



Moving Beyond Set-It-And-Forget-It Privacy Settings on Social Media

Mainack Mondal^{†*}

Mohammad Taha Khan[§]

Chris Kanisch[§]

[†]University of Chicago

Günce Su Yilmaz[†]

Michael Tang[†]

Blase Ur[†]

[§] University of Illinois at Chicago

Noah Hirsch[†]

Christopher Tran[§]

Elena Zheleva[§]

^{*}IIT Kharagpur

Need to revisit old privacy settings

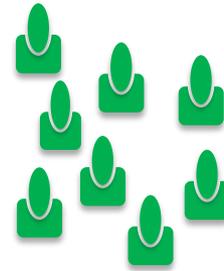
2009



Privacy setting: **“all friends”**

Need to revisit old privacy settings

2009



Undergraduate friends

Privacy setting: **“all friends”**

Need to revisit old privacy settings

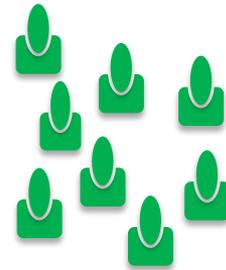
2019



Privacy setting: **“all friends”** !!

Need to revisit old privacy settings

2019

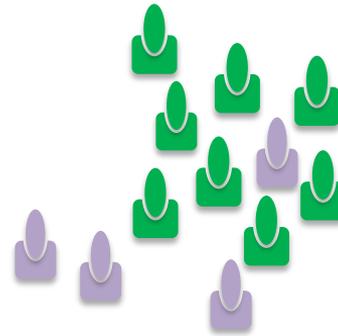


Undergraduate friends

Privacy setting: **“all friends”** !!

Need to revisit old privacy settings

2019



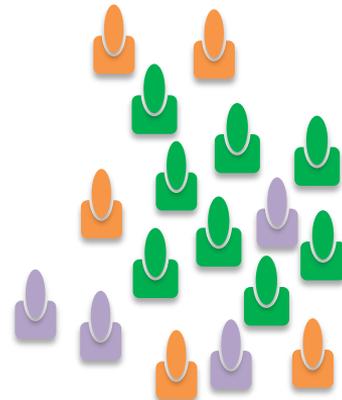
Undergraduate friends

Graduate school friends

Privacy setting: **“all friends”** !!

Need to revisit old privacy settings

2019



Undergraduate friends

Graduate school friends

Professional colleagues

Privacy setting: **“all friends”** !!

Need to revisit old privacy settings

2019



Privacy setting: **“all friends”** !!

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

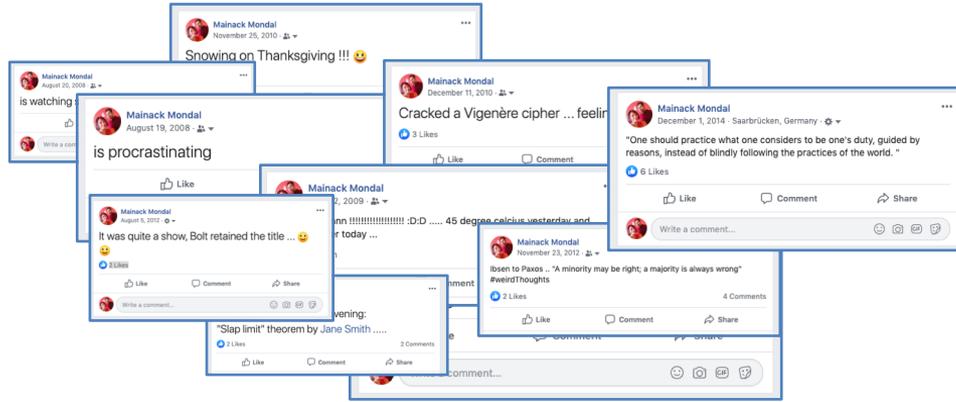
Need: Retrospectively manage privacy

Retrospective privacy management is difficult



Undergraduate friends

Retrospective privacy management is difficult



Students



Undergraduate friends

Graduate school friends

Professional colleagues



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](#)

Limit Past 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?

No Thanks

Let's Do It!

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](#)

Limit Past 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

Outline

Our data collection approach

Privacy settings and friend network over time

Preferences for changing privacy settings

Automated classifiers

Outline

Our data collection approach

Privacy settings and friend network over time

Preferences for changing privacy settings

Automated classifiers

Study overview

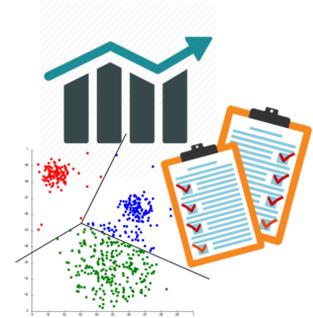


Privacy-preserving
data-collection
Infrastructure



amazon
mechanical turk

78 Facebook users



Two surveys

Generic survey



Overall Facebook usage over time

- —
- —
- —



Use of Facebook's privacy features

- —
- —

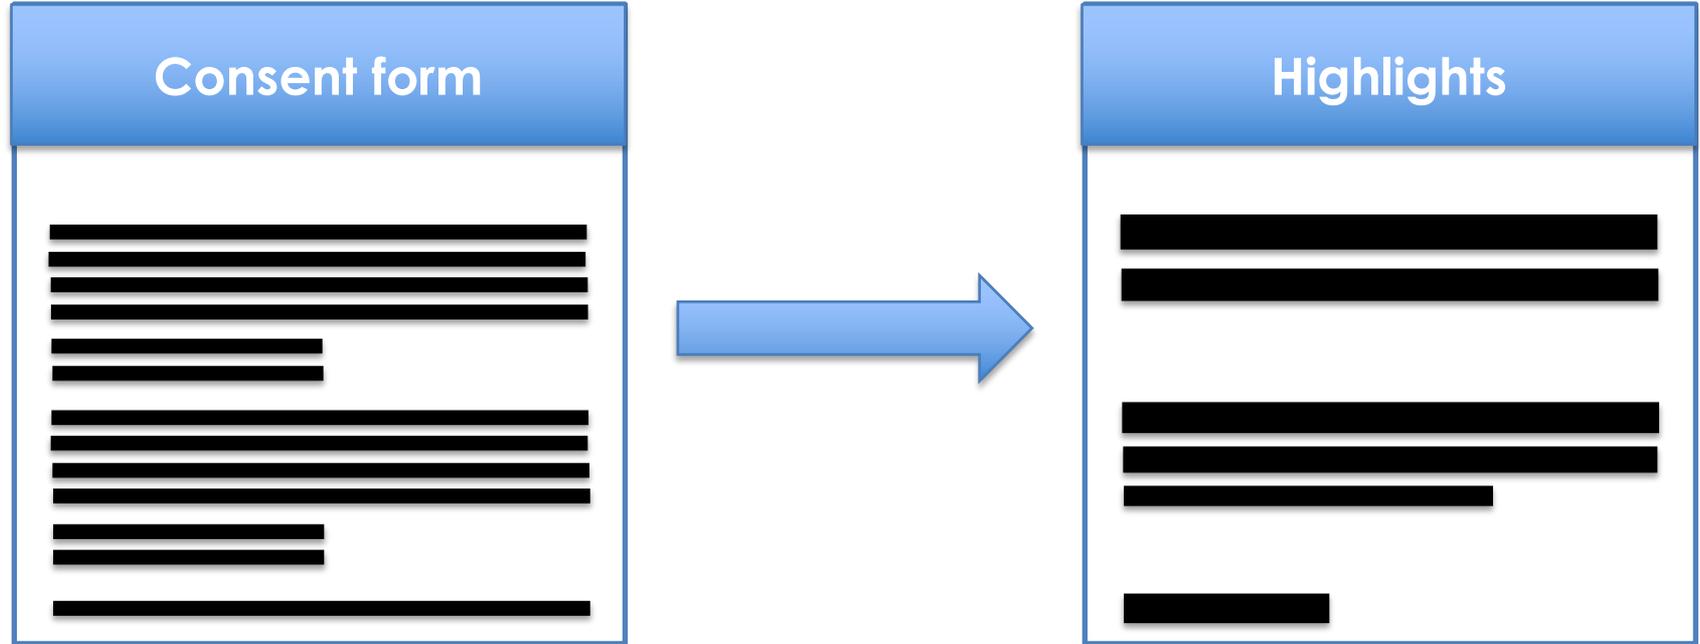


Participant demographics

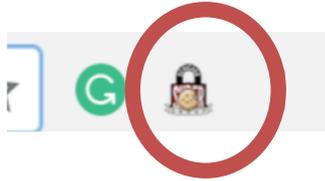
- —



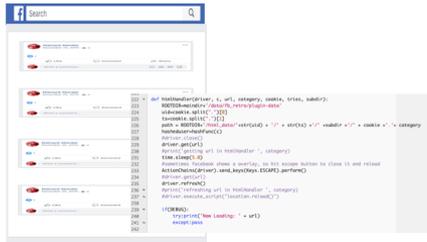
Consent process



Data collection process



Data collection process



Programmatic
No humans ever view raw HTML



Hash names and IDs;
No images collected



Never access
friends' profiles

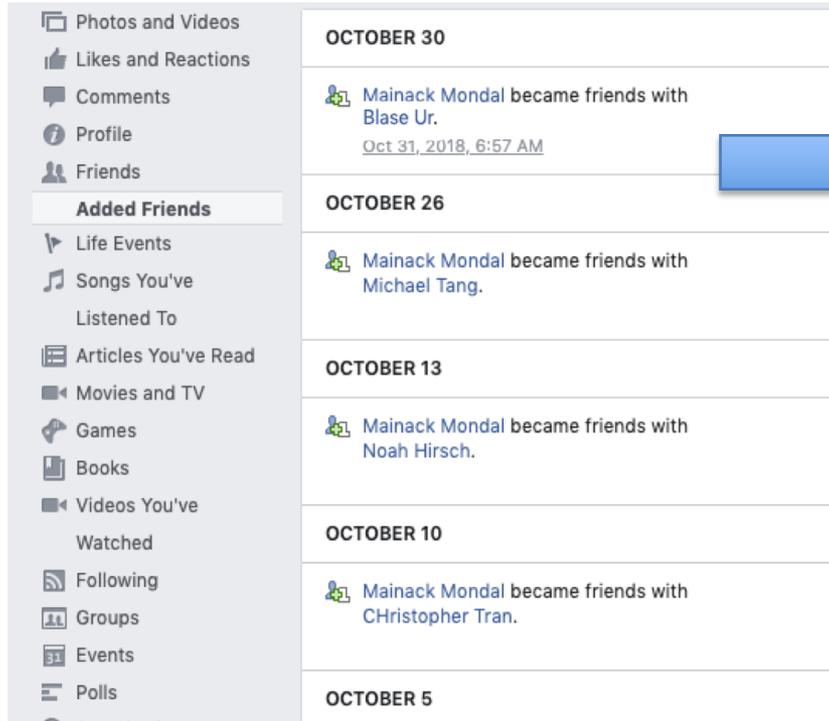
Facebook Timeline data



```
1 {
2   "user": "23765ae45...",
3   "timestamp": "May 26, 2009",
4   "privacy": "friends",
5   "numLikes": 5,
6   "numComments": 18,
7   "Text": "When we were young",
8   "post_url": https://facebook.com/...,
9   "commentObjects" : [
10    {
11     "user": "877326d4f...",
12     "text": "You look funny",
13     "timestamp": "...",
14     ...
15    }
16  ]
17  .....
18 }
```

Chose not to store images!

Facebook Activity Log data



The screenshot shows the Facebook Activity Log interface. On the left is a navigation menu with categories like Photos and Videos, Likes and Reactions, Comments, Profile, Friends, and Added Friends. The main content area displays a list of activities, with a blue arrow pointing from the 'Added Friends' section to the JSON data on the right. The activities listed are:

- 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**: (Activity partially visible)

```
1 {
2   "activityType": "addfriends",
3   "user": "23765ae45...",
4   "friend": "3264325ef...",
5   "timestamp": "Oct 31, 2018"
6 }
7 {
8   "activityType": "addfriends",
9   "user": "23765ae45...",
10  "friend": "85e47873...",
11  "timestamp": "Oct 26, 2018"
12 }
13 ...
```

**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

Outline

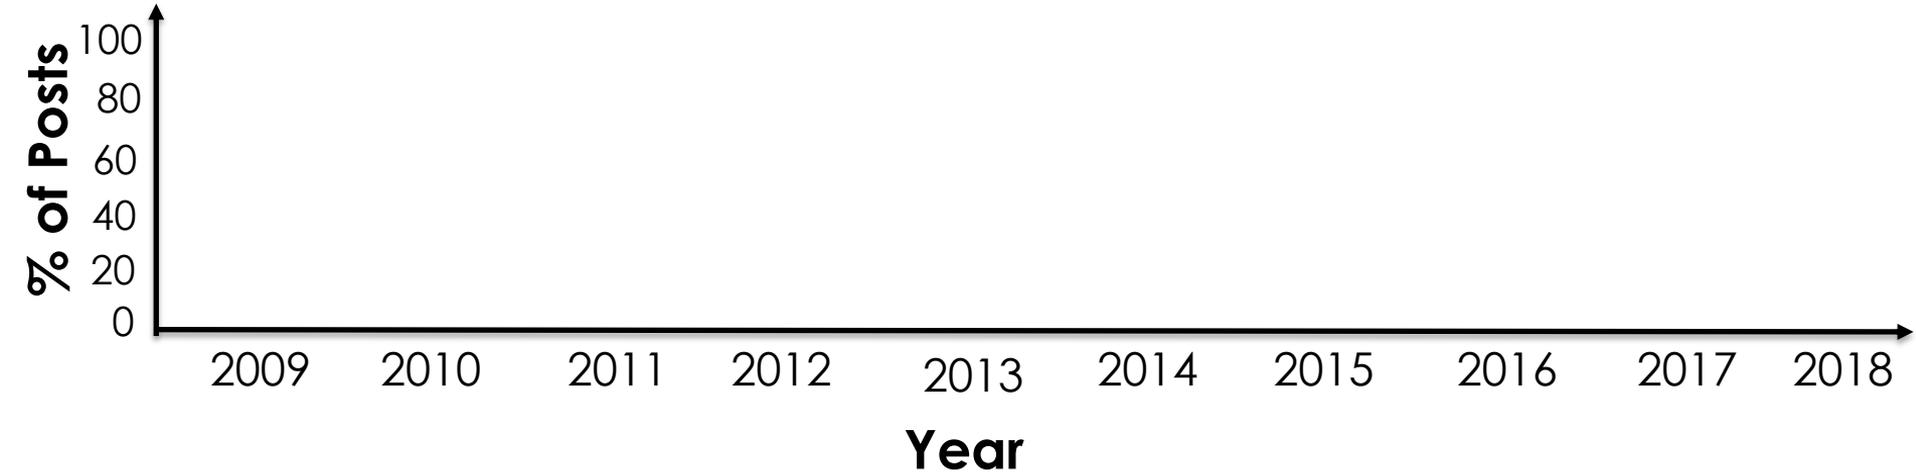
Our data collection approach

Privacy settings and friend network over time

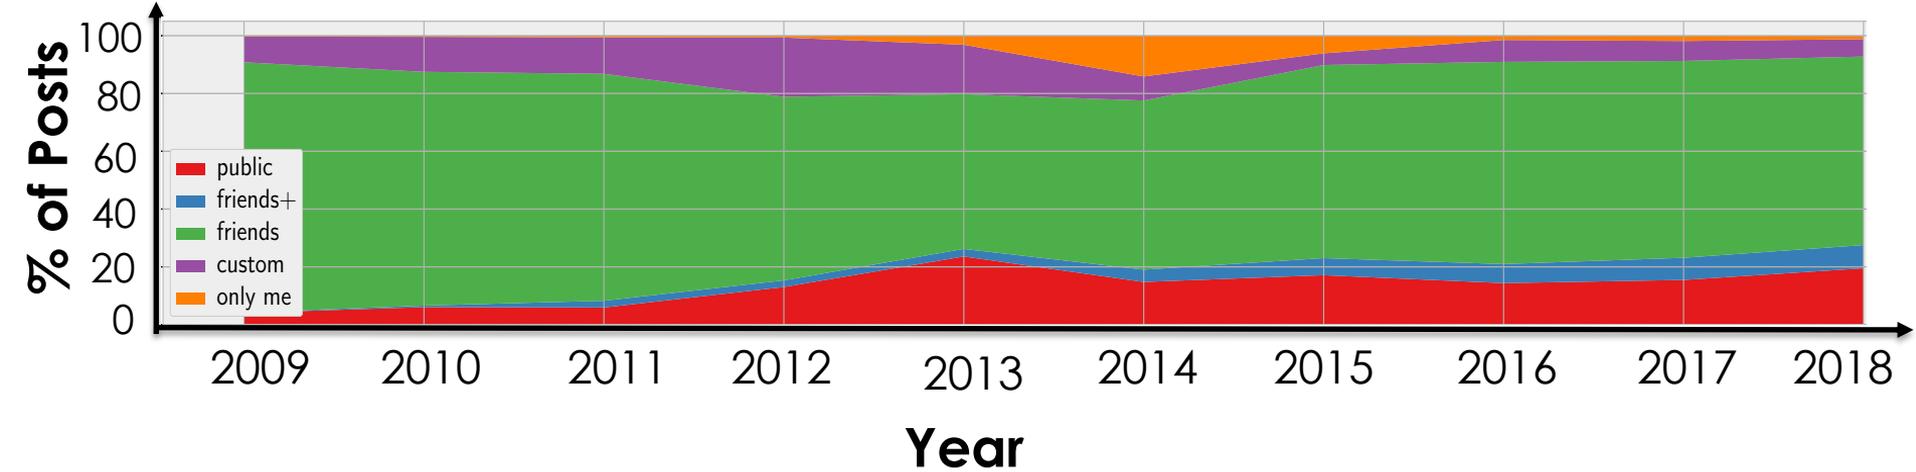
Preferences for changing privacy settings

Automated classifiers

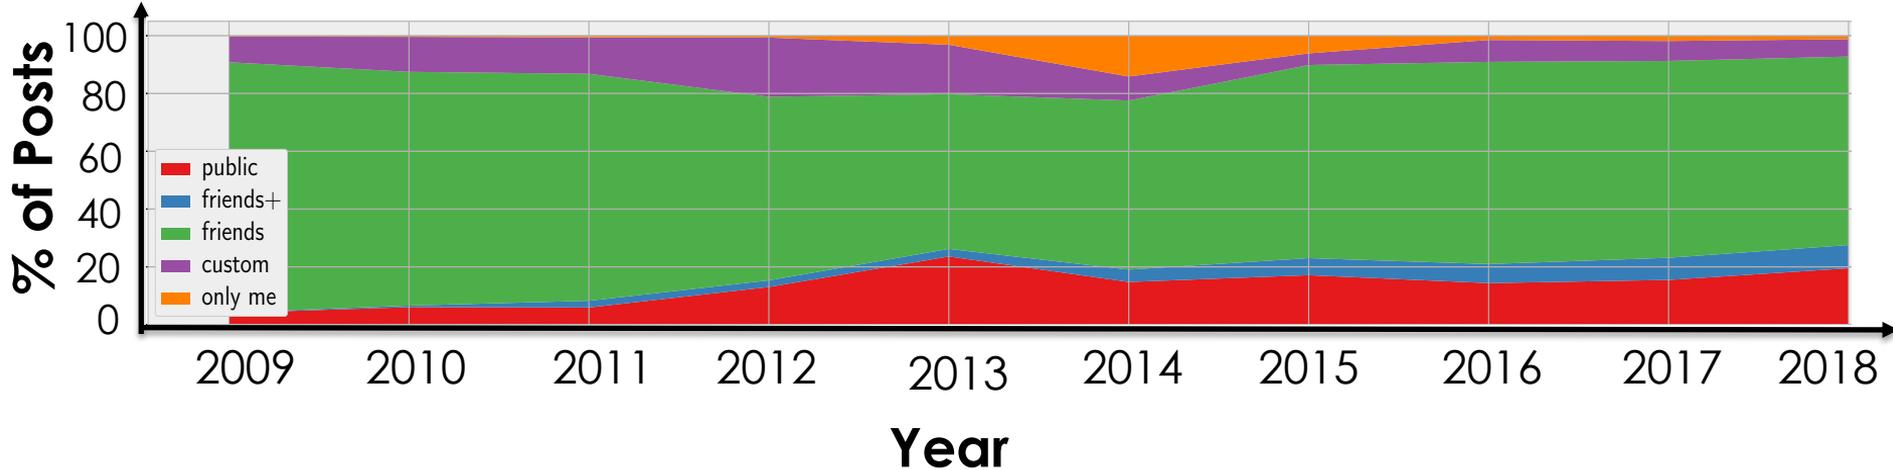
Privacy settings over time



Privacy settings over time

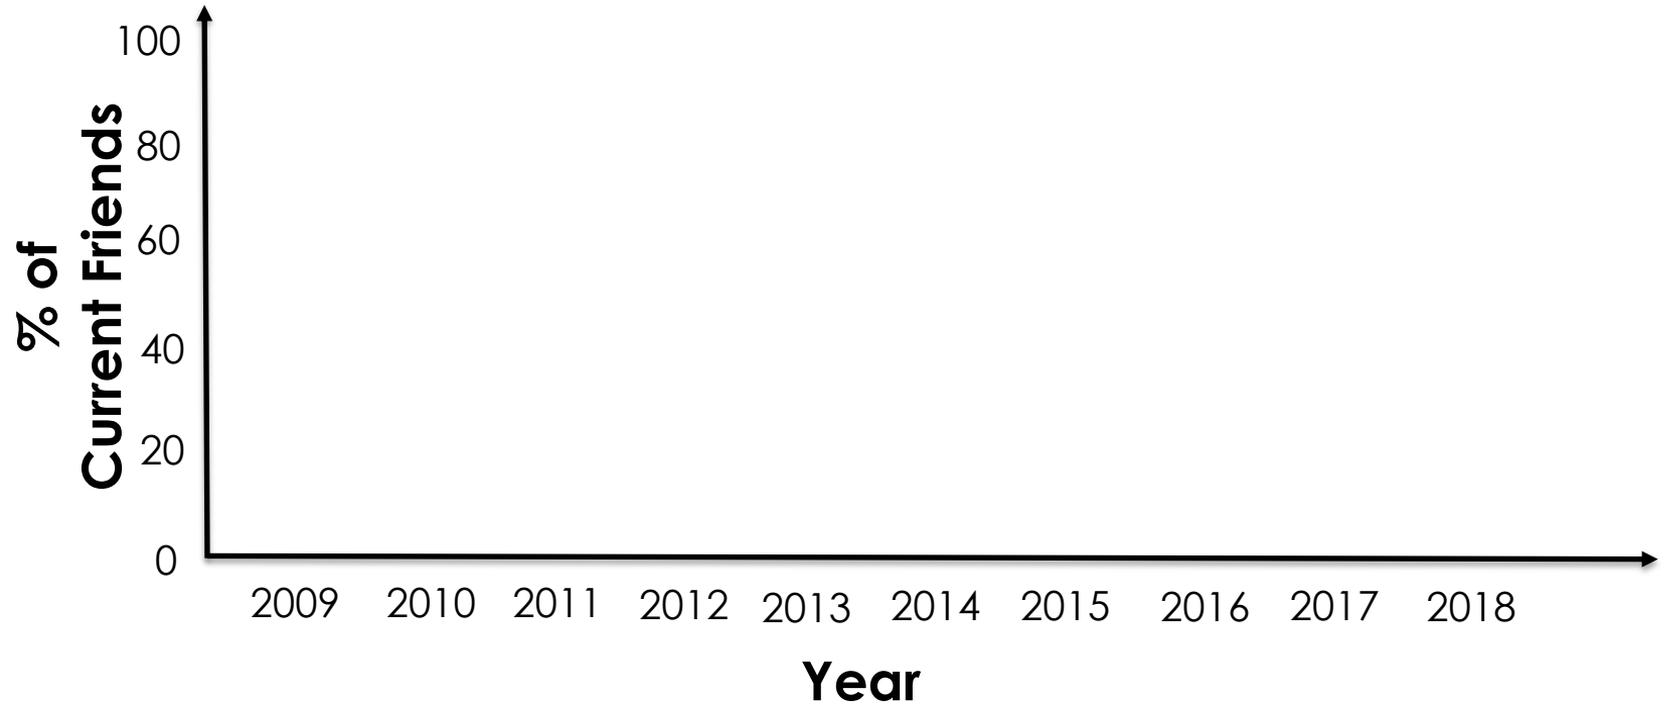


Privacy settings over time

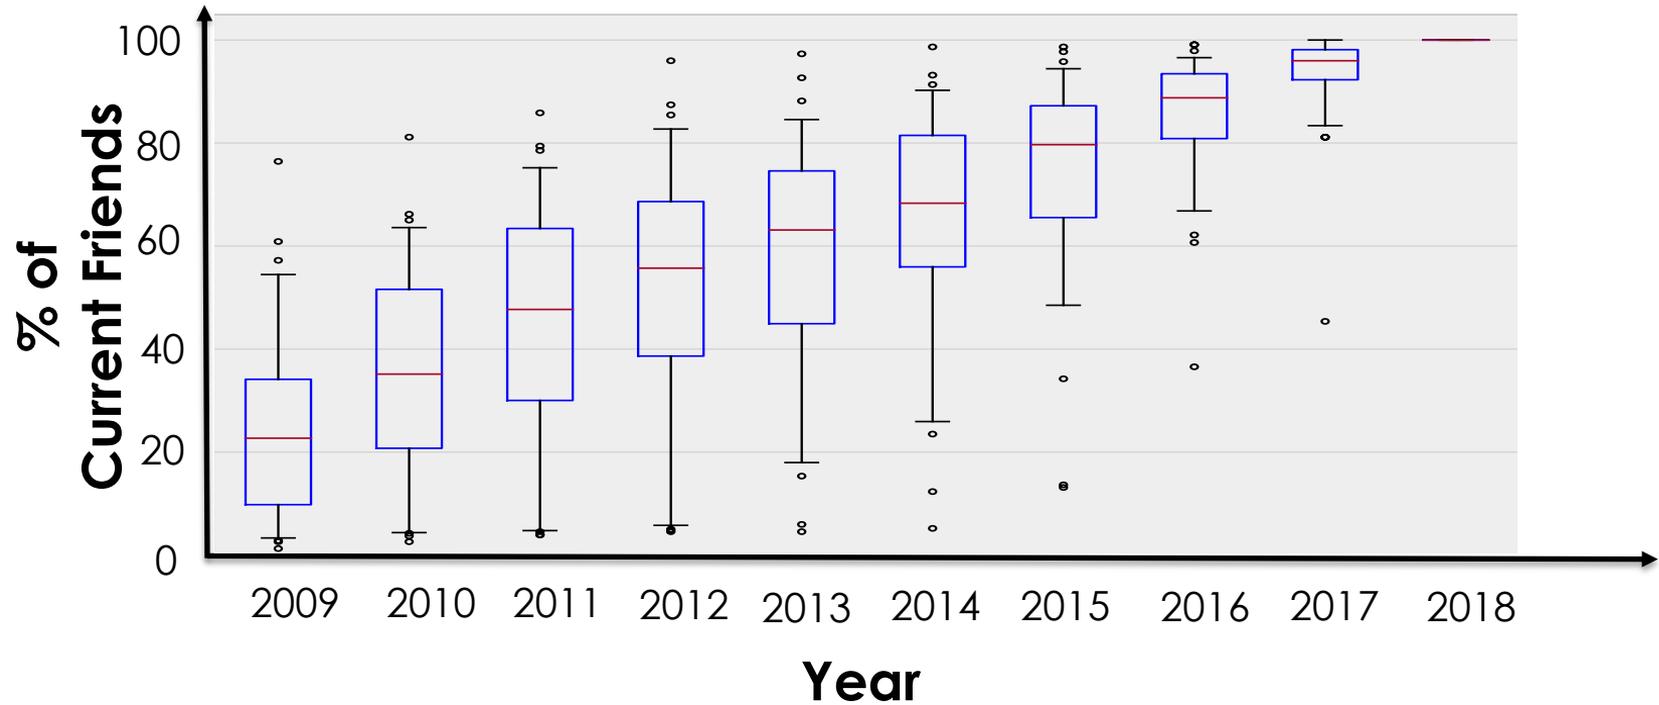


Majority of old posts are shared with all “friends”

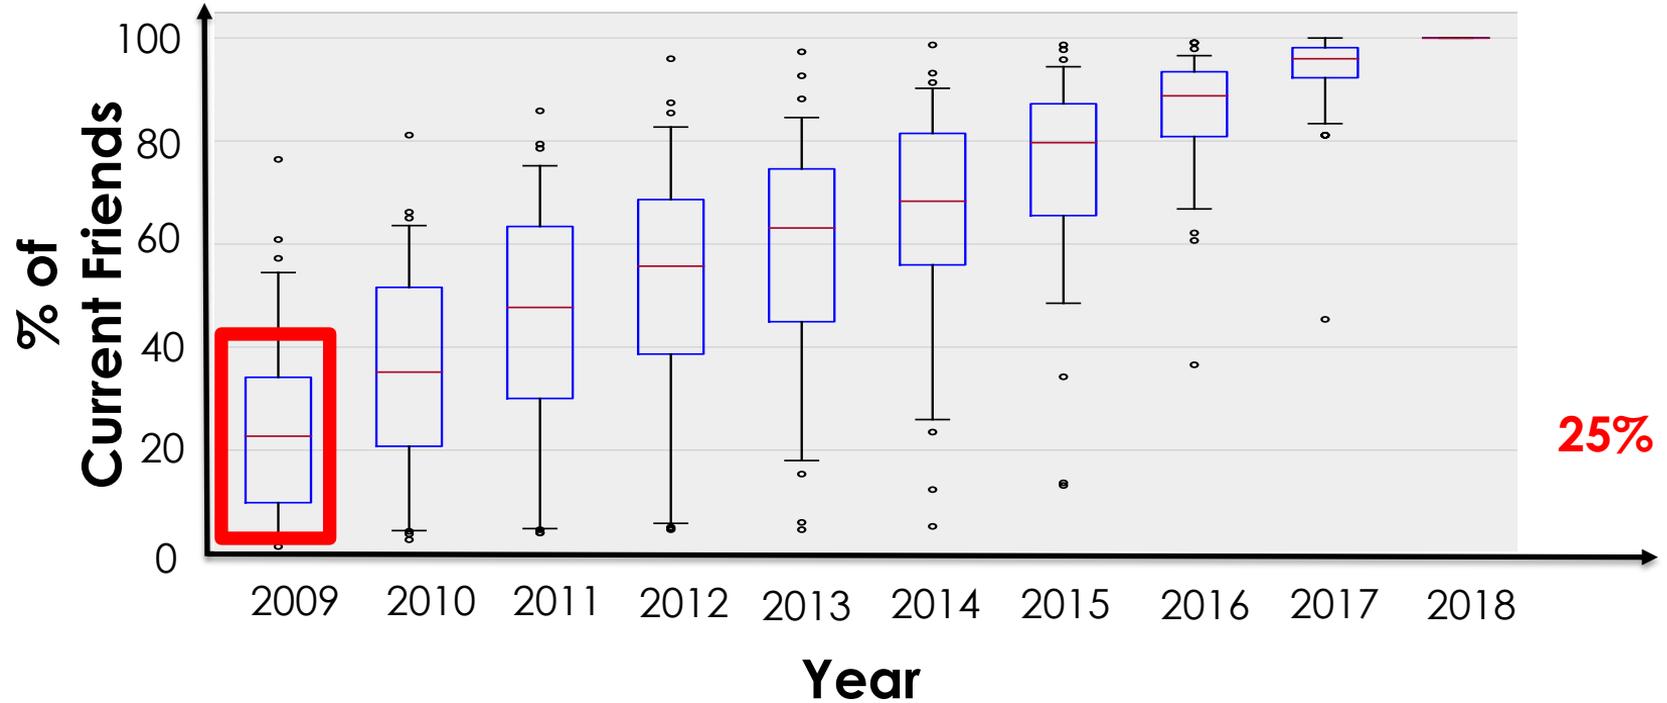
Change in number of friends



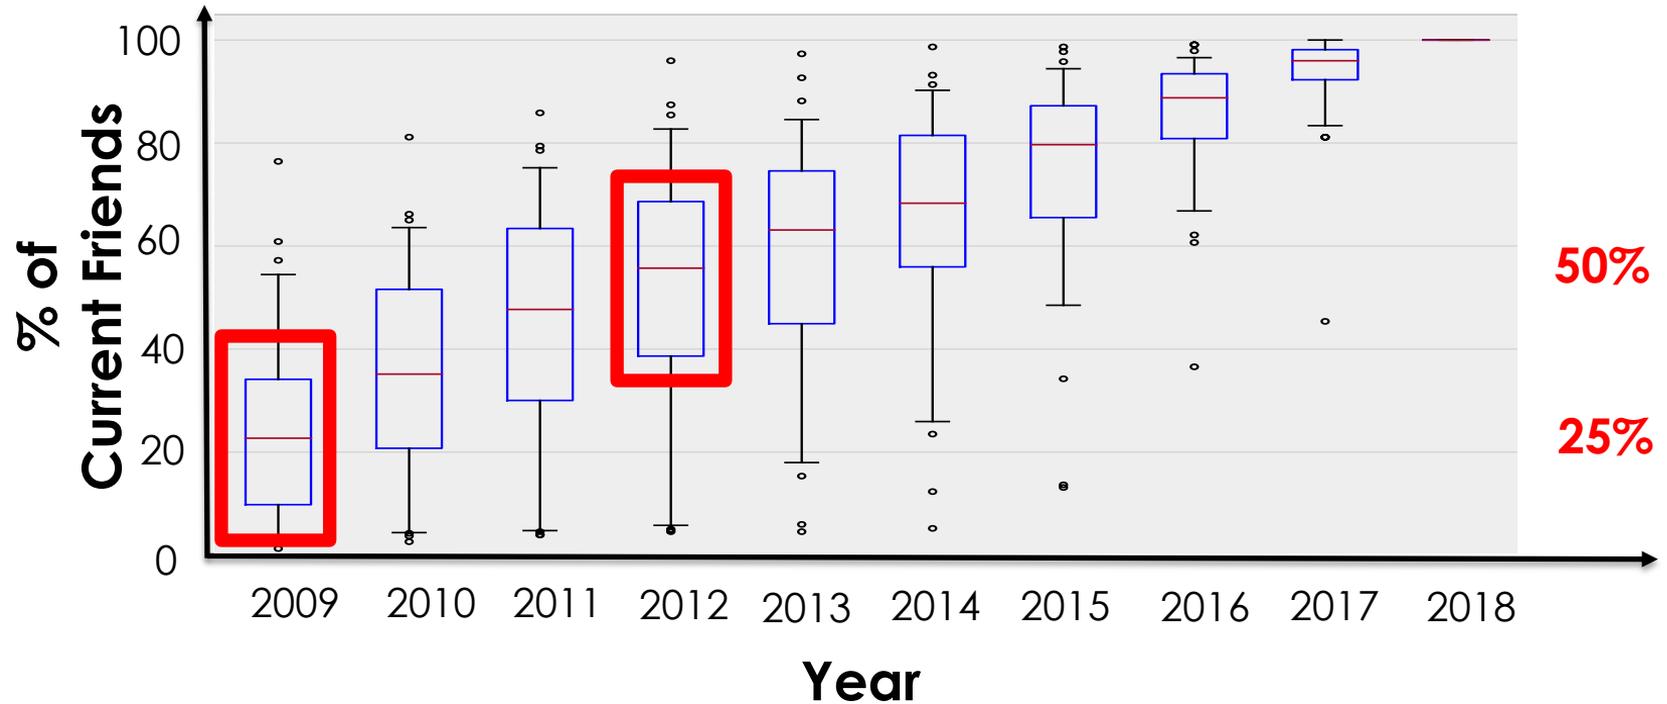
Change in number of friends



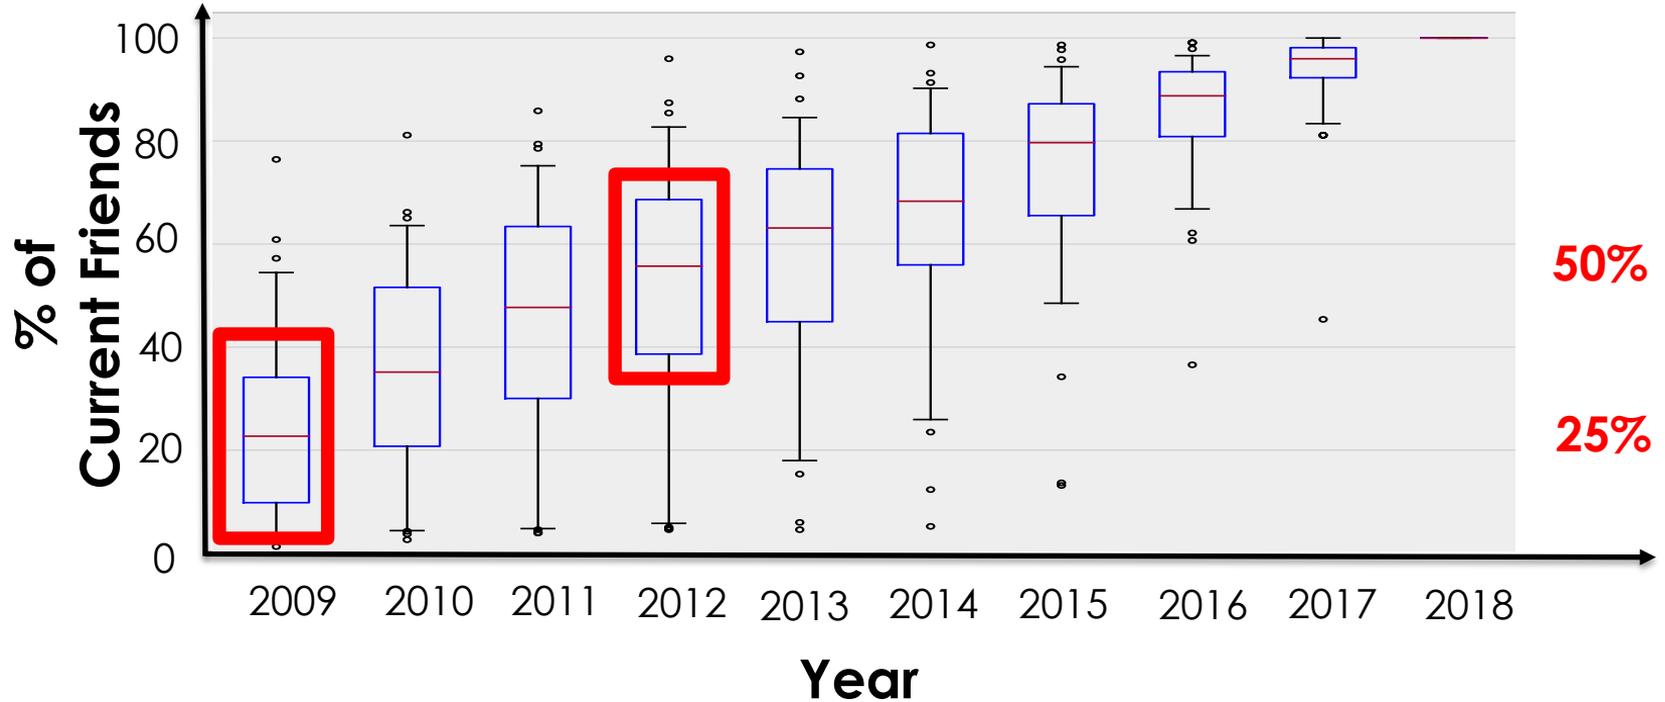
Change in number of friends



Change in number of friends



Change in number of friends



Substantial change in the meaning of “friends” privacy setting

Outline

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**

Current setting	<u>Desired setting</u>						
	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**

Current setting	<u>Desired setting</u>						
	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](#)

[Limit Past 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?

No Thanks

Let's Do It!

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

Outline

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



with



Stop sharing



with



Stop sharing



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

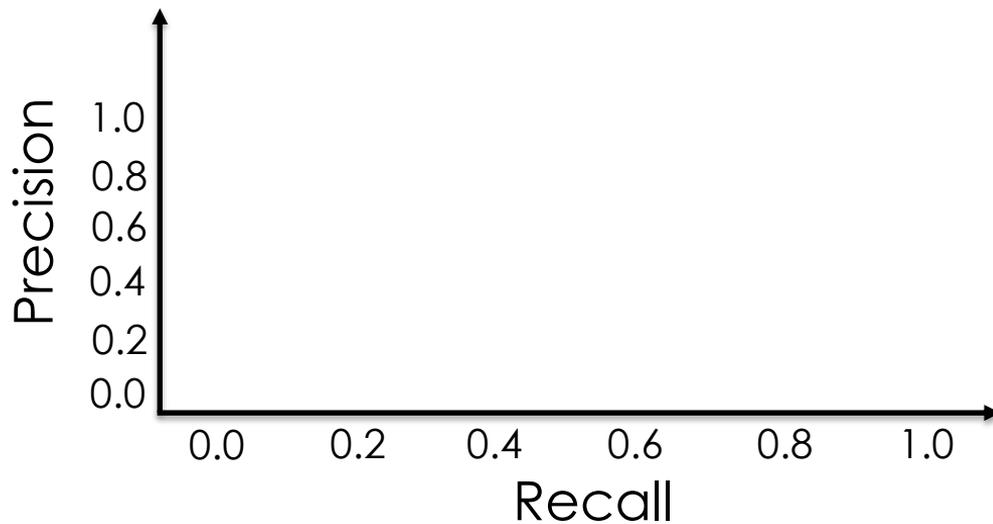
Random Forests, XGBoost, Decision Trees, Logistic Regression, Support Vector Machines, Deep Neural Networks

Baselines

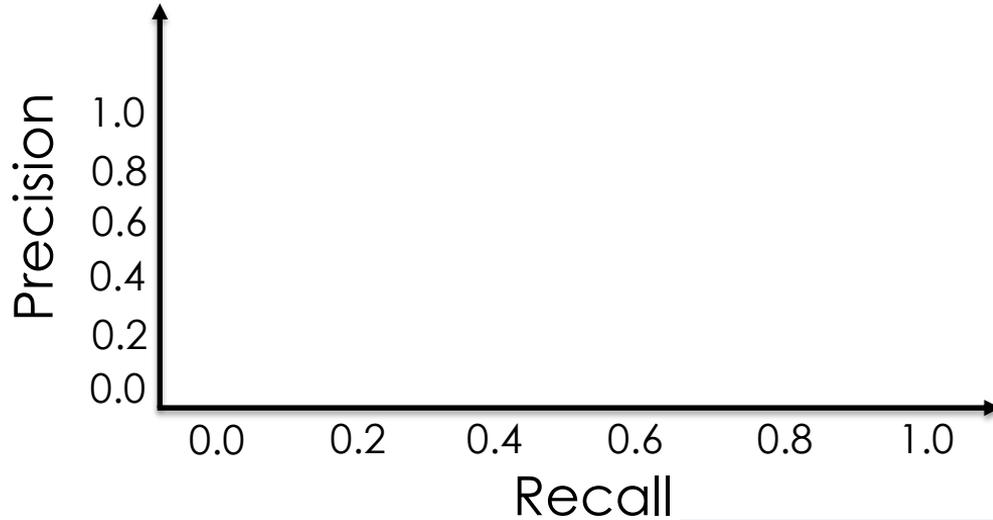
Random: Randomly predicts “stop sharing” for a pair

Interaction: Low interaction level → “stop sharing”

Are our models better than the baselines?

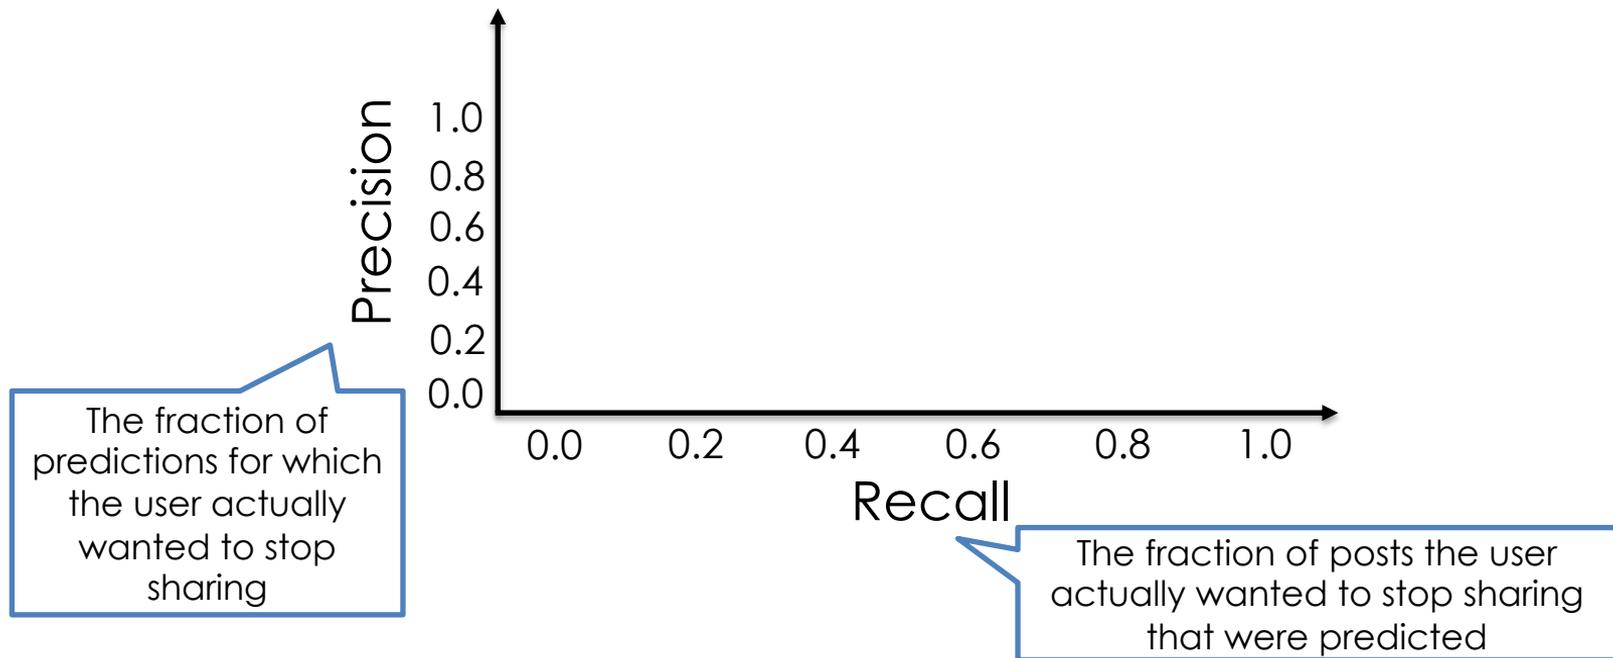


Are our models better than the baselines?

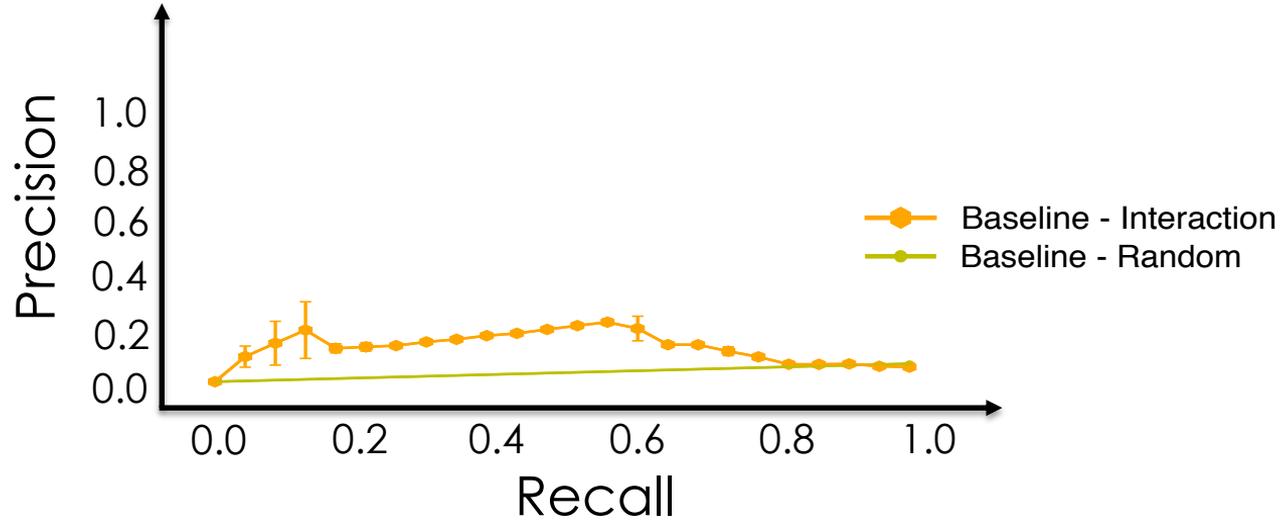


The fraction of posts the user actually wanted to stop sharing that were predicted

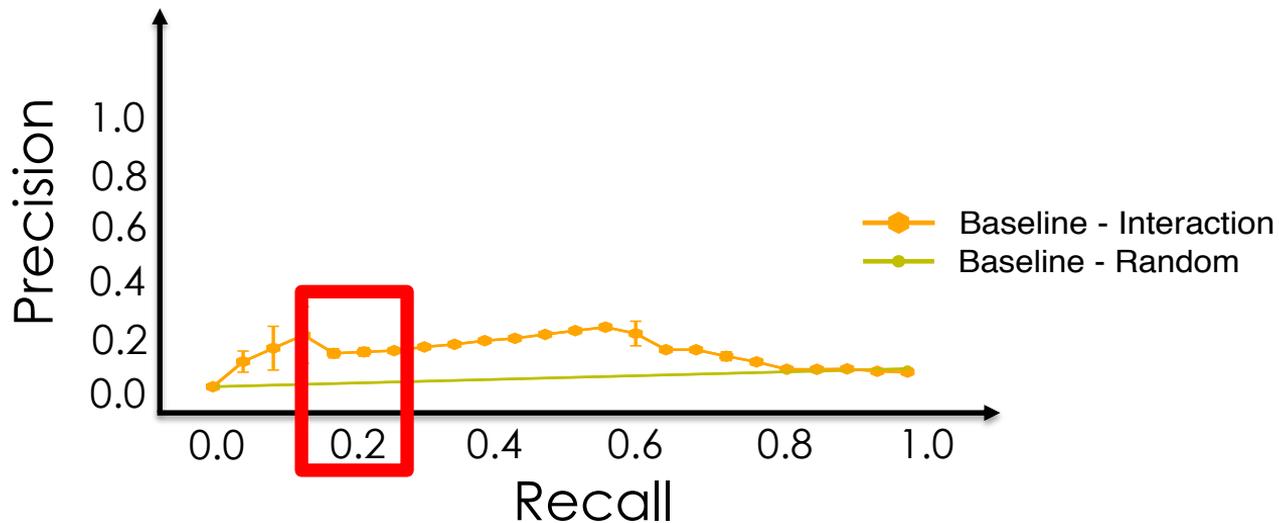
Are our models better than the baselines?



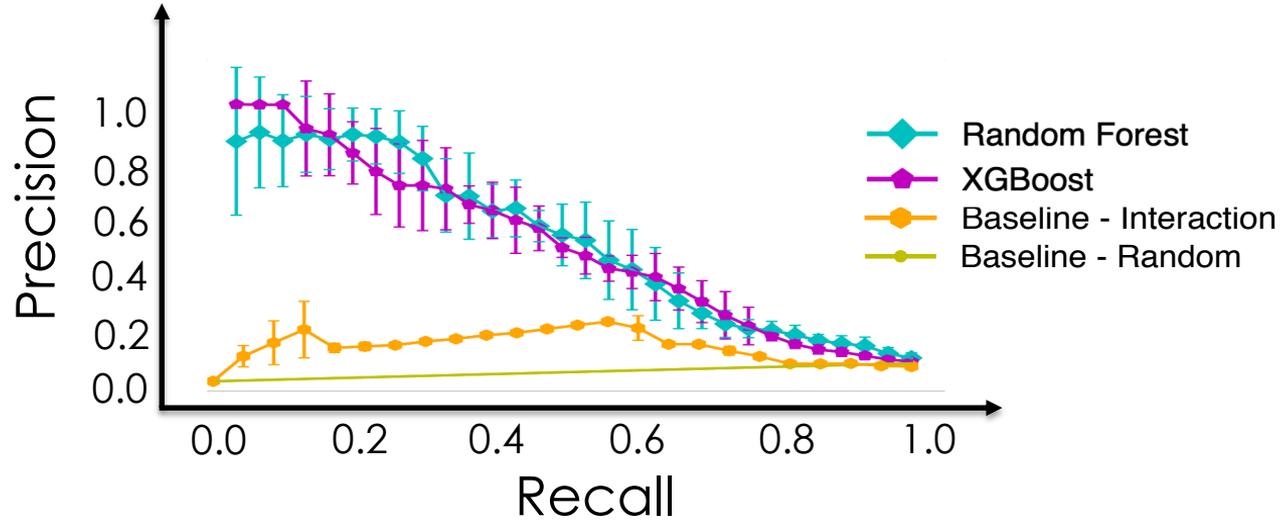
Are our models better than the baselines?



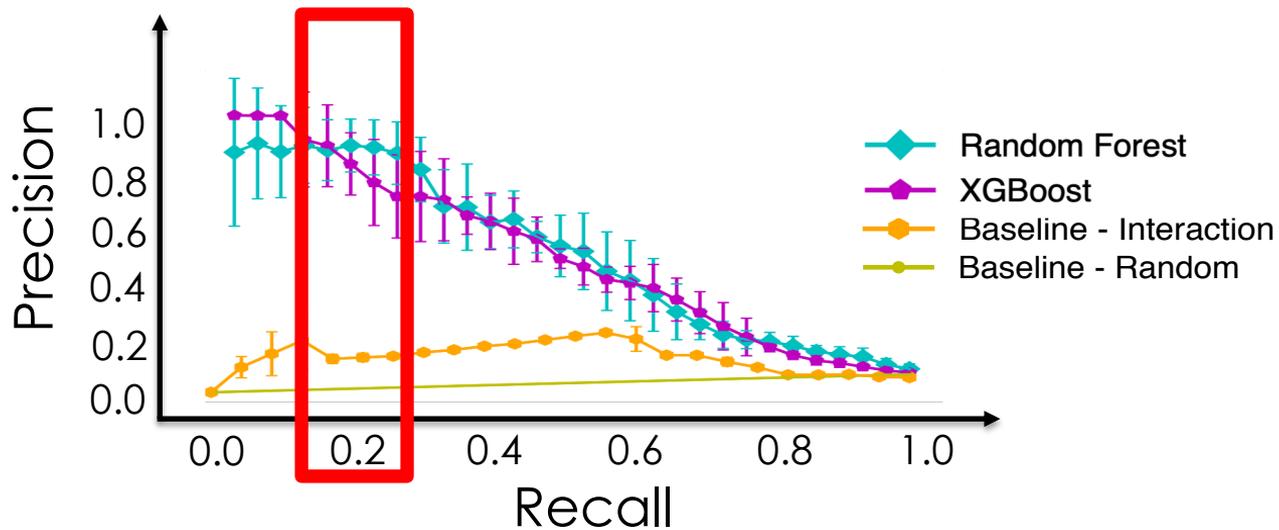
Are our models better than the baselines?



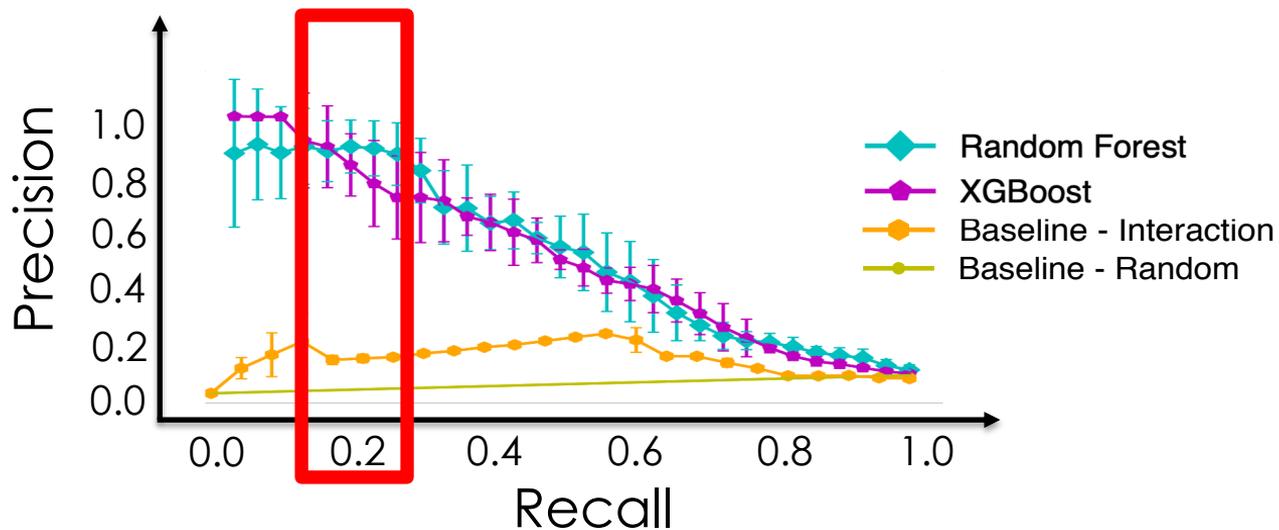
Are our models better than the baselines?



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

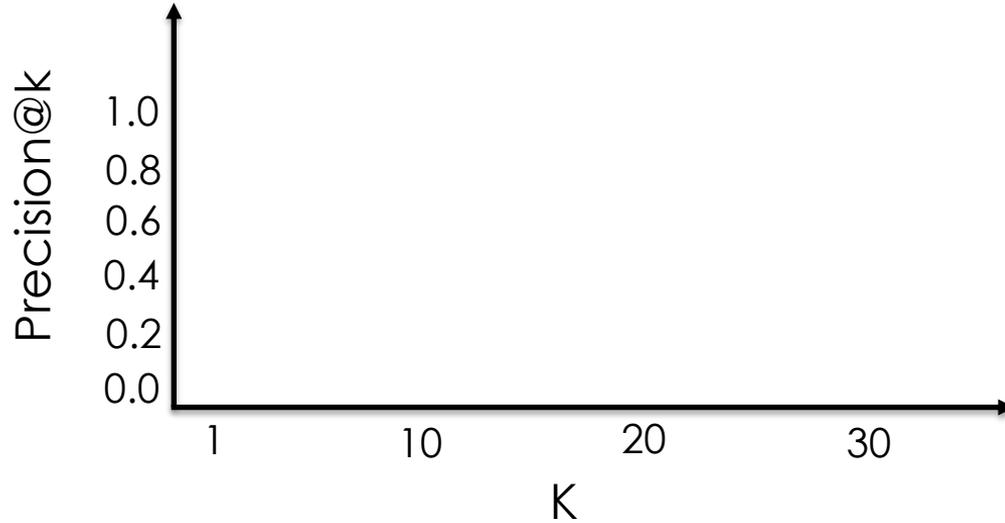
Output

List of friend-post pairs ordered by probability

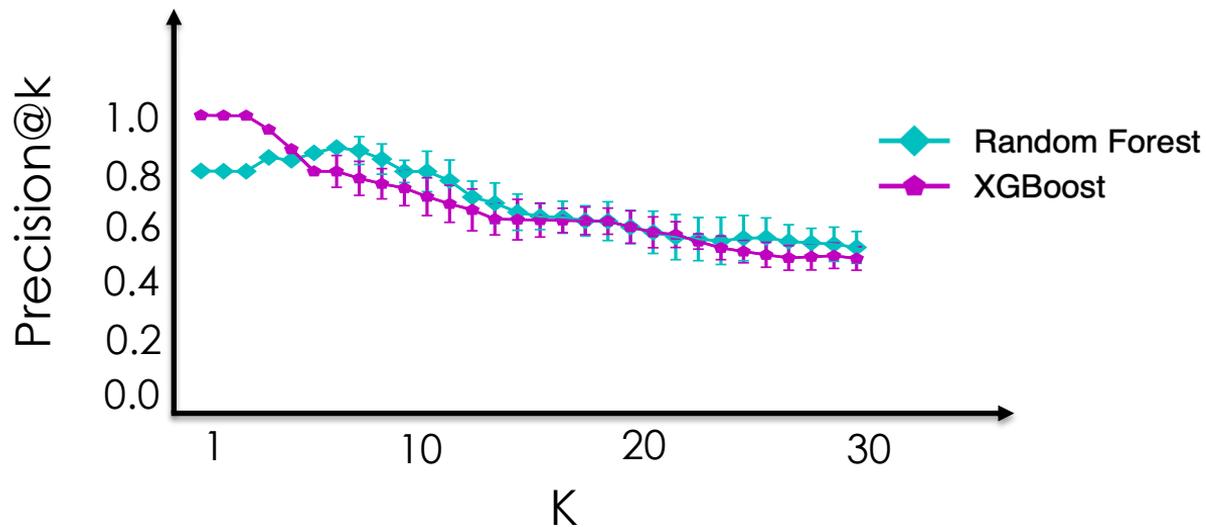
Ground truth

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

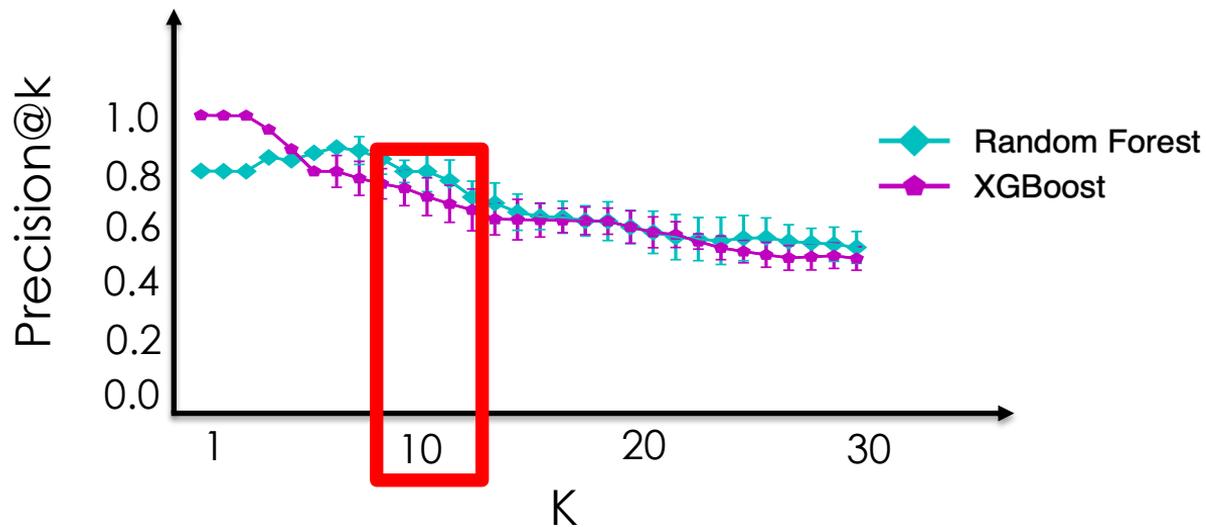
Recommendation accuracy of our models



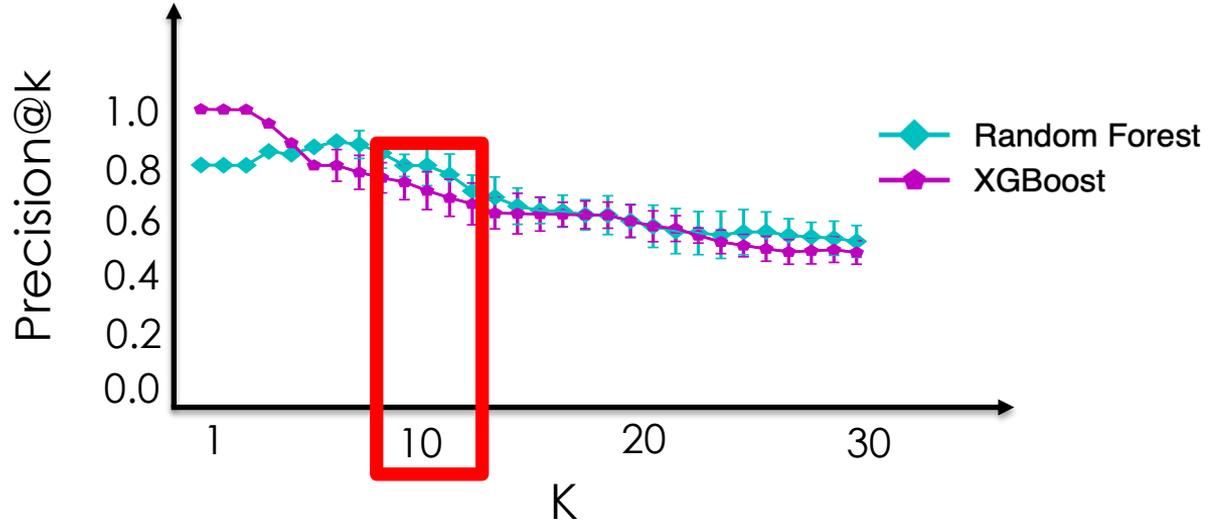
Recommendation accuracy of our models



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

Summary

Median 1,840 #old posts, 400% increase in audience size within 10 years

Users want to change privacy for 25% of old posts

Predictive recommender system achieves > 80% precision identifying some friend-post pairs needing privacy change

Summary

Median 1,840 #old posts, 400% increase in audience size within 10 years

Users want to change privacy for 25% of old posts

Predictive recommender system achieves $> 80\%$ precision identifying some friend-post pairs needing privacy change

Moving Beyond Set-It-And-Forget-It Privacy Settings on Social Media

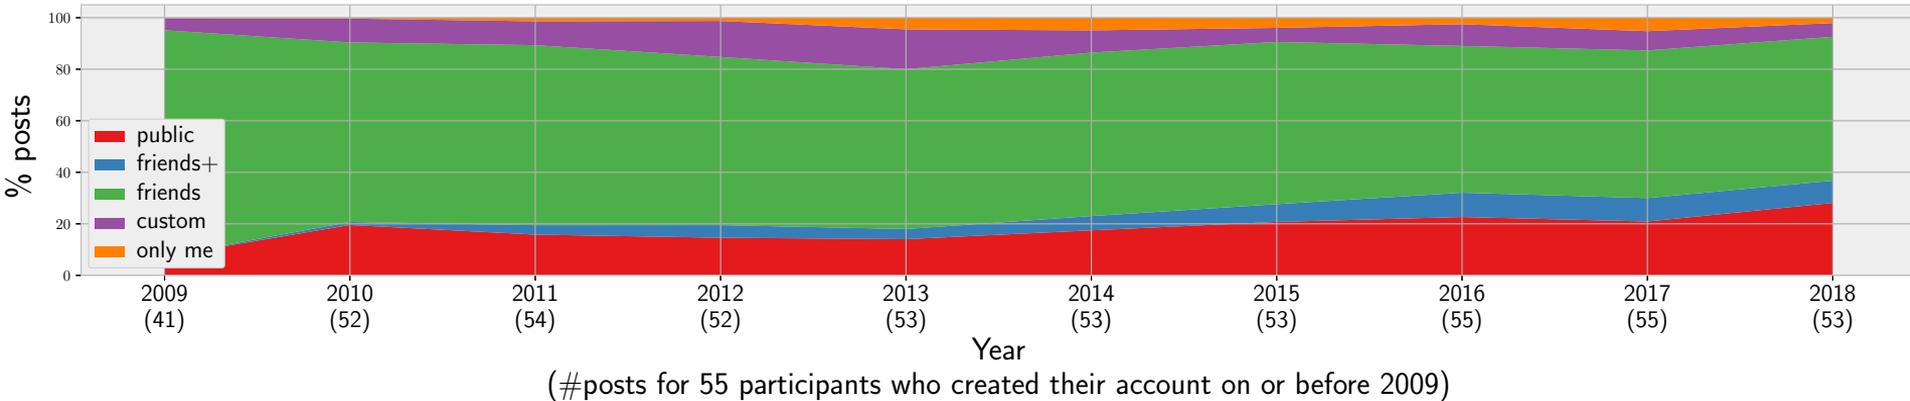
Mainack Mondal, Günce Su Yilmaz, Noah Hirsch, Mohammad Taha Khan, Michael Tang, Christopher Tran, Chris Kanich, Blase Ur, Elena Zheleva.



Extra slides!

Privacy setting change over time

Normalized by users



Worry about others checking old posts

86% expect other people check their old posts

55% reported checking others' old posts

Purposes: checking past relationships, family history

Might partially **motivate changes in privacy settings**



Correlating interaction / friend-addition time

Post-specific survey: privacy setting change for friend-post pairs
Friends chosen based on interaction frequency and addition time

Hypothesis: Desired privacy changes correlate with interaction frequency
Found only weak correlation
Mostly offline communication with some close relations

Hypothesis: Desired privacy changes correlate with when friend was added
Found no correlation