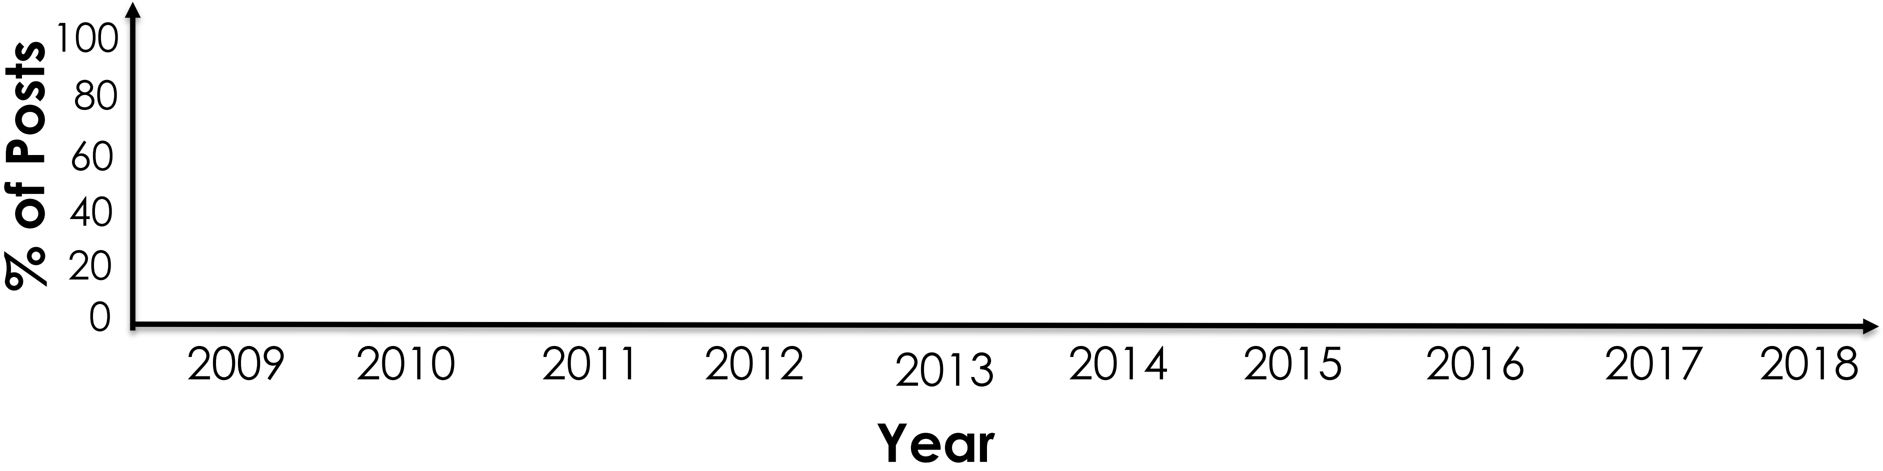

Privacy settings over time

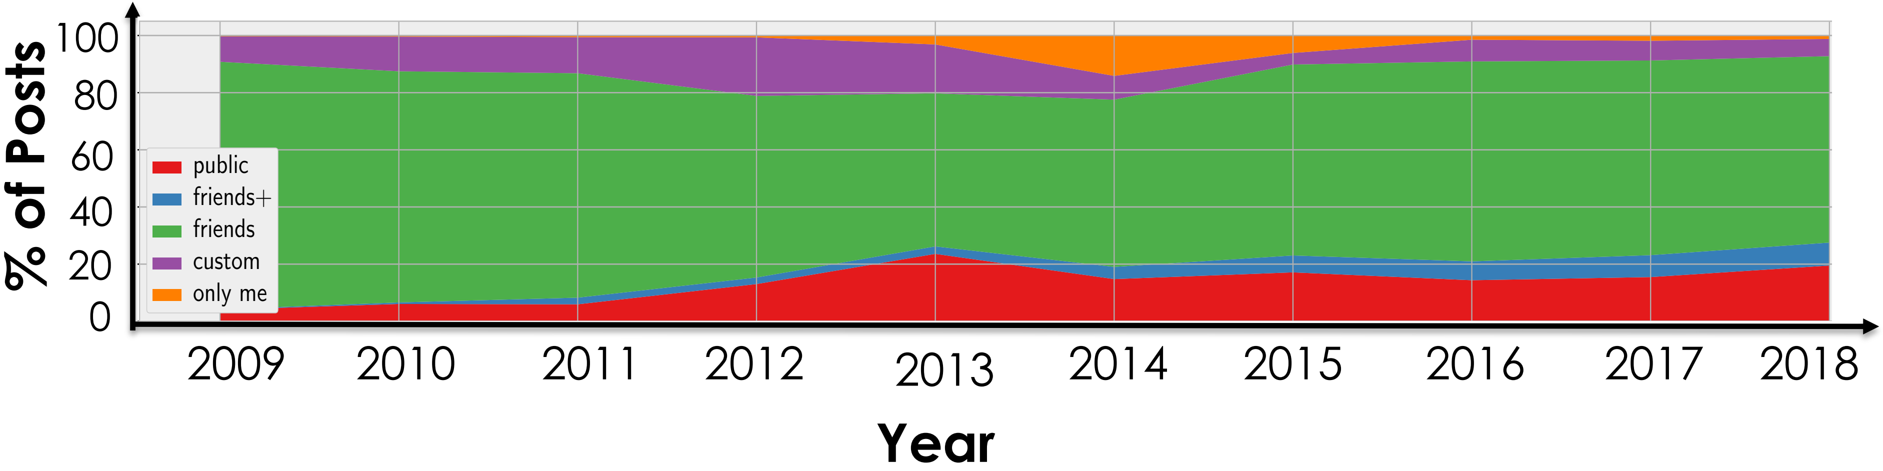

Privacy settings over time

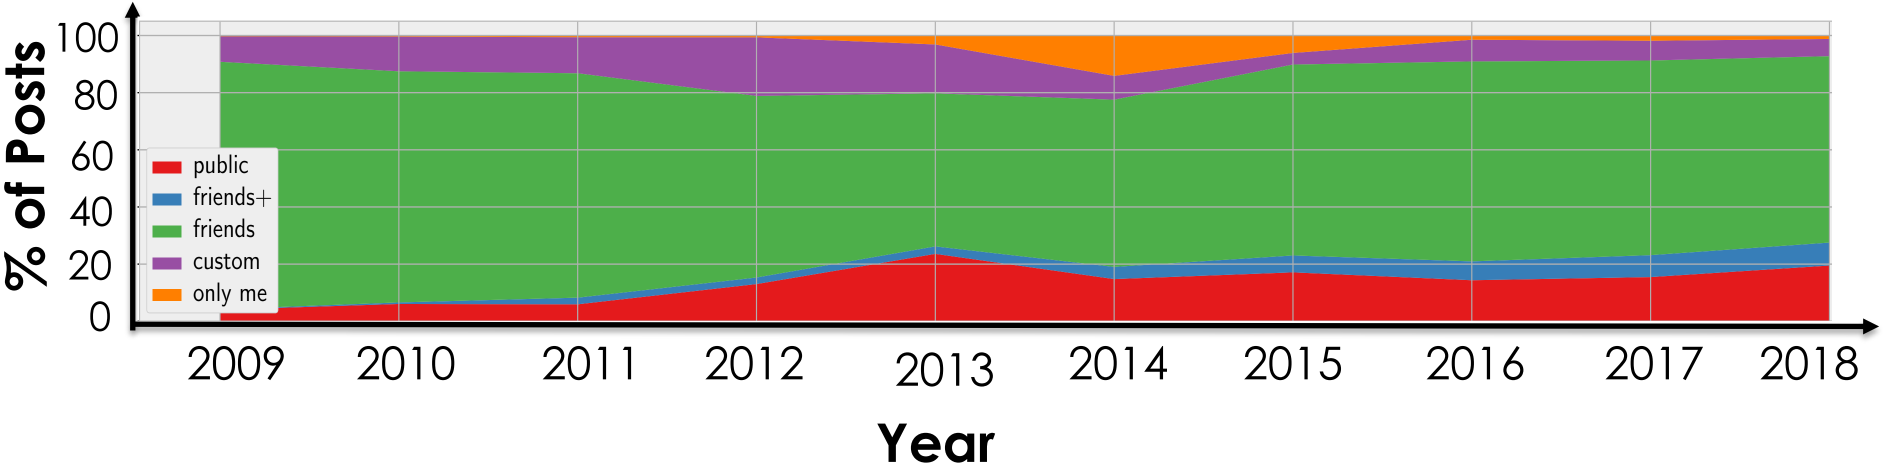

Majority of old posts are shared with all "friends"

Change in number of friends

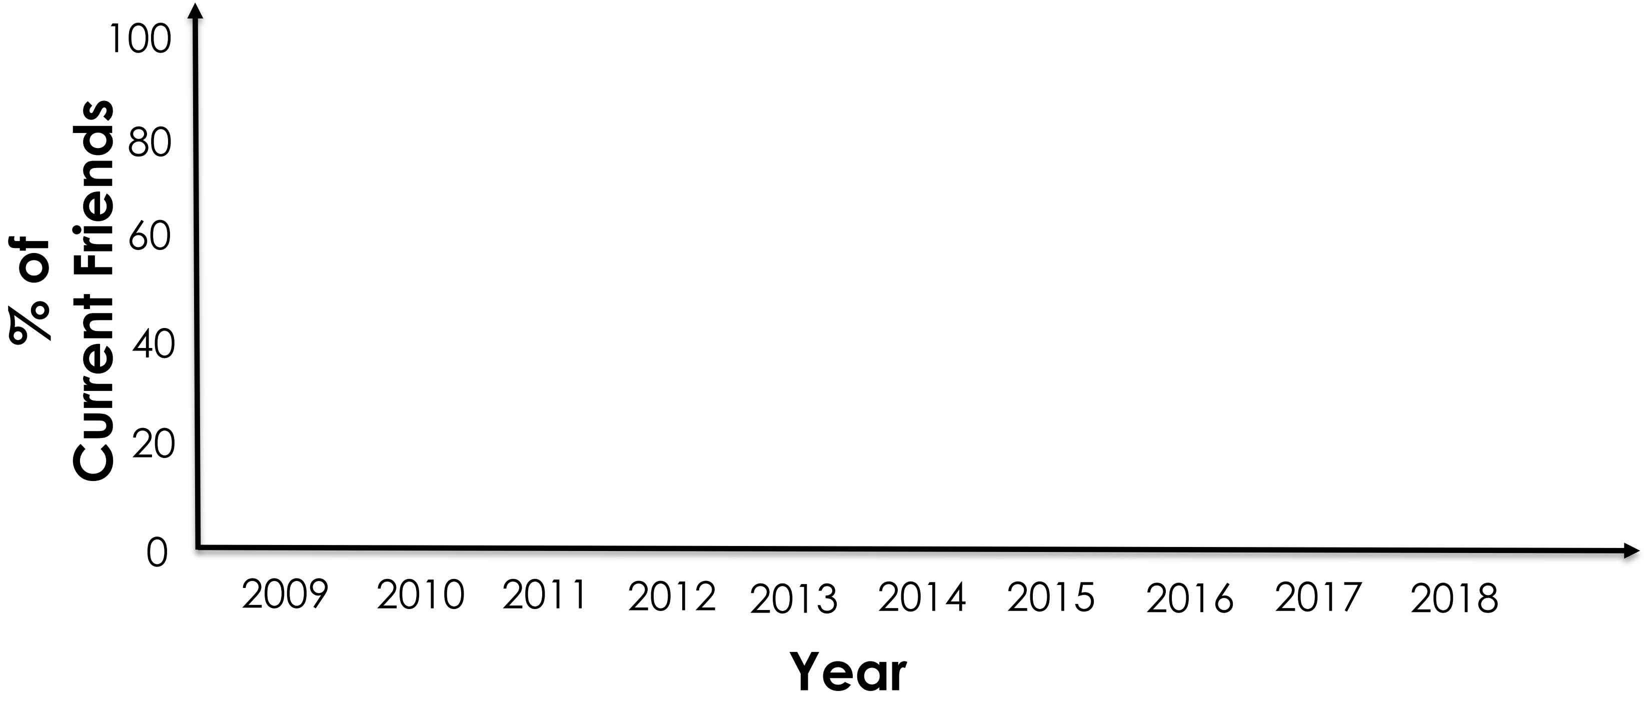

Change in number of friends

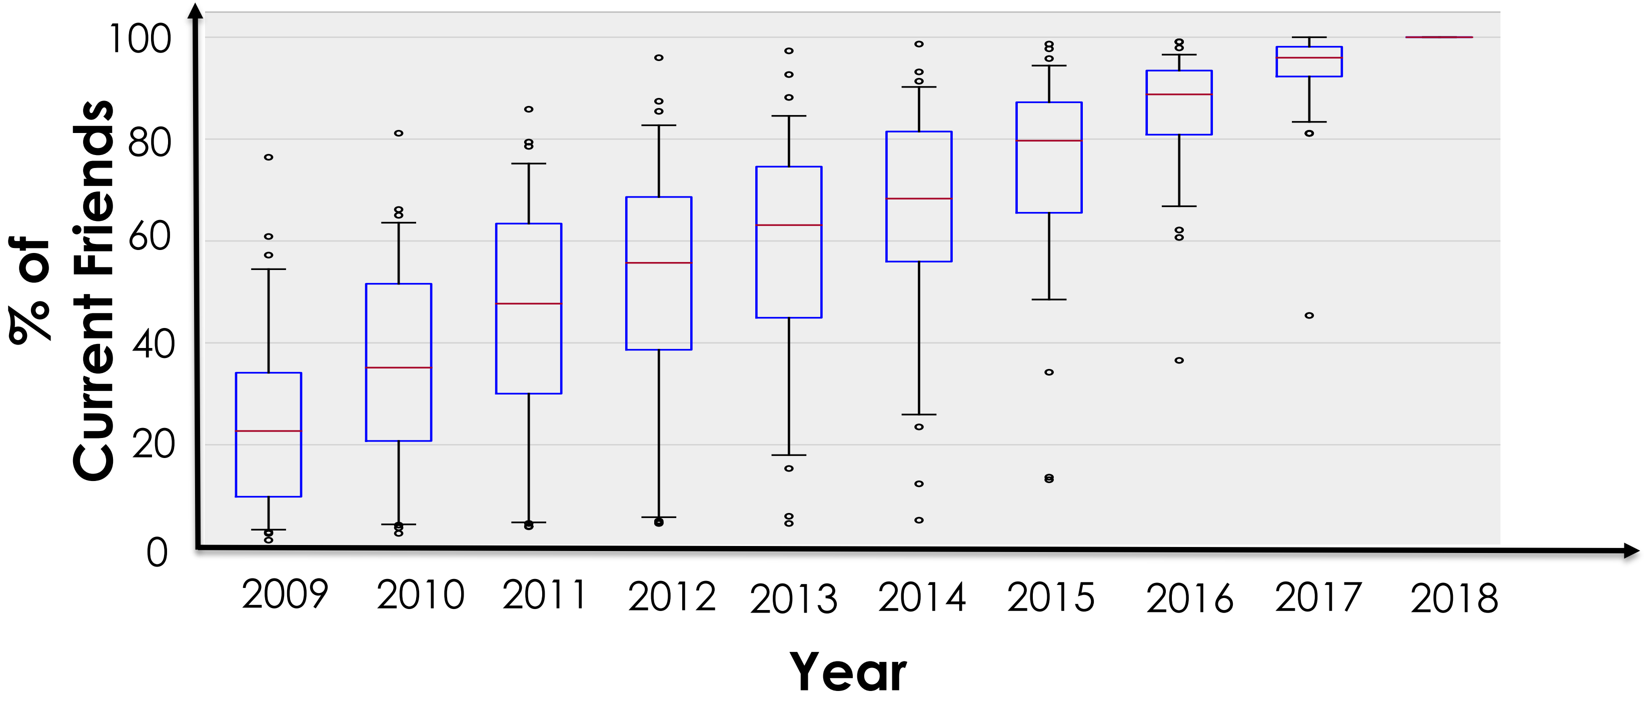

Change in number of friends

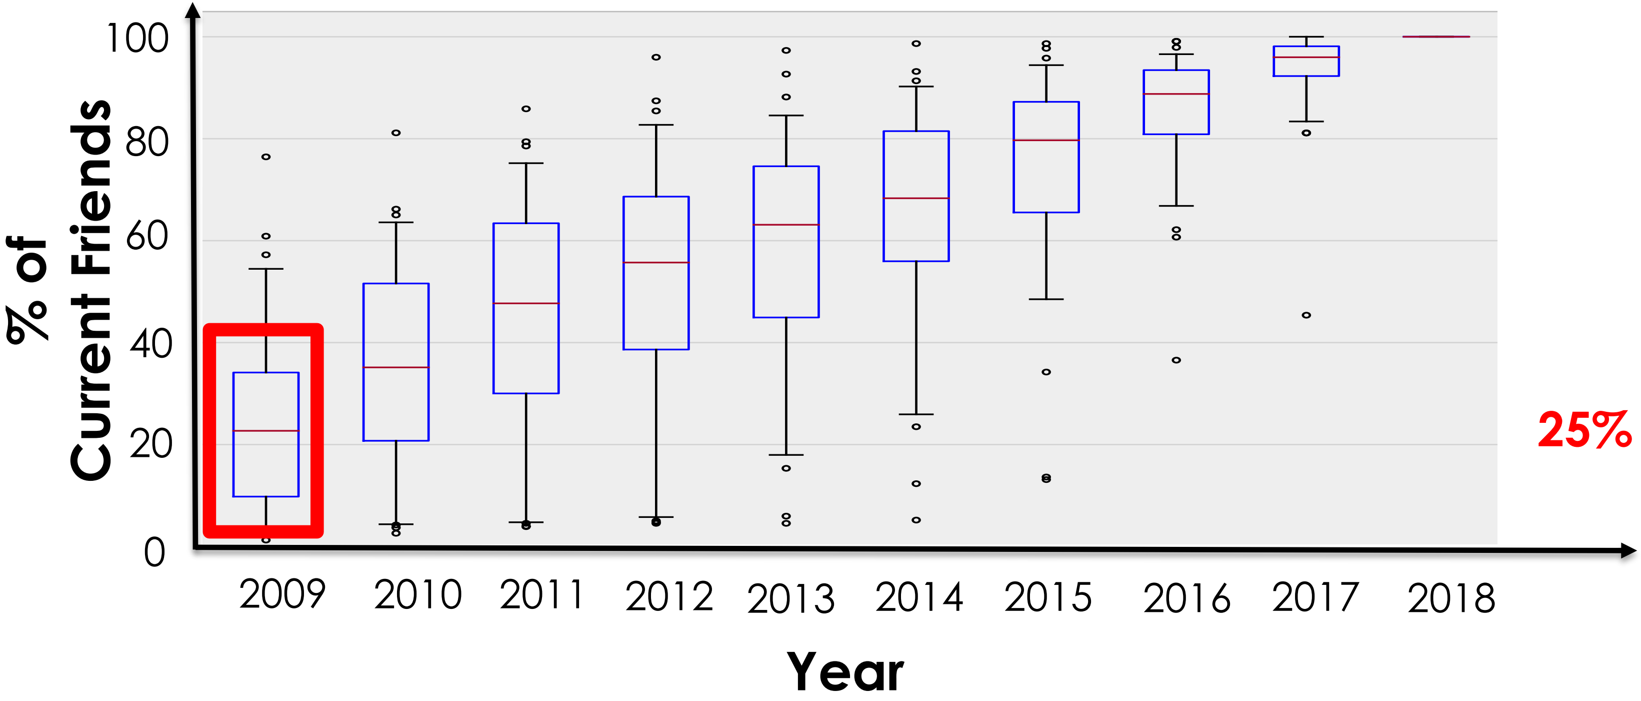

Change in number of friends

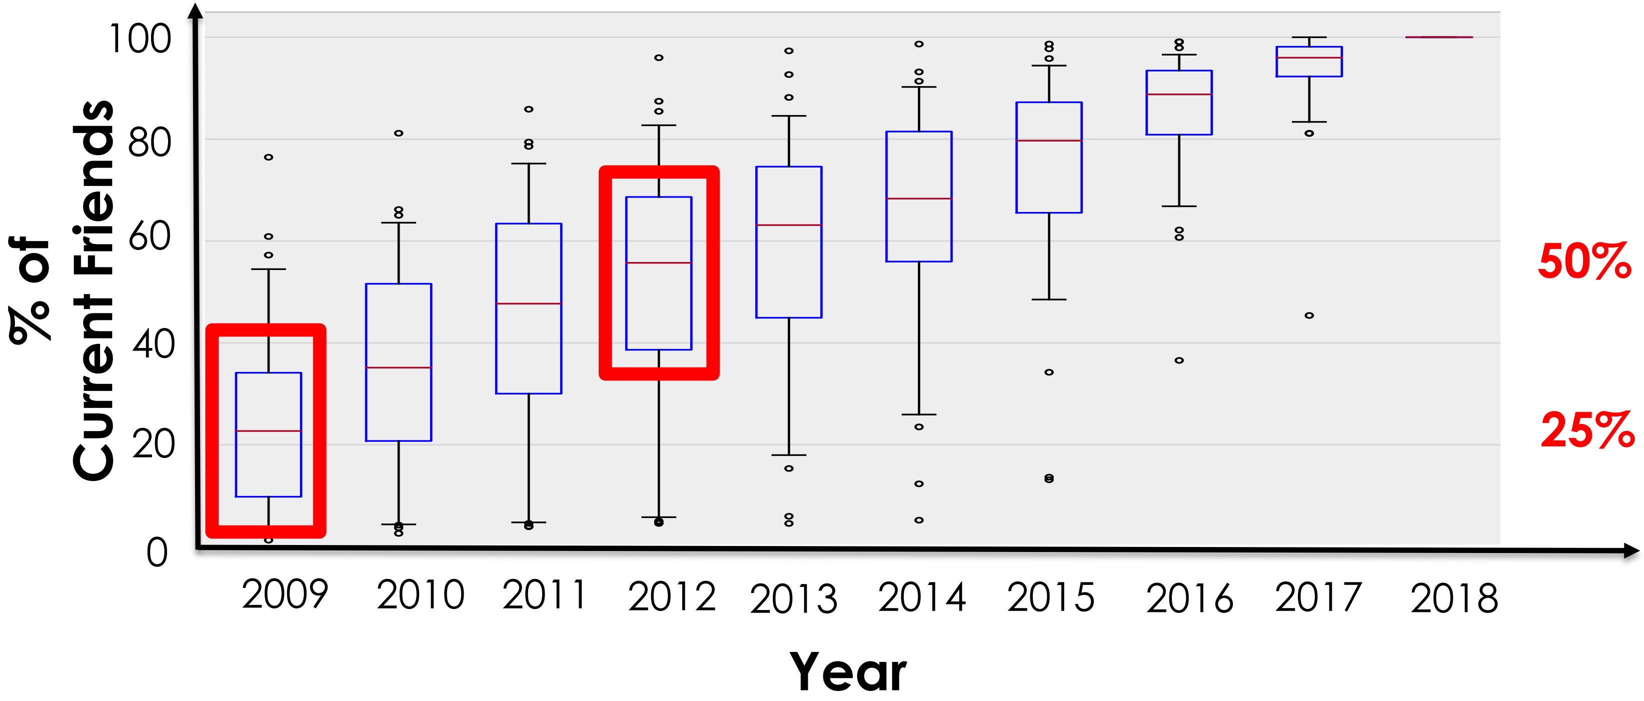

Change in number of friends

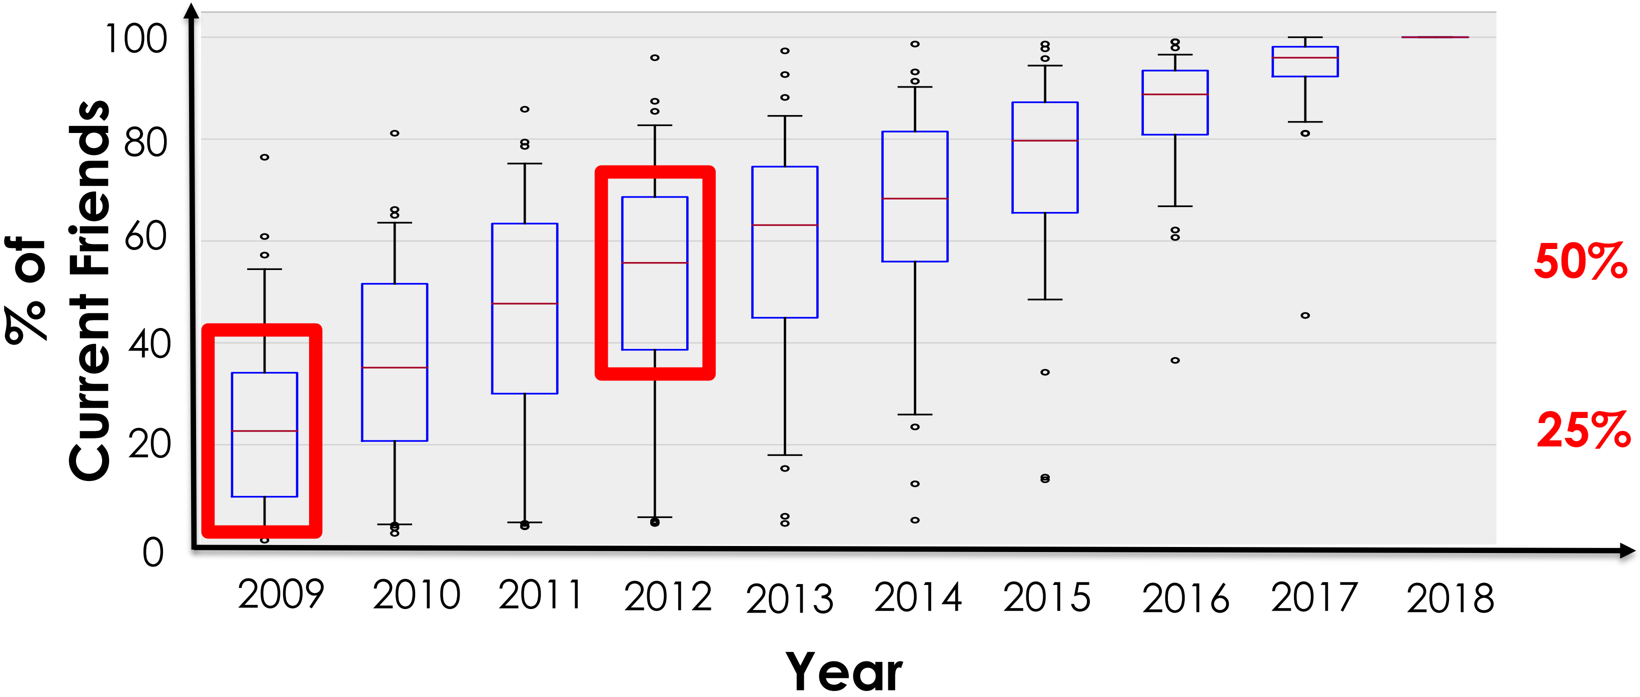

Substantial change in the meaning of “friends” privacy setting

Assisting users in temporal privacy management

Our data collection approach

Privacy settings and friend network over time

Preferences for changing privacy settings

Automated classifiers

Desired privacy setting for old posts

Post-specific survey: **Desired privacy setting for 390 random posts**

Desired privacy setting for old posts

Post-specific survey: **Desired privacy setting for 390 random posts**

	Desired setting						
Current setting	**Public**	**Friends+**	**Friends**	**Custom**	**Only Me**	**Custom (Decreased)**	**Delete**
Public	**58**	-	**3**	-	-	-	**1**
Friends+	**3**	**27**	**3**	-	-	-	-
Friends	**21**	**4**	**177**	**3**	**5**	-	**31**
Custom	**6**	**2**	**9**	**19**	**1**	**2**	**4**
Only Me	-	-	-	-	**9**	-	**1**

Desired privacy setting for old posts

Post-specific survey: **Desired privacy setting for 390 random posts**

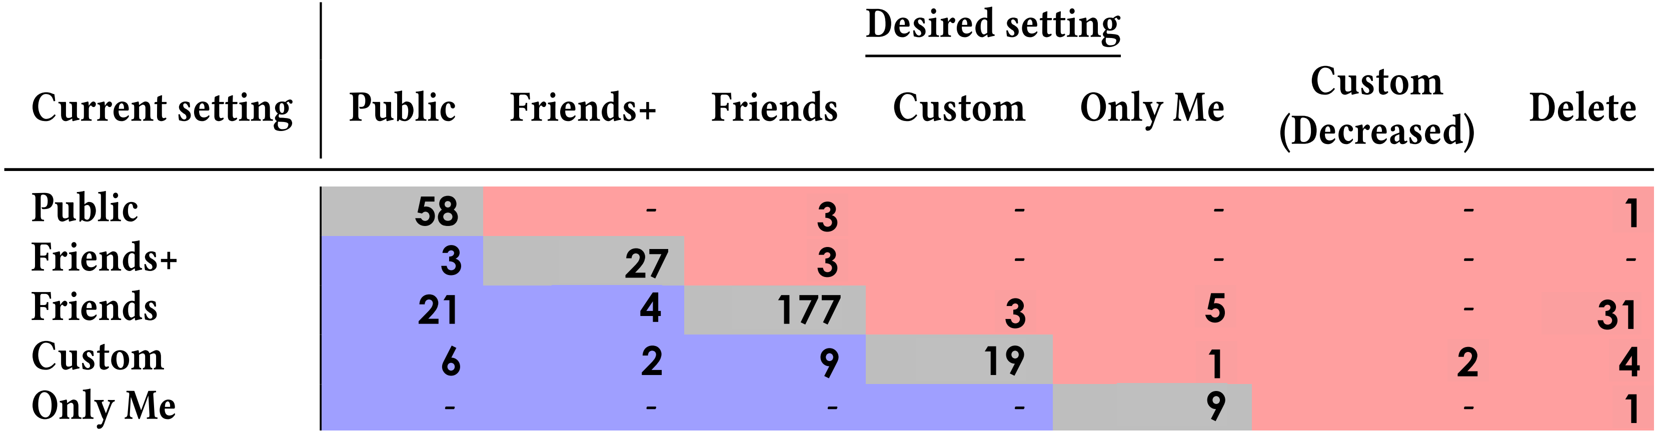

	Desired setting						
Current setting	Public	Friends+	Friends	Custom	Only Me	Custom (Decreased)	Delete
Public	58	-	3	-	-	-	1
Friends+	3	27	3	-	-	-	-
Friends	21	4	177	3	5	-	31
Custom	6	2	9	19	1	2	4
Only Me	-	-	-	-	9	-	1

Participants **desire to change audience for 25%** of old posts!

Desired privacy setting for old posts

Post-specific survey: **Desired privacy setting for 390 random posts**

Desire to limit audience: 54 posts

Desired setting

Current setting	Public	Friends+	Friends	Custom	Only Me	Custom (Decreased)	Delete
Public	58	-	3	-	-	-	1
Friends+	3	27	3	-	-	-	-
Friends	21	4	177	3	5	-	31
Custom	6	2	9	19	1	2	4
Only Me	-	-	-	-	9	-	1

Desire to increase audience: 45 posts

Participants **desire to change audience for 25%** of old posts!

Effectiveness of Facebook's privacy tools

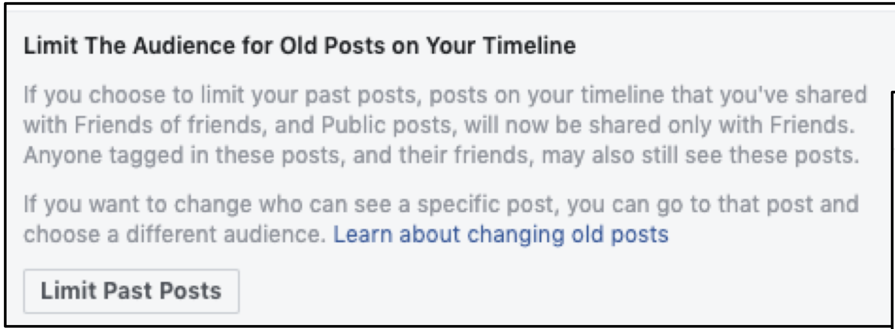

Limit The Audience for Old Posts on Your Timeline

If you choose to limit your past posts, posts on your timeline that you've shared with Friends of friends, and Public posts, will now be shared only with Friends. Anyone tagged in these posts, and their friends, may also still see these posts.

If you want to change who can see a specific post, you can go to that post and choose a different audience. Learn about changing old posts

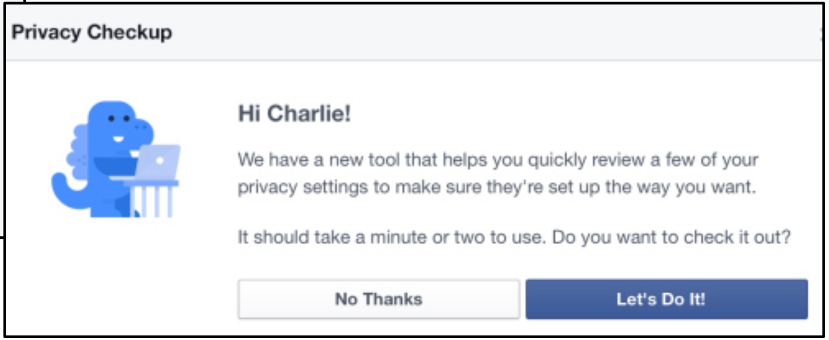

Privacy Checkup

Hi Charlie!

We have a new tool that helps you quickly review a few of your privacy settings to make sure they're set up the way you want.

It should take a minute or two to use. Do you want to check it out?

Found no significant correlation between usage of these tools and the desire to change posts' privacy settings

Assisting users in temporal privacy management

Our data collection approach

Privacy settings and friend network over time

Preferences for changing privacy settings

Automated classifiers

A human-in-the-loop design

Inspiration

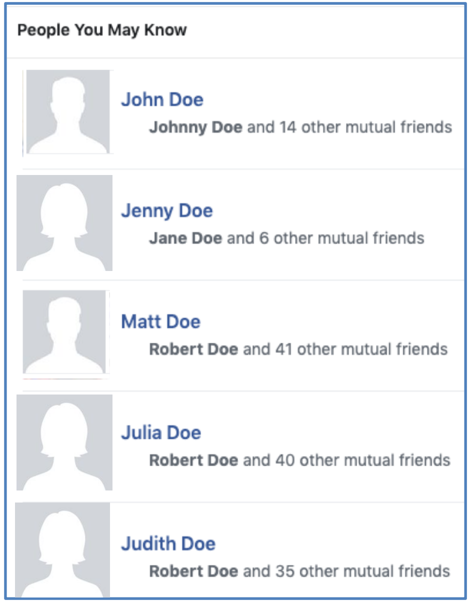

People You May Know

John Doe
Johnny Doe and 14 other mutual friends

Jenny Doe
Jane Doe and 6 other mutual friends

Matt Doe
Robert Doe and 41 other mutual friends

Julia Doe
Robert Doe and 40 other mutual friends

Judith Doe
Robert Doe and 35 other mutual friends

A human-in-the-loop design

Inspiration

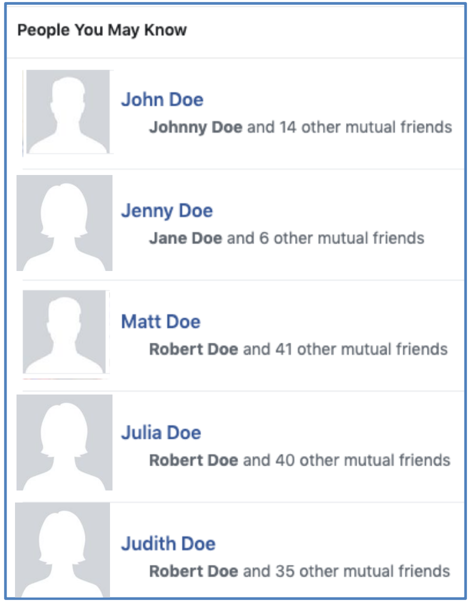

People You May Know

John Doe
Johnny Doe and 14 other mutual friends

Jenny Doe
Jane Doe and 6 other mutual friends

Matt Doe
Robert Doe and 41 other mutual friends

Julia Doe
Robert Doe and 40 other mutual friends

Judith Doe
Robert Doe and 35 other mutual friends

Our vision

Stop sharing

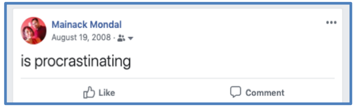

Mainack Mondal
August 19, 2008 ·

is procrastinating

with

?

Stop sharing

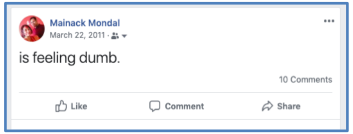

Mainack Mondal
March 22, 2011 ·

is feeling dumb.

10 Comments

with

?

Stop sharing

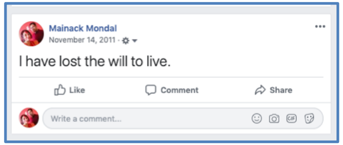

Mainack Mondal
November 14, 2011 ·

I have lost the will to live.

with

?

Prediction task

Prediction task

Predict if a user wants to “stop sharing” a given post with a given friend

Output

List of friend-post pairs ordered by probability

Ground truth

Privacy decisions for 78 participants x 5 posts X 6 friends = 2,340 pairs

Features for prediction

User-specific	#friends, age of the account, life change, Facebook privacy tool usage, user age, CS-background
Post metadata	Age of the post, #likes, #comments, previous change in privacy setting, type of post, tagged friend
Post content	Word2vec embeddings, Google content-classification categories, sentiment
Friend-specific	Days since first and last communication, #wall words exchanged, #likes from friend to user

Prediction algorithms

Supervised learning algorithms with cross validation

- **Random Forests, XGBoost**, Decision Trees, Logistic Regression, Support Vector Machines, Deep Neural Networks

Baselines

- Random: Randomly predicts “stop sharing” for a pair
- Interaction: Low interaction level → “stop sharing”

Are our models better than the baselines?

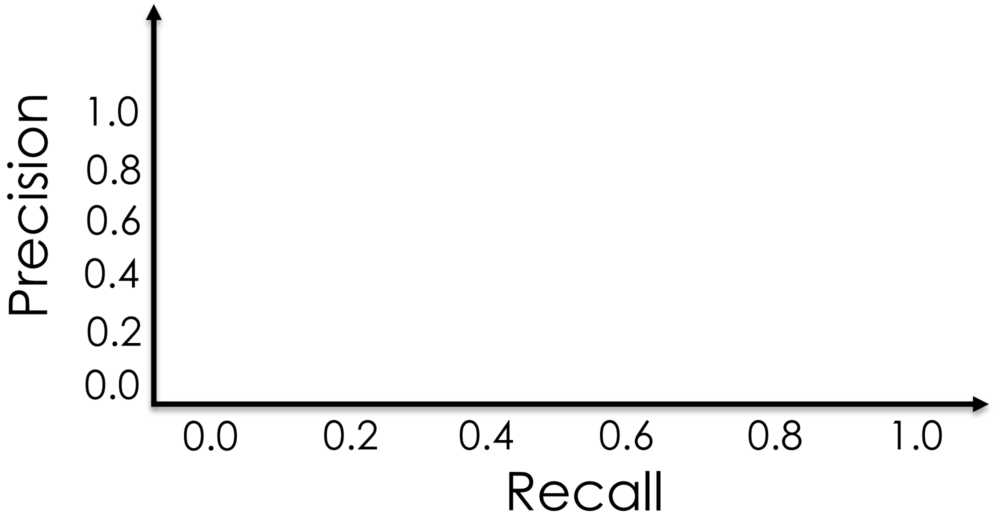

Are our models better than the baselines?

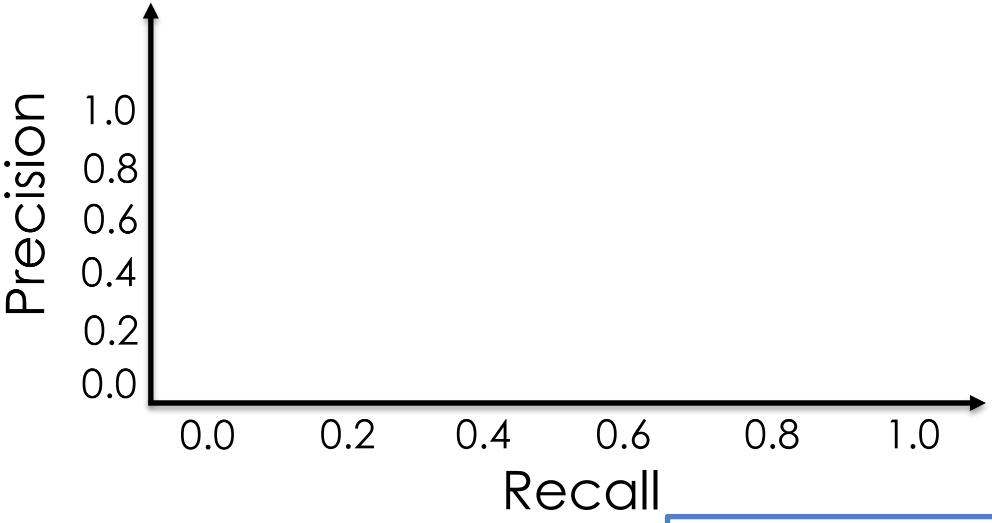

The fraction of posts the user actually wanted to stop sharing that were predicted

Are our models better than the baselines?

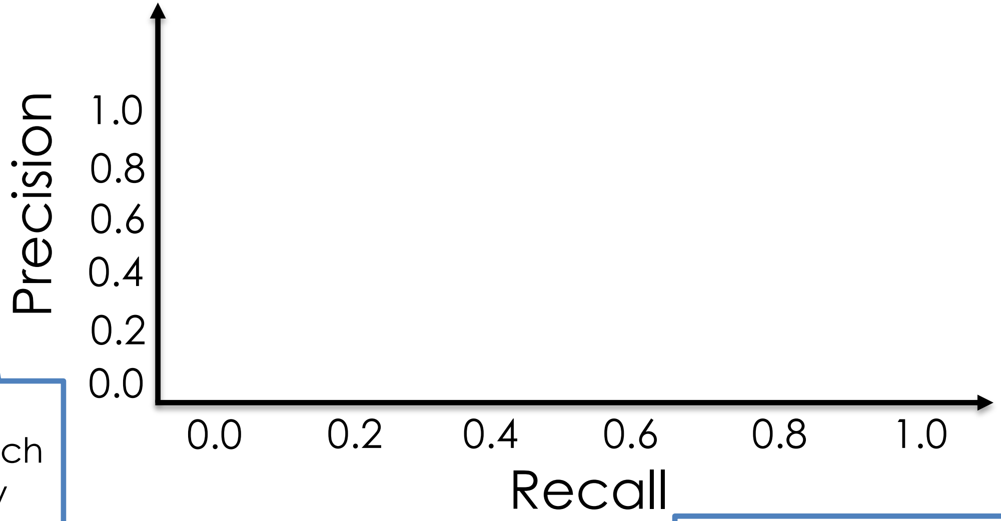

The fraction of predictions for which the user actually wanted to stop sharing

The fraction of posts the user actually wanted to stop sharing that were predicted

Are our models better than the baselines?

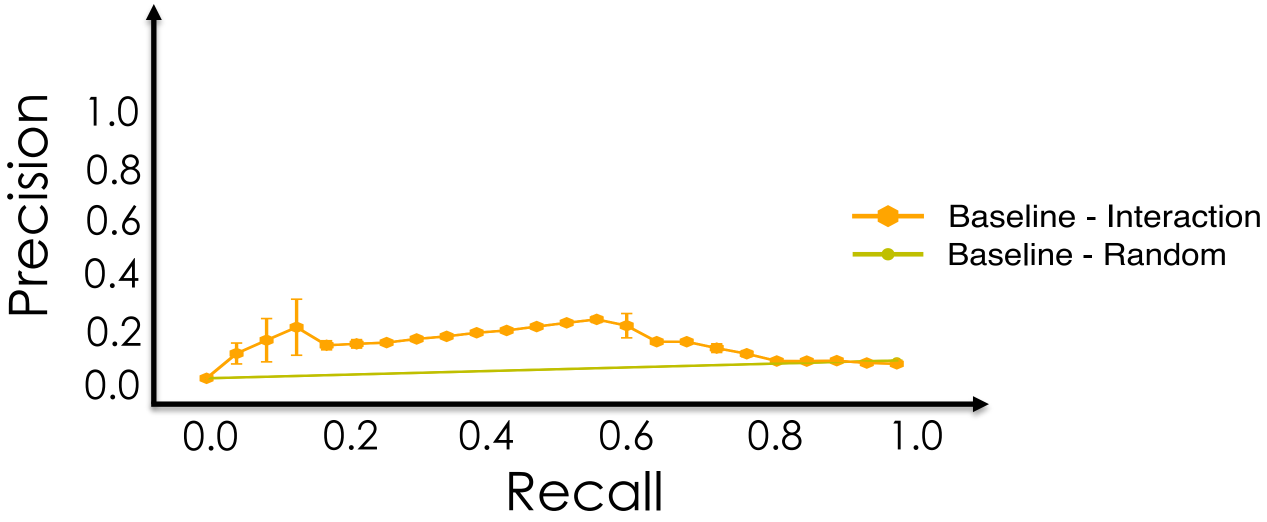

Are our models better than the baselines?

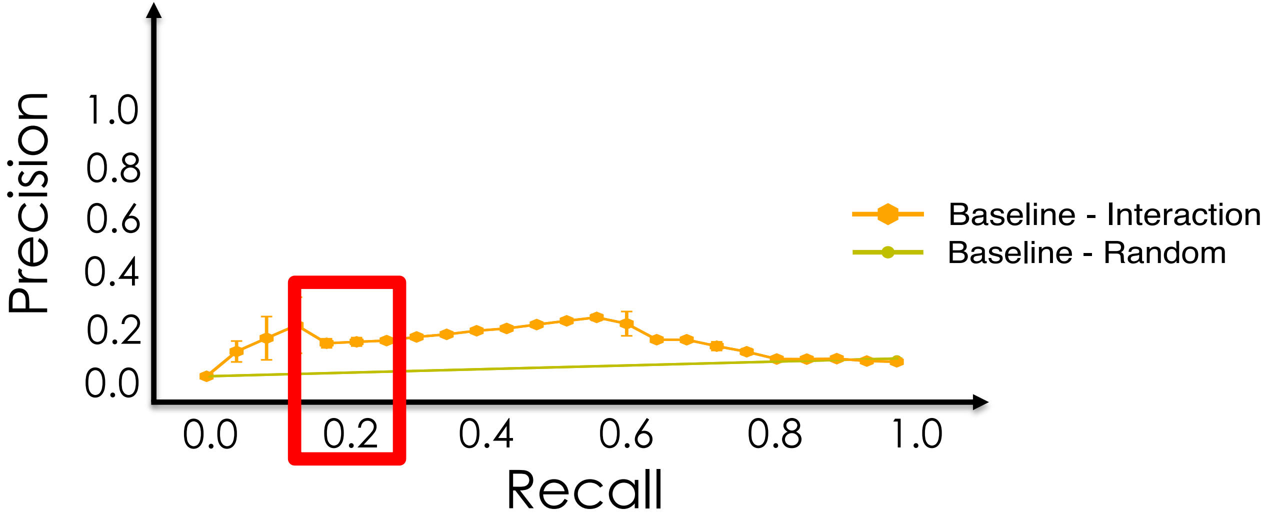

Are our models better than the baselines?

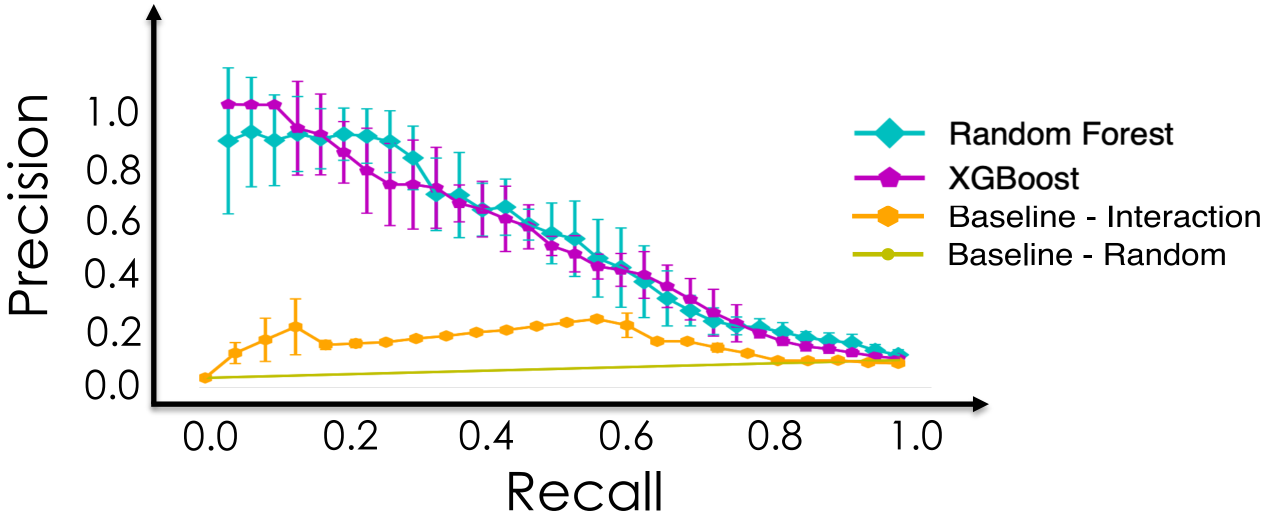

Are our models better than the baselines?

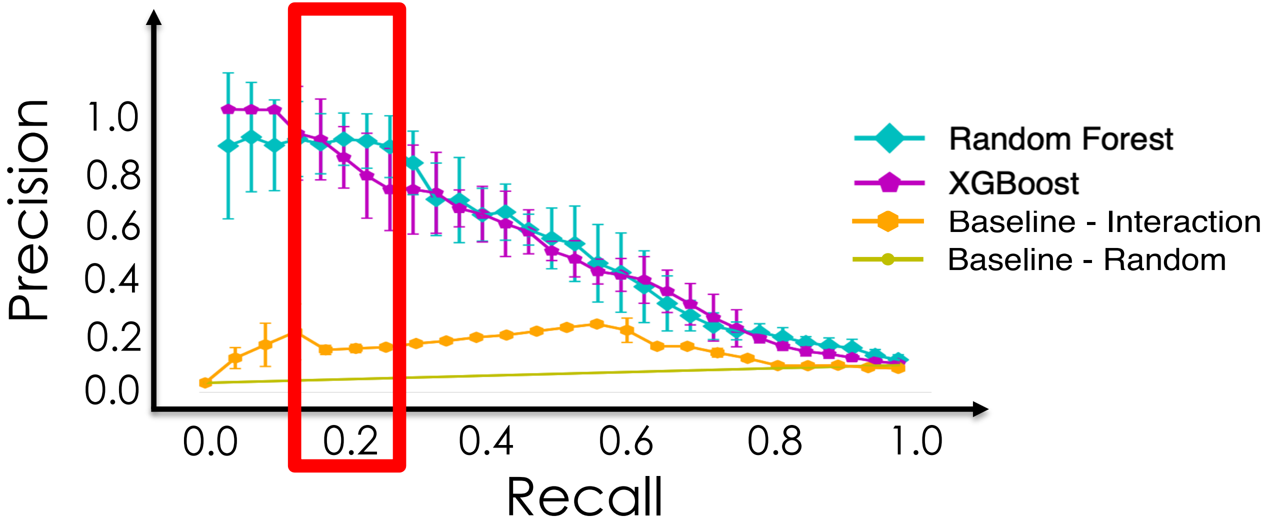

Are our models better than the baselines?

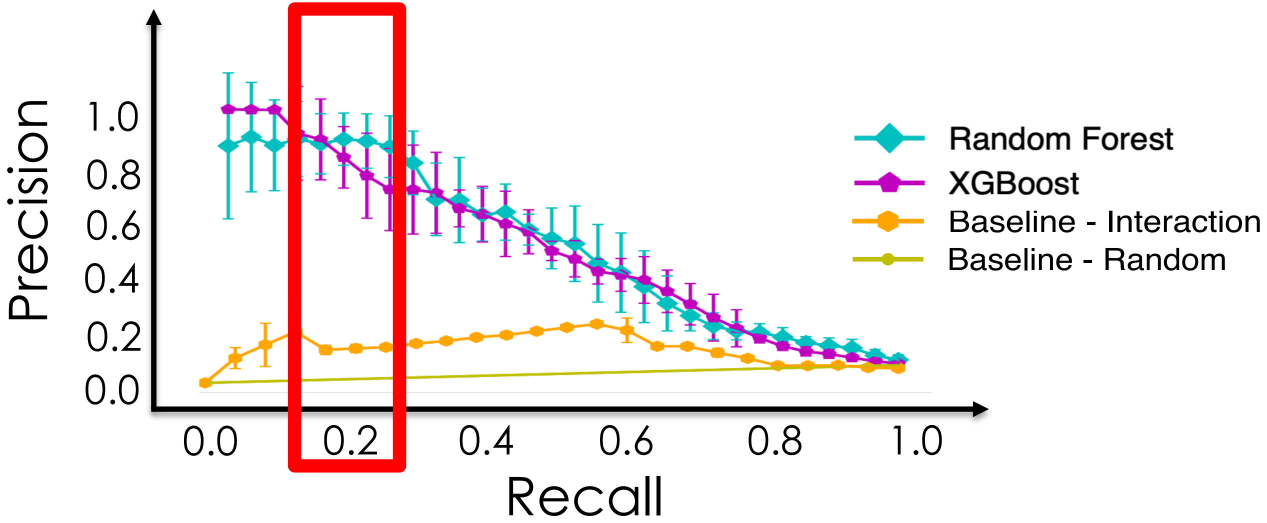

Substantial improvement over baselines

Prediction task

Prediction task

Predict if a user wants to "stop sharing" a given post with a given friend

Output

List of friend-post pairs ordered by probability

Ground truth

Privacy decisions for 78 participants x 5 posts X 6 friends = 2,340 pairs

Recommendation accuracy of our models

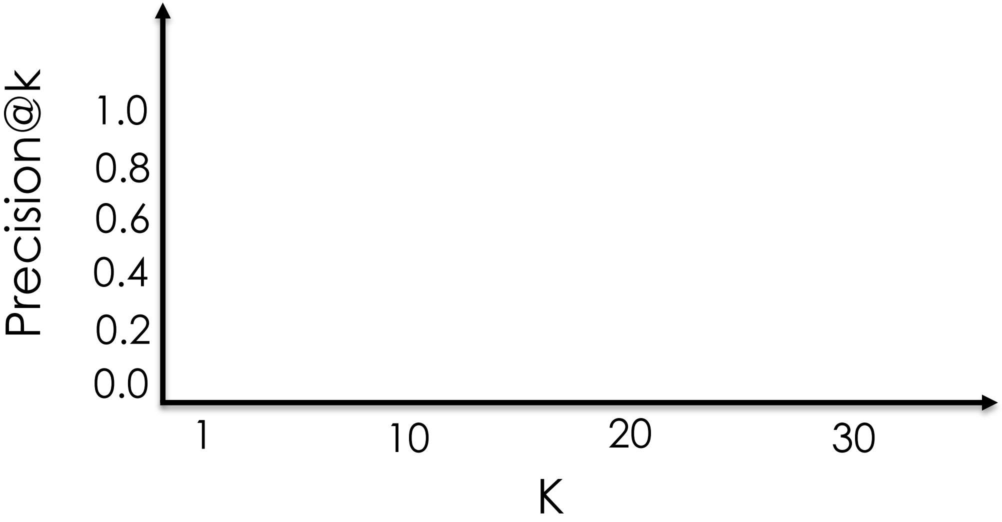

Recommendation accuracy of our models

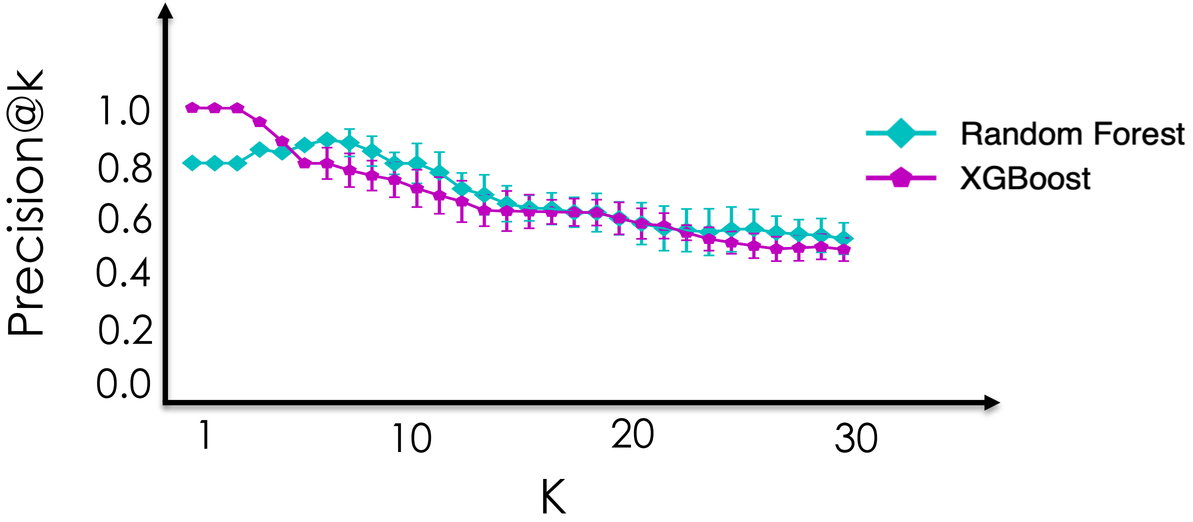

Recommendation accuracy of our models

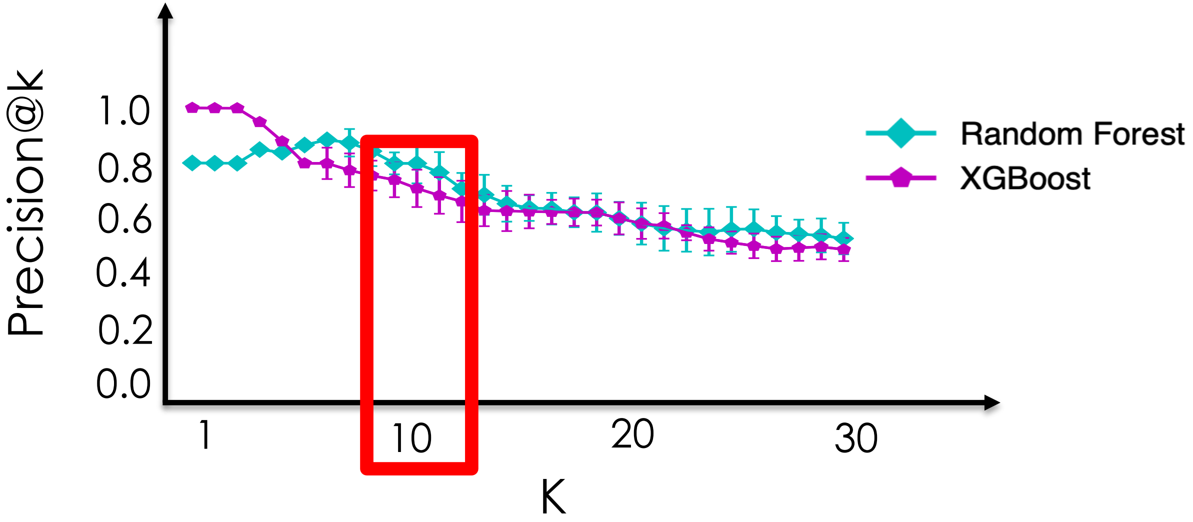

Recommendation accuracy of our models

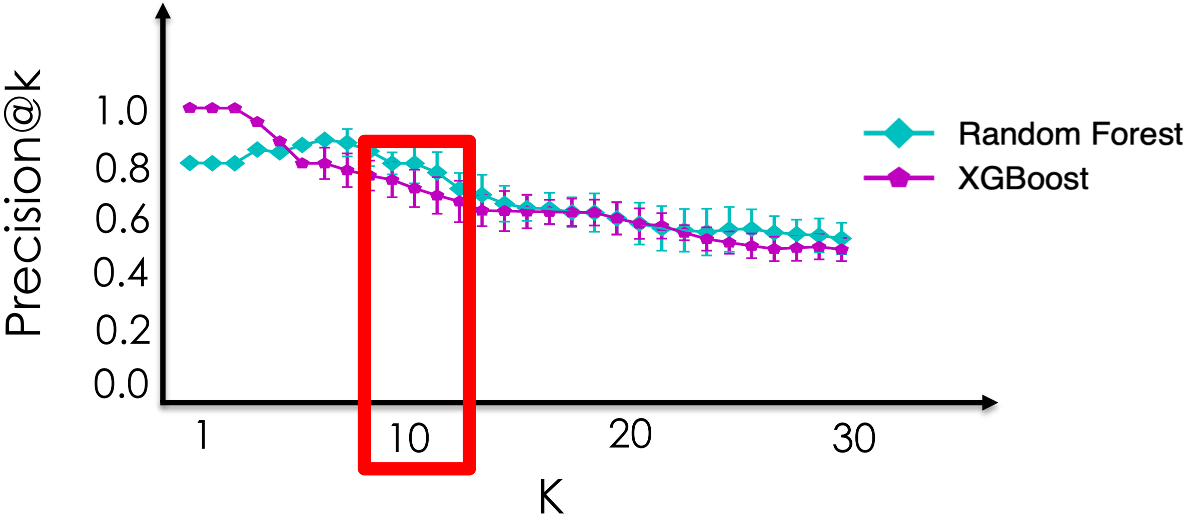

30 recommendations with good precision!

Understanding inaccurate predictions

Qualitative data from survey: “Why” did desired setting change?

Understanding inaccurate predictions

Qualitative data from survey: “Why” did desired setting change?

“I no longer participate in these activities and don’t find them appropriate any longer.”

Understanding inaccurate predictions

Qualitative data from survey: "Why" did desired setting change?

"I no longer participate in these activities and don't find them appropriate any longer."

"Because the people I feel close to has changed in the years since that post."

Understanding inaccurate predictions

Qualitative data from survey: “Why” did desired setting change?

“I no longer participate in these activities and don’t find them appropriate any longer.”

“Because the people I feel close to has changed in the years since that post.”

“ it shows a time that I was upset and i would rather not relive that.”

Understanding inaccurate predictions

Qualitative data from survey: “Why” did desired setting change?

“I no longer participate in these activities and don’t find them appropriate any longer.”

“Because the people I feel close to has changed in the years since that post.”

“ it shows a time that I was upset and i would rather not relive that.”

Coded this data to identify additional predictive features for future efforts

Future features to collect

Features of posts	Features from external content (image/video)
	Classes of sensitive information (e.g., children)
	Similarity of content with user’s current interest

Future features to collect

Features of posts	Features from external content (image/video)
	Classes of sensitive information (e.g., children)
	Similarity of content with user’s current interest
Features of friends	Interests, likes and dislikes of specific friends
	If particular friends are close family or related
	Frequency of offline interaction

Users change privacy preferences over time

2009

Content posted in freshman year: shared with everybody on internet

2012

3 years later: Hiring manager and colleagues **should not** see this

Temporal privacy management: control **who can see old content (e.g., via deletion)**

Temporal Privacy: Deleting content

Collecting data on privacy preference change

In this study we focus on Twitter

Simple privacy preferences

- Either publicly visible to everyone
- Or withdrawn from public domain (by deletion or making account private)

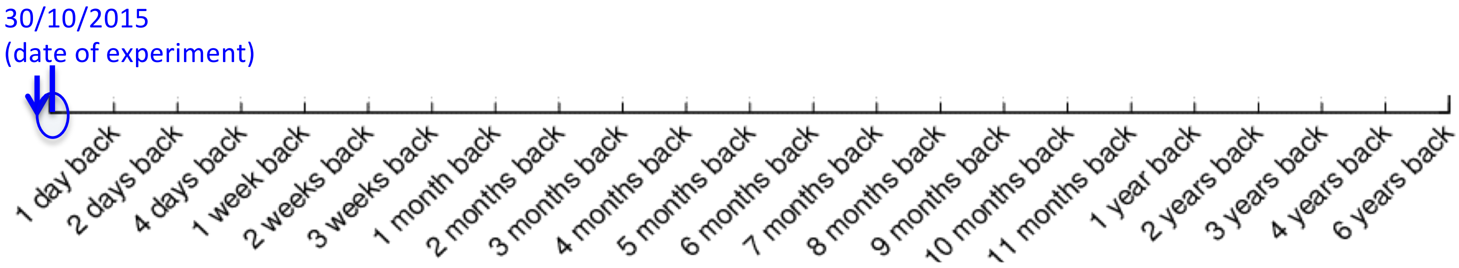

30/10/2015
(date of experiment)

Time in past when the tweets were posted (relative to the date of experiment**)**

All of these past tweets were **public when they were posted**

If **inaccessible** on experiment date, privacy preferences **changed** over time

Do users change privacy preferences over time?

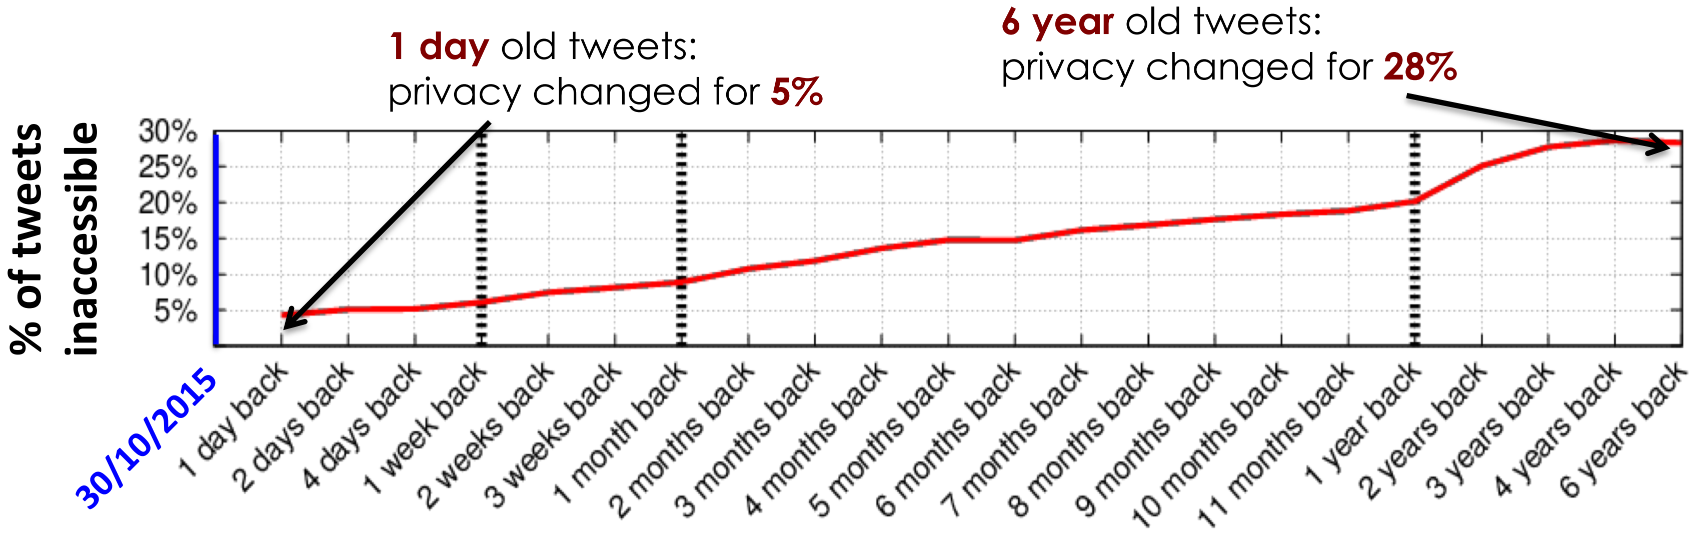

Users change privacy for increasing amount of old data with time

How do these users change privacy of this content?

Mechanisms to change privacy on Twitter

Three ways users change privacy of old content in Twitter

They are the temporal privacy control mechanisms

Mechanism	Description
Selective deletion	Selectively withdraw some old tweets to control exposure
Account deletion	Withdraw all old tweets to control exposure in bulk
Making account private	Withdraw all old tweets to control exposure in bulk

How do users change privacy preferences?

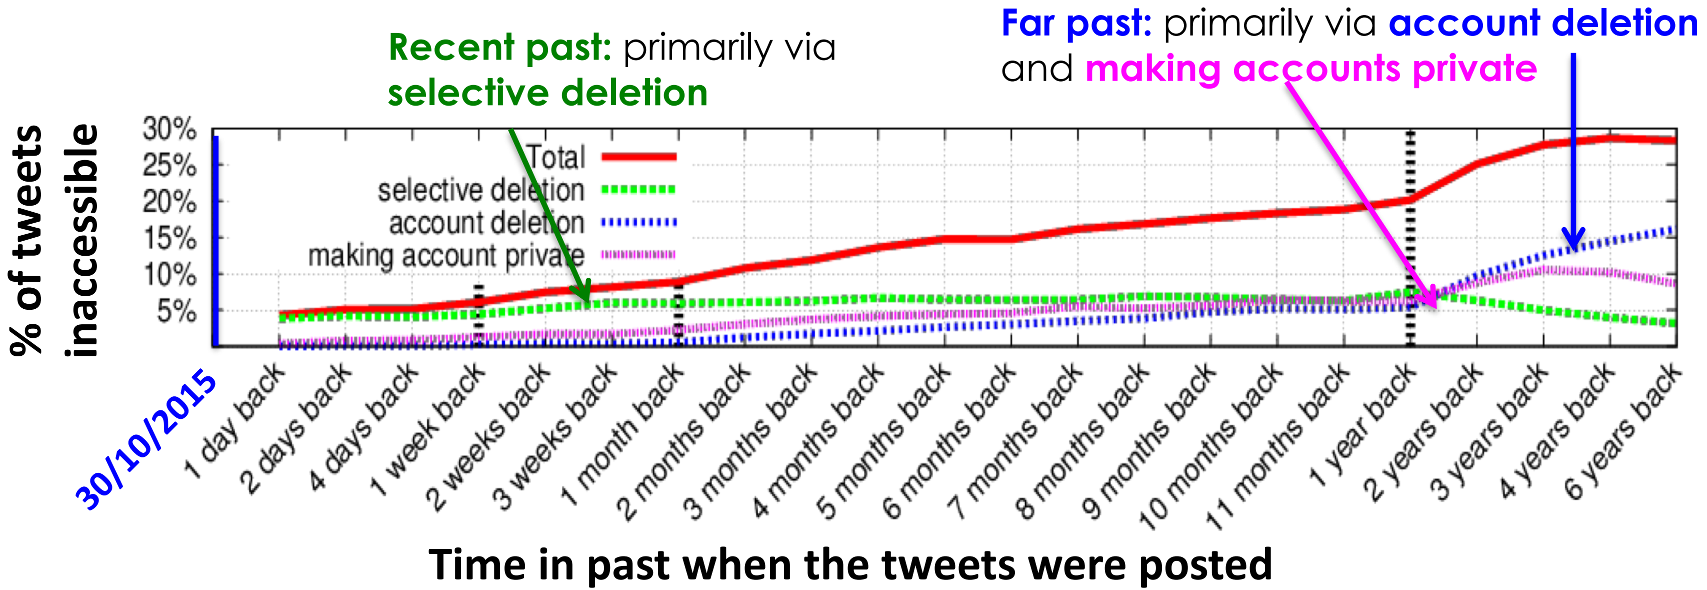

Changing privacy for content from **far past compared to recent past**

Very different mechanisms

Do many users change privacy of old content?

We randomly sample **100k** active users from 2009

Out of 8.9m random old tweets from these users 29.1% is inaccessible

What fraction of users change privacy of their content?

User type	% of all users
Selectively deleted tweets	8.3%
Deleted their account	15.9%
Made their account private	10.4%
Users who take actions that changes privacy of their content	**34.6%**

A significant fraction of users change privacy of their old content

However there is a problem …

Issue with content withdrawal

Posts from others (e.g., replies, tags) **leak information about withdrawn content**

We call them residual activities

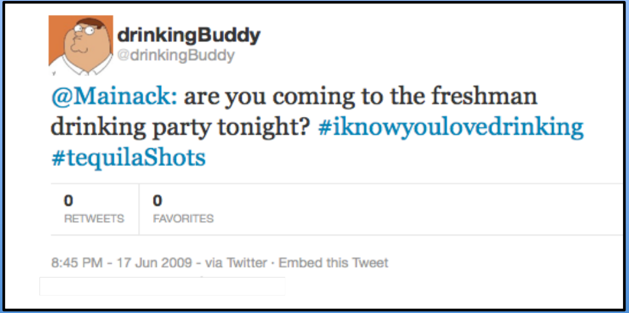

drinkingBuddy
@drinkingBuddy

@Mainack: are you coming to the freshman drinking party tonight? #iknowyoulovedrinking #tequilaShots

0 RETWEETS 0 FAVORITES

8:45 PM - 17 Jun 2009 - via Twitter · Embed this Tweet

Created an app to rai on leak

http://twitter-app.mpi

Need for temporal privacy: Summary

Twitter users indeed withdraw 28% of their 6 year old posts

Residual activities leak a lot of information about withdrawn content

Created a web application to raise user awareness about the information leak

Deletion Privacy

Courtaey for some slides: Mohsen Minaei

Enormous amount of social content is deleted

Long-term exposure of the shared data raises numerous longitudinal privacy concerns

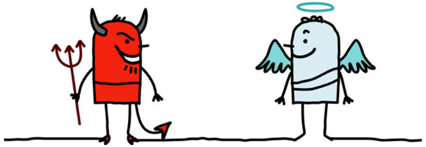

Deletions are common on social platforms

> 30% of posts are deleted within a 6 year period

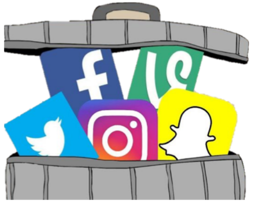

Do deletions hide the unwanted information?

Case 2: Fallait Pas Supprimer

“Should not Delete”

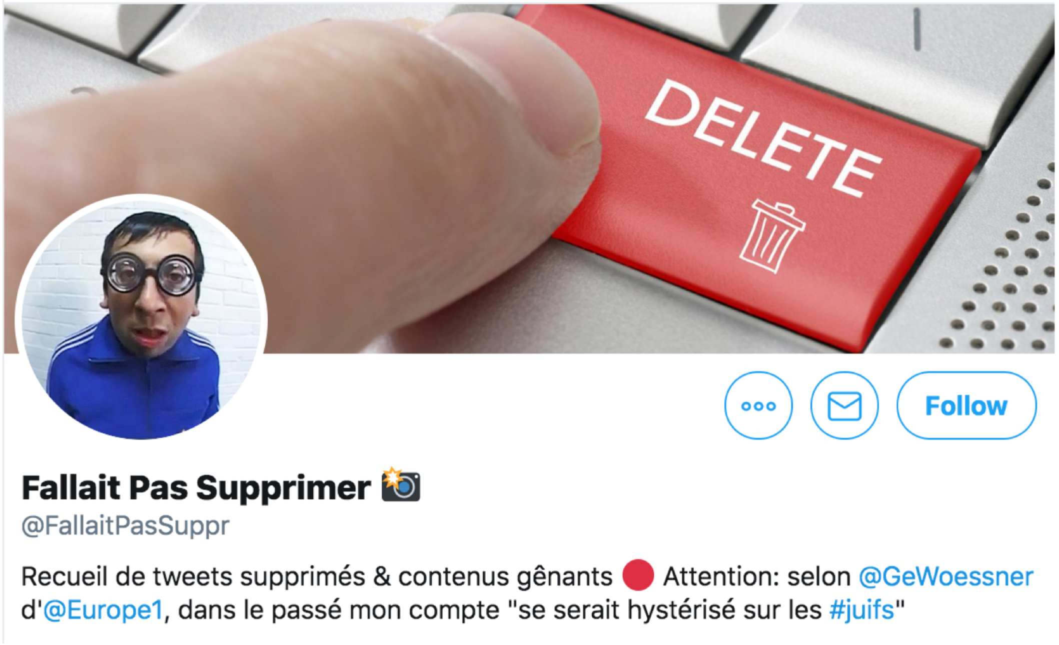

Fallait Pas Supprimer 📸

@FallaitPasSuppr

Recueil de tweets supprimés & contenus gênants 🔴 Attention: selon @GeWoessner d'@Europe1, dans le passé mon compte "se serait hystérisé sur les #juifs"

Case 2: Fallait Pas Supprimer

Deletion of normal daily users are noticed

92 Following 919 Followers

Fallait Pas Supprimer 📸

@FallaitPasSuppr

Recueil de tweets supprimés & contenus gênants 🔴 Attention: selon @GeWoessner d'@Europe1, dans le passé mon compte "se serait hystérisé sur les #juifs"

Web Services Hoard Deleted Content

Removeddit

Uneddit

StackPrinter-Deleted

YouTomb

Politwoops

reddit

stackoverflow

YouTube

twitter

Lethe: Intuition

A simple but drastic solution:

Hide and resurrect the non-deleted posts!!

Confuse the adversary: is a post hidden or deleted?

A trade-off between Privacy and Availability

Twitter example

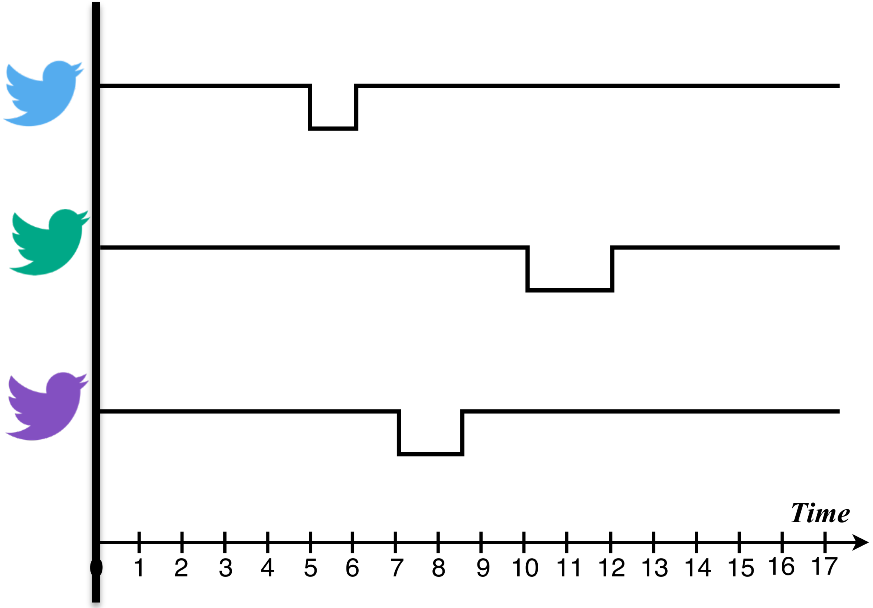

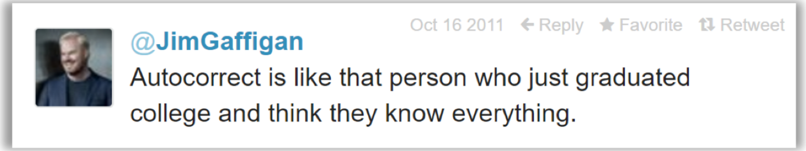
Oct 16 2011 ← Reply ★ Favorite ⇅ Retweet

@JimGaffigan

Autocorrect is like that person who just graduated college and think they know everything.

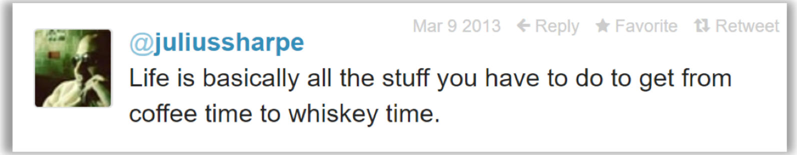
Mar 9 2013 ← Reply ★ Favorite ⇅ Retweet

@juliussharpe

Life is basically all the stuff you have to do to get from coffee time to whiskey time.

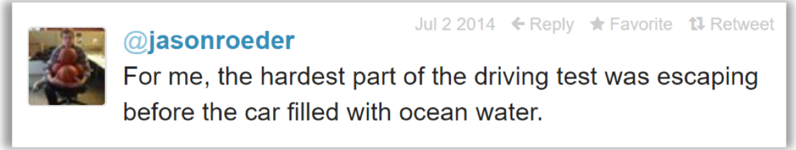
Jul 2 2014 ← Reply ★ Favorite ⇅ Retweet

@jasonroeder

For me, the hardest part of the driving test was escaping before the car filled with ocean water.

Key idea of the design

Intermittent withdrawal mechanism

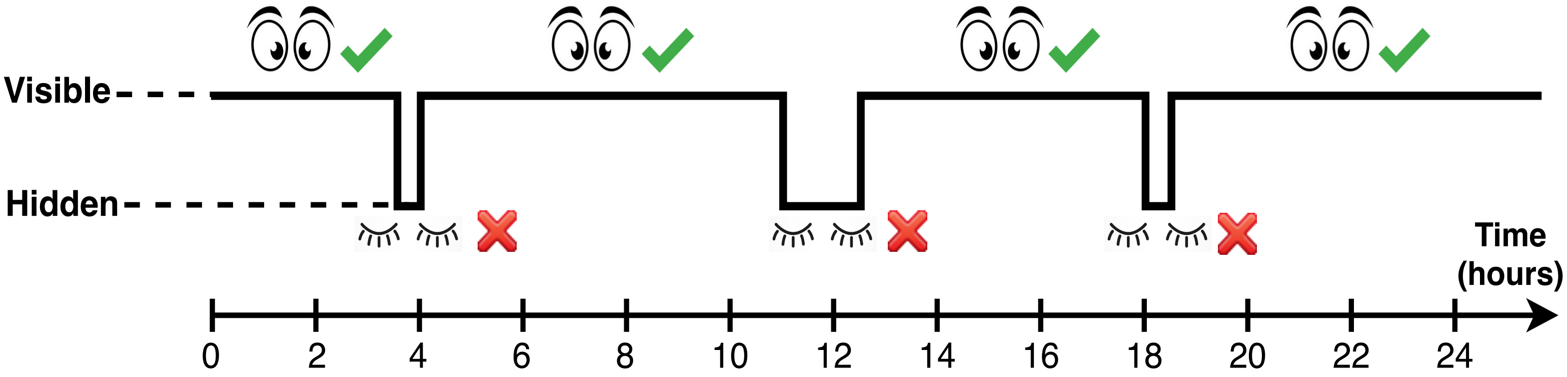

Example of a non-deleted post for a day with 90% availability

Threat Model

Persistently observes the platform and takes snapshot of it at different times

Act as normal users

Large-scale analysis of data

System & Security Goals

Deletion Privacy

Adversarial overhead

Availability

Deletion privacy: Our definition

Uncertainty about a post being deleted or just temporarily withdrawn at a given point of time

$$LR = \frac{Pr(\text{observation of all periods} \mid \text{post deleted at time } t_c)}{Pr(\text{observation of all periods} \mid \text{post not deleted at time } t_c)}$$

Observed states

$$LR = \frac{Pr(\mathcal{O}(\Delta t_u, \Delta t_d) \mid \mathcal{R}(t_c) = 0)}{Pr(\mathcal{O}(\Delta t_u, \Delta t_d) \mid \mathcal{R}(t_c) = 1)}.$$

Real State of the post

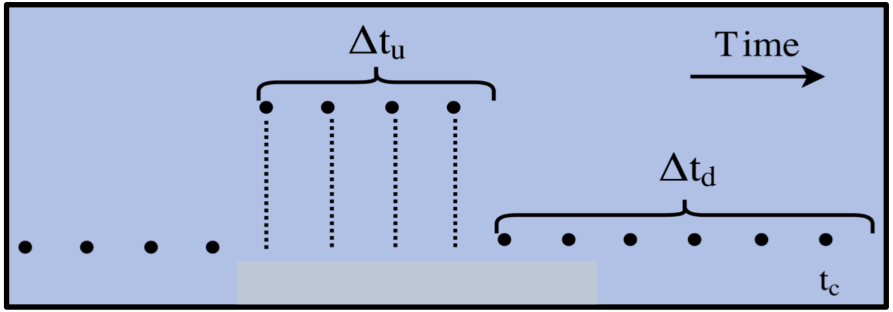

Likelihood ratio (LR)

Analyzing the LR

$$LR = \left(\frac{\overline{F_{T_u}}(\Delta t_u)}{f_{T_u}(\Delta t_u)} + 1 \right) \cdot \frac{1}{\overline{F_{T_d}}(\Delta t_d - 1)}$$

LR is dependent on the PMF and CCDF of the up distribution as well as the CCDF of the down distribution

Quantifying the success of adversary

Adversarial overhead: precision and recall

$$Precision = \frac{TP}{TP+FP}$$

$$Recall = \frac{TP}{TP+FN}$$

TP: correctly detected deleted posts	FP: falsely detected non-deleted posts
FN: falsely not detected deleted posts	TN: correctly not detected non-deleted posts

Platform Availability: avg. availability of a post within a period

Choice of the up/down distributions

$$LR = \left(\frac{\overline{F_{T_u}}(\Delta t_u)}{f_{T_u}(\Delta t_u)} + 1\right) \cdot \frac{1}{\overline{F_{T_d}}(\Delta t_d - 1)}$$

Up Distribution: memoryless Geometric distribution

It has a constant inverse hazard rate for all up time periods

Down Distribution: heavy tailed Negative binomial distribution

lowest inverse CCDF value via empirical exploration

Deletion Privacy = Adversary Decision Threshold

$$LR = \left(\cancelto{c}{\frac{\overline{F_{T_u}}(\Delta t_u)}{f_{T_u}(\Delta t_u)}} + 1 \right) \cdot \frac{1}{\overline{F_{T_d}}(\Delta t_d - 1)}$$

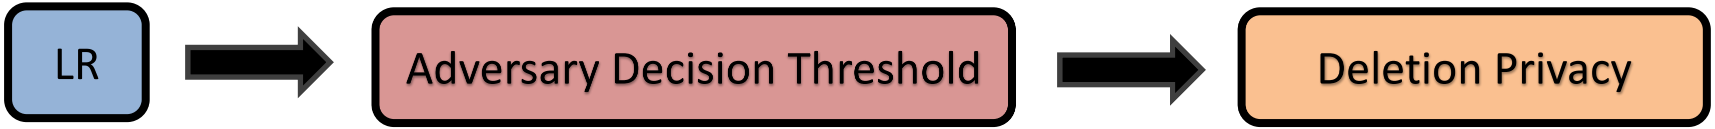
LR

Adversary Decision Threshold

Deletion Privacy

System Evaluation

What is the adversarial overhead for identifying deleted posts with Lethe?

Experiment set up

Dataset:

- 1% random sample of daily tweets (Oct 15 – Mar 17)
- 100 million tweets deleted from the one billion collection

Parameters

- Mean down time: 1 hour
- Mean up time: 6, 9, 19, hours
- availability 85, 90, 95%

Adversarial overhead with increasing precision

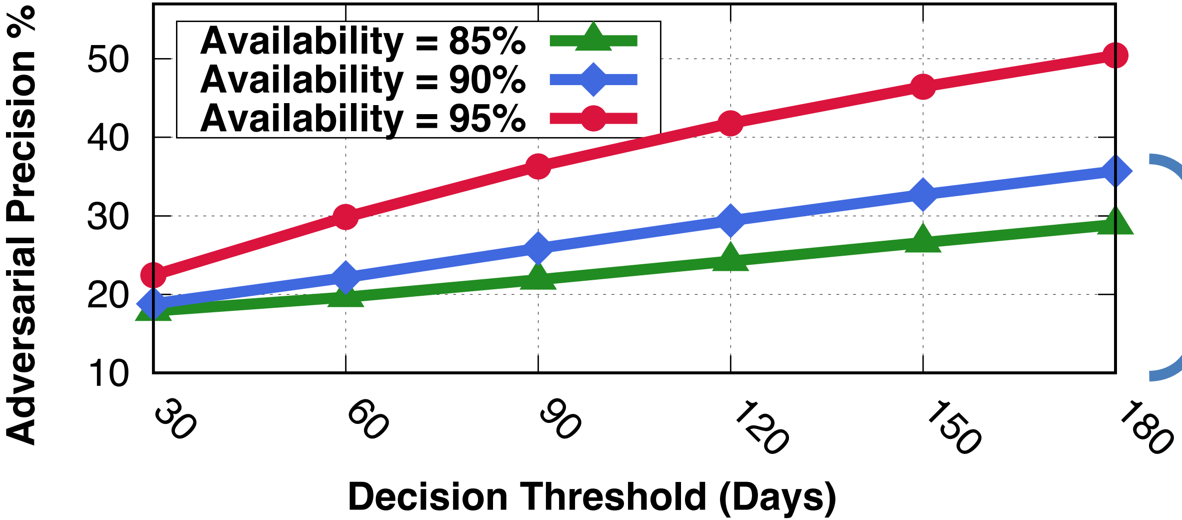

Adversarial precision less than 40%

Adversary has a low precision in identifying deleted content for different thresholds for all values of platform availability