THE UNIVERSITY OF CHICAGO

Enforcing Contextual Integrity With Exposure Control

Mainack Mondal and Blase Ur

University of Chicago

CI Symposium, September 2018

Understanding privacy

There are a number of definitions

- Warren and Brandeis (1890)
- Westin's definition (1967)
  - •
  - •
  - •
- Solove's taxonomy of privacy (2008)
- Nissenbaum's privacy as contextual integrity (2010)

Understanding privacy

There are a number of definitions

- Warren and Brandeis (1890)
- Westin's definition (1967)
  - •
  - •
  - •
- Solove's taxonomy of privacy (2008)
- **Nissenbaum's privacy as contextual integrity (2010)**

Privacy as contextual integrity

A framework to argue about privacy violation

Privacy is preserved by **appropriate flows of information**

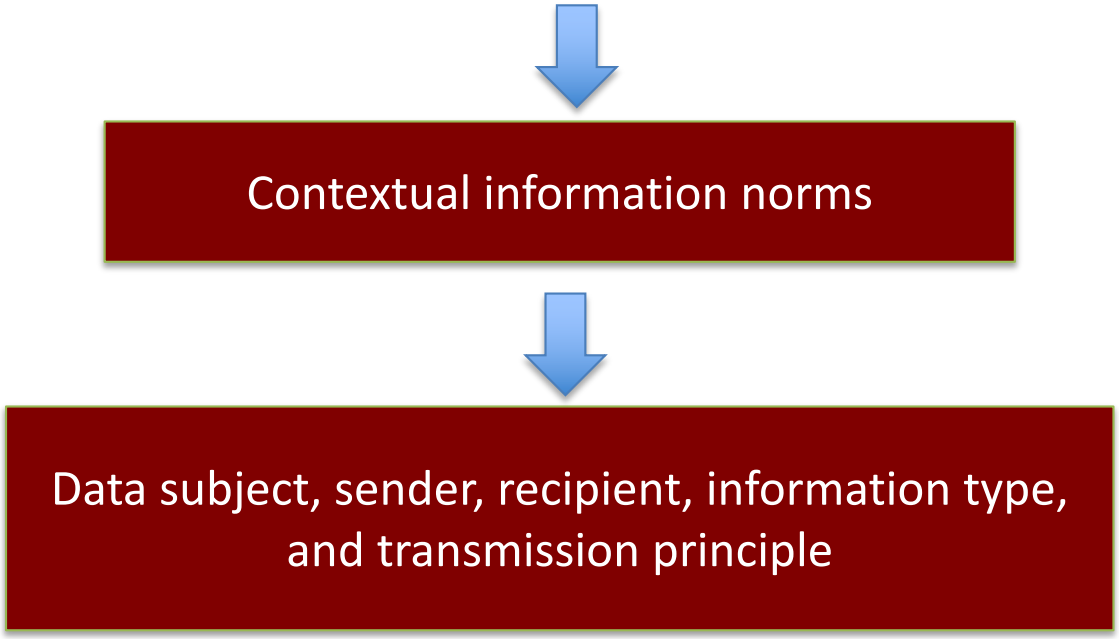

Conceptions of privacy are based on **dynamic ethical concerns**

Contextual integrity explains “what of privacy”

A **subsequent** step is to **build privacy preserving mechanisms**

State of the art: Access control model

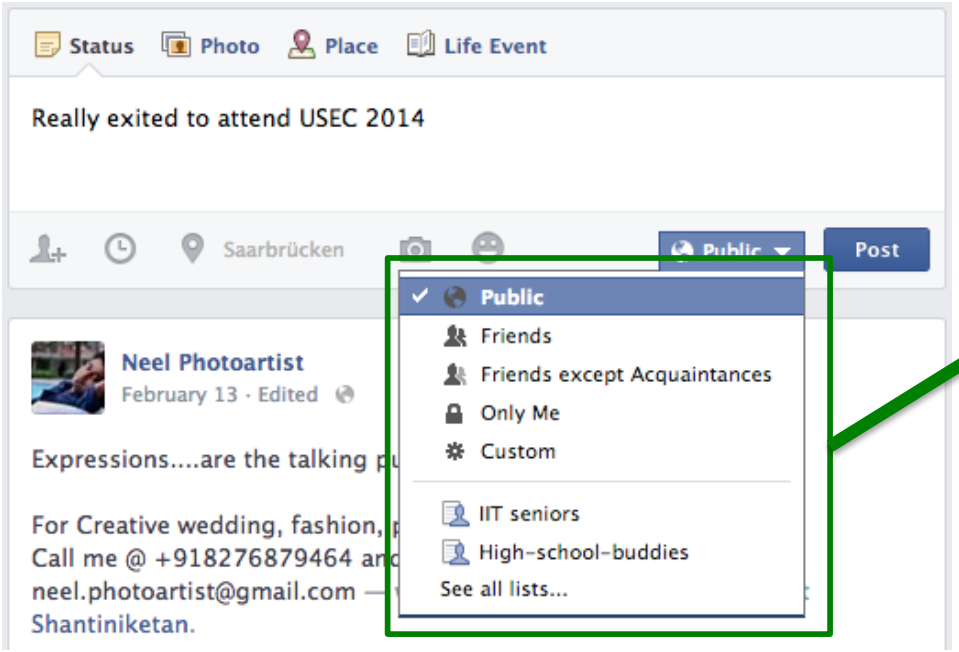

Access control lists (ACLs):
Allow others
access to content

Privacy violation from access control point of view:

If someone accesses content who the user did not allow

Access control model is useful to enforce contextual integrity where **all five parameters are explicitly expressed** via ACLs

Privacy violations in the real world

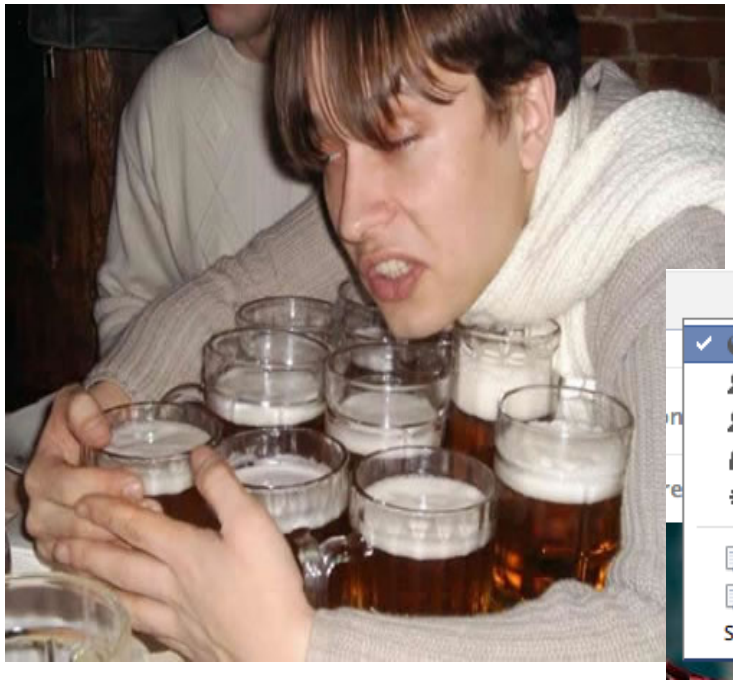

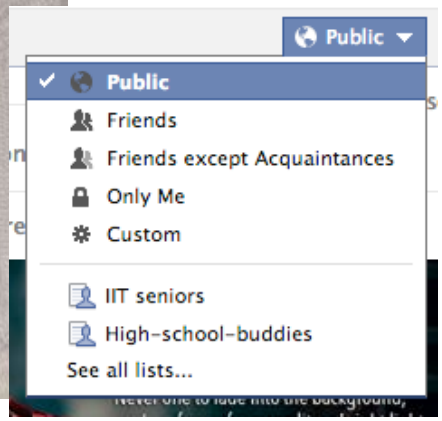

reddit

Privacy violation happened due to **increased accessibility**

Recipient and transmission principle violated

Privacy violation in real world from user's point of view:

If someone **views** content who the **user did not expect**

Access control is **inadequate** to capture **many** such violations!

Scenario 1: Facebook newsfeed

Facebook pushes your content as updates

Others automatically get your content
when they login to their Facebook page

After Newsfeed: More people actually saw the content

Users complained of privacy violation [Boyd et al. '08]

Contextual integrity is violated

Before and after Newsfeed: **access control did not change**!

Scenario 2: Facebook timeline

Sort your content by upload time

Others can search by time

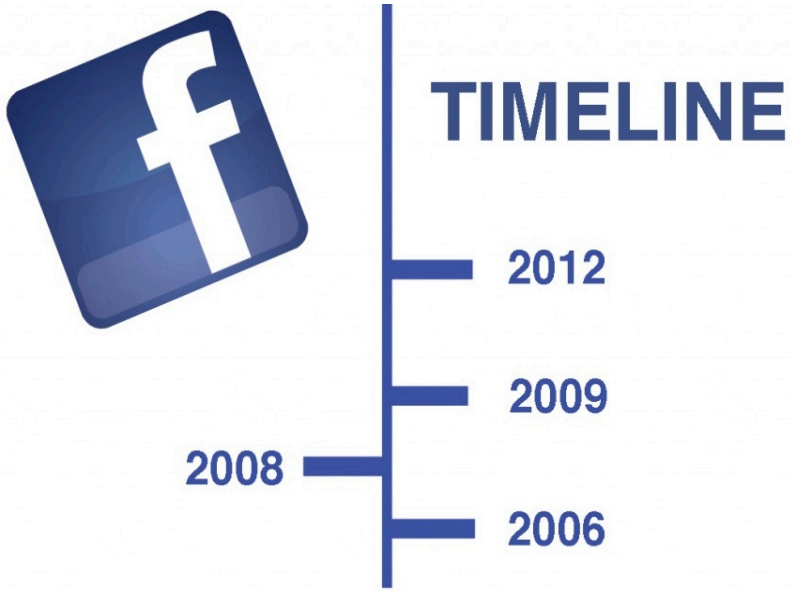

After timeline: Old content became easily searchable

Users felt privacy was violated

Contextual integrity is violated

readwrite

Before and after Timeline: **access control did not change**!

Scenario 3: Spokeo

Service aggregating public data from web

Others get all of this data by searching Spokeo

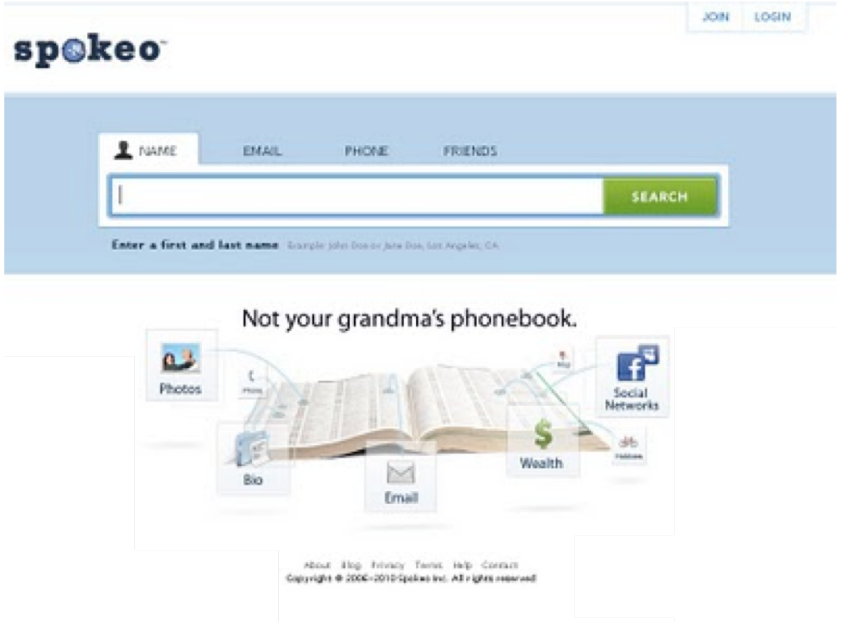

After aggregation: Inferring non public data become easier

Users complained of privacy violation

Contextual integrity is violated

Before and after aggregation: **access control did not change!**

Each of the cases violate contextual integrity

However access control was not violated in any of the cases

Take away 1: Access control is inadequate to capture user intention and enforce contextual integrity

Outline

Access control is inadequate to enforce contextual integrity

Exposure: An extension of access control to better enforce contextual integrity

Discussion: How to **manage privacy via exposure**

Exposure : Definition

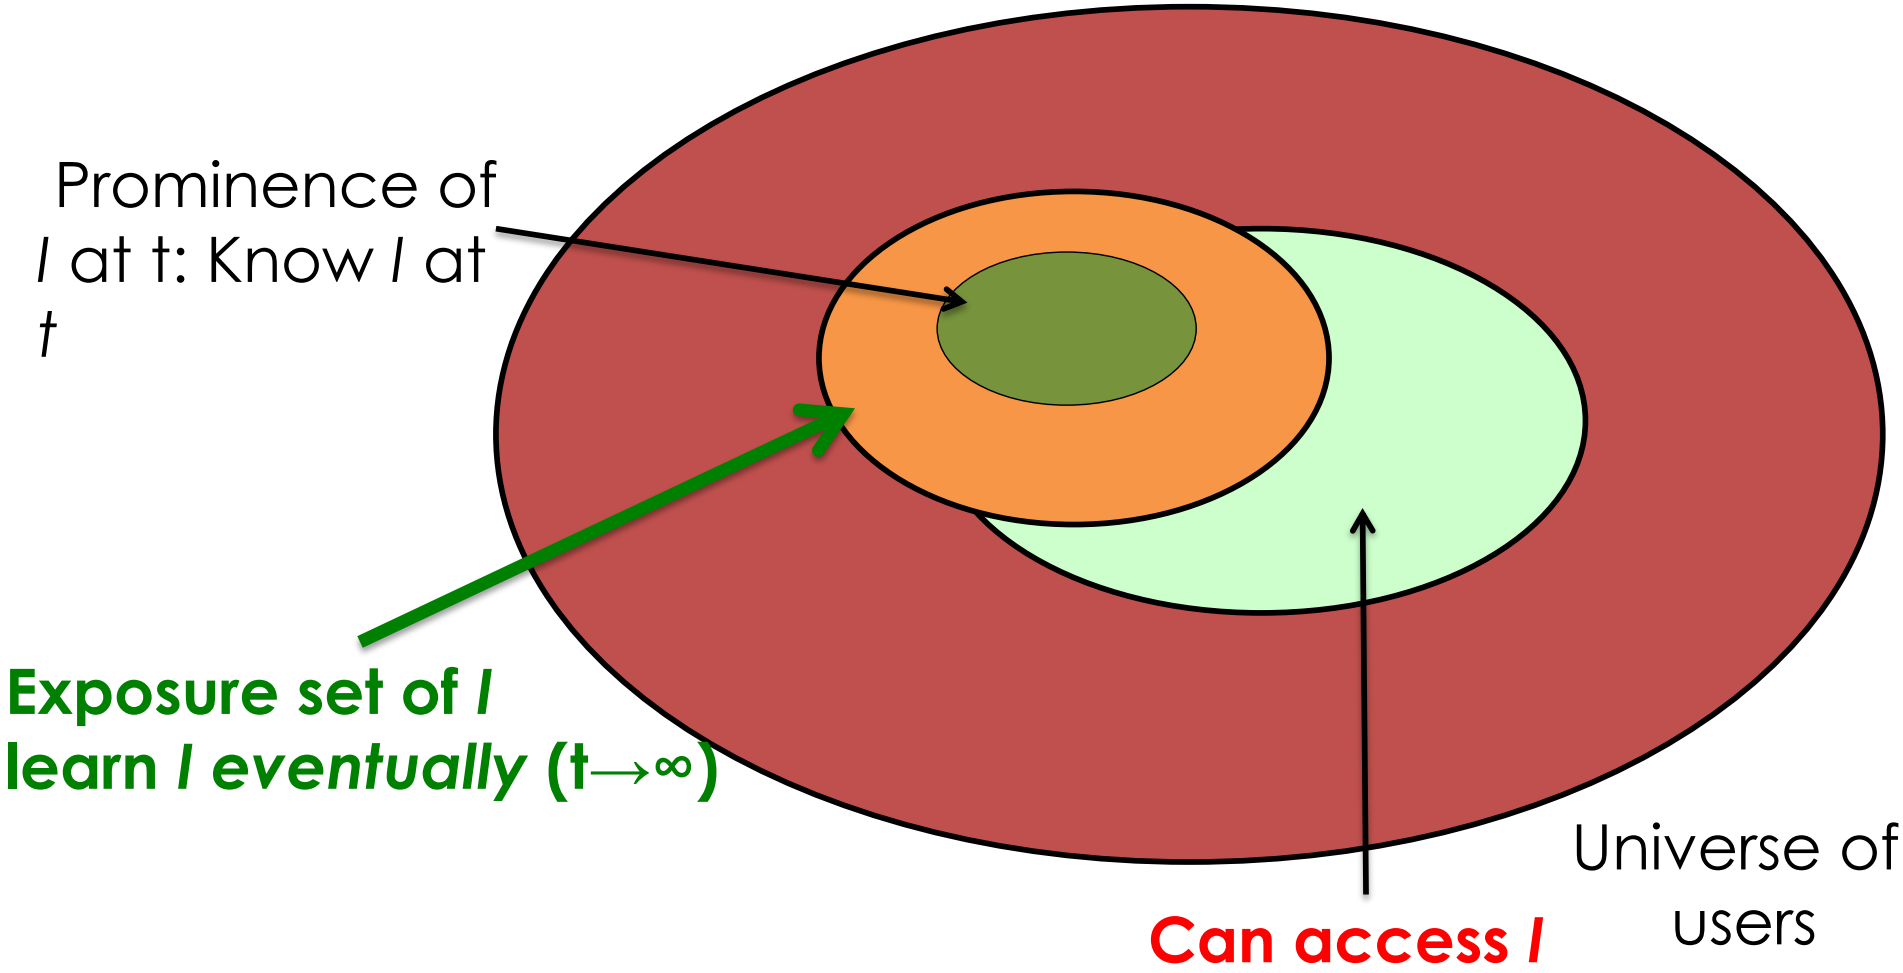

Exposure for content *I*

The set of people who will learn *I* eventually

How accurately do users estimate exposure?

Facebook researchers did a study with 589 users

[Bernstein et al. 2013]

Perceived exposure grossly underestimates actual exposure

There may be a feeling of privacy violation when actual exposure is different from perceived exposure

Exposure in more detail

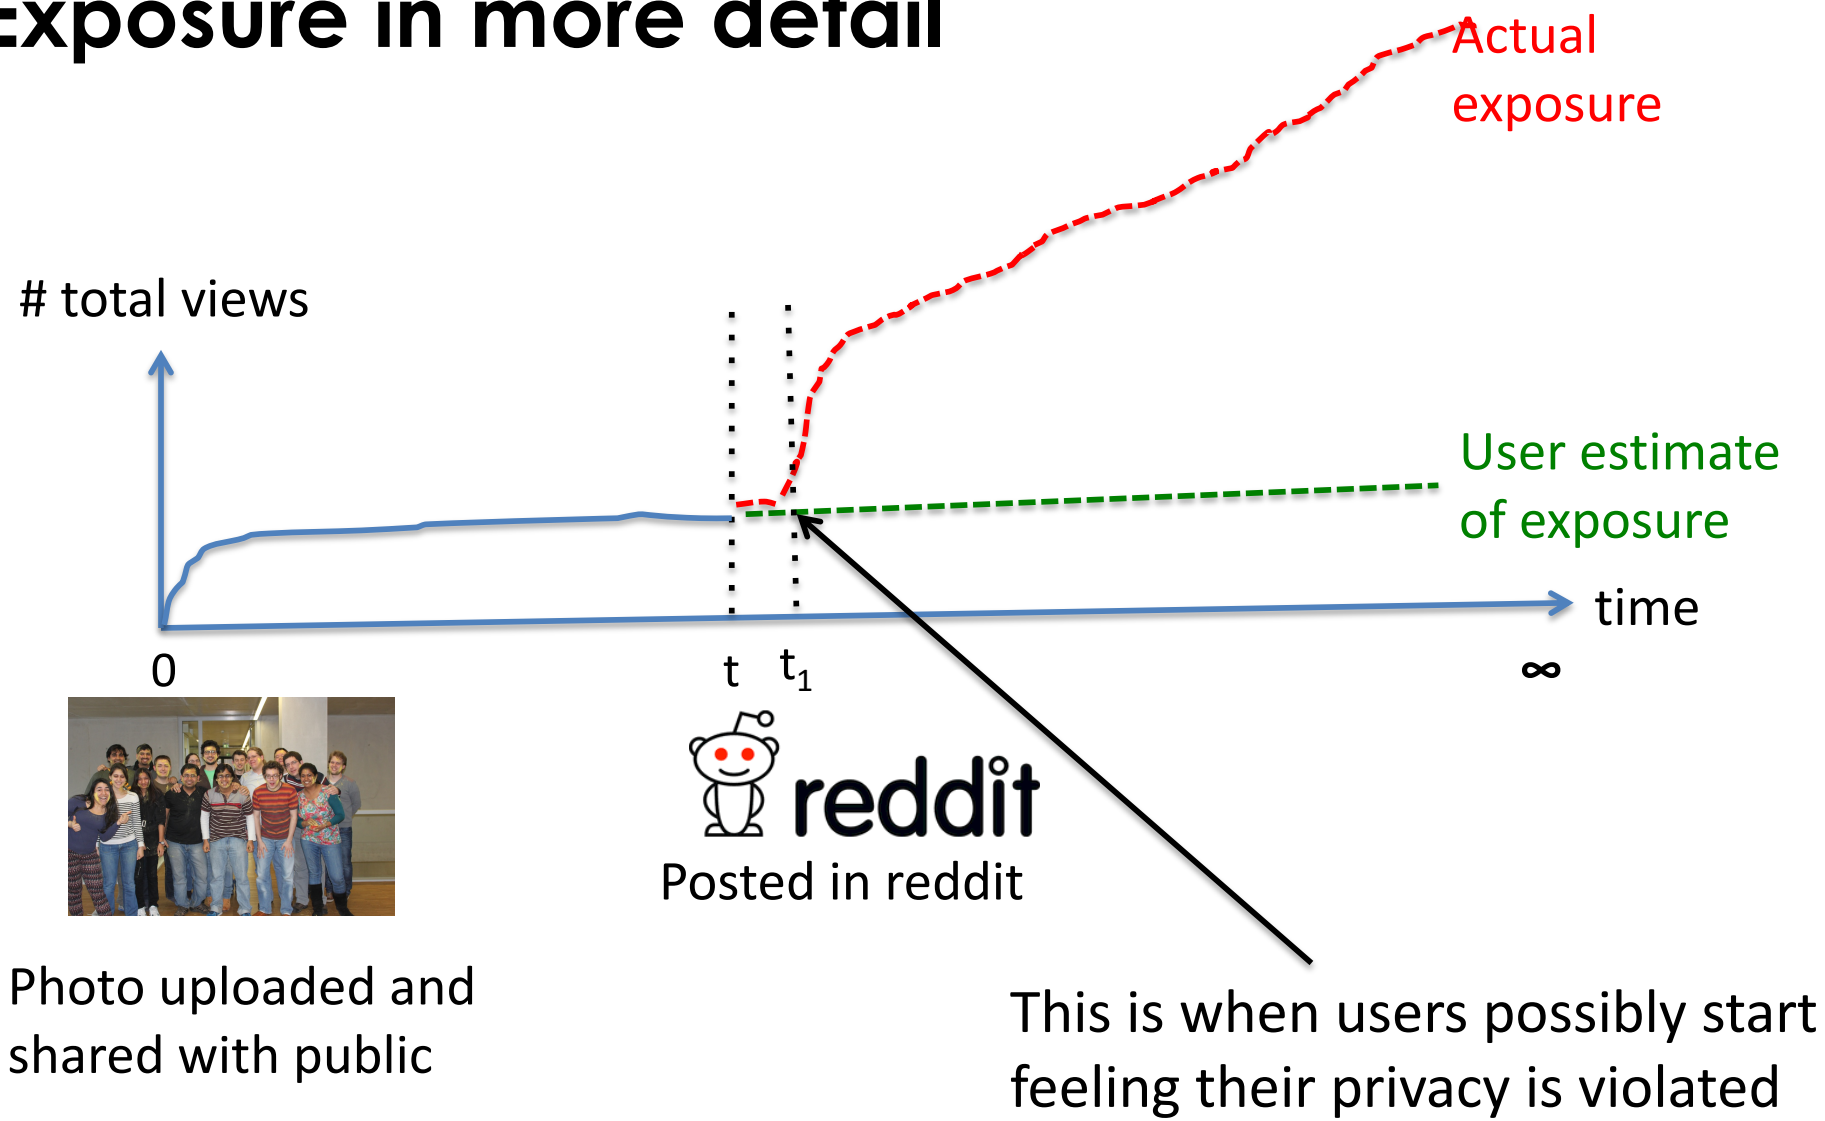

Revisiting scenario 1: Facebook newsfeed

Exposure before newsfeed

Friends who visit profile

Exposure after newsfeed

All the friends who are logged into Facebook

Exposure of uploaded information **after newsfeed** > **Exposure** of uploaded information **before newsfeed**

Revisiting scenario 2: Facebook timeline

Exposure of old content before timeline

Users who will scroll down

thousands of content

Exposure of old content after timeline

All users who search by time

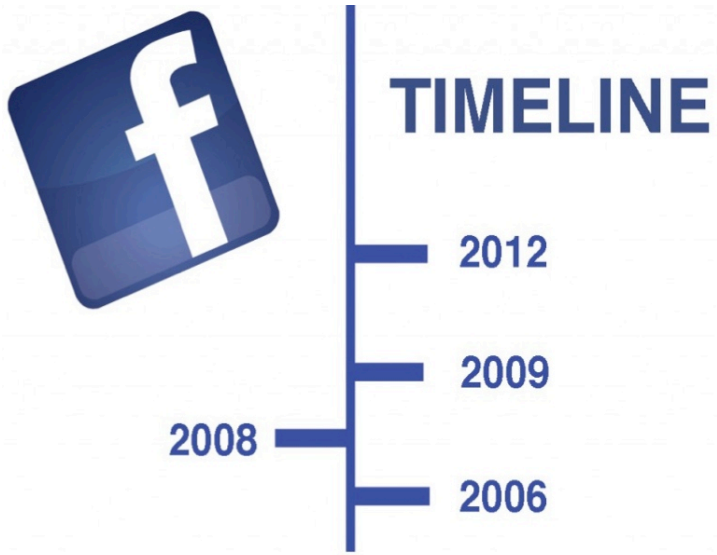

Exposure of old information **after timeline** > **Exposure** of old information **before timeline**

Revisiting scenario 3: Spokeo

Exposure before aggregation

Users who collect content themselves from multiple sources

Exposure after aggregation

Any user who searches in Spokeo

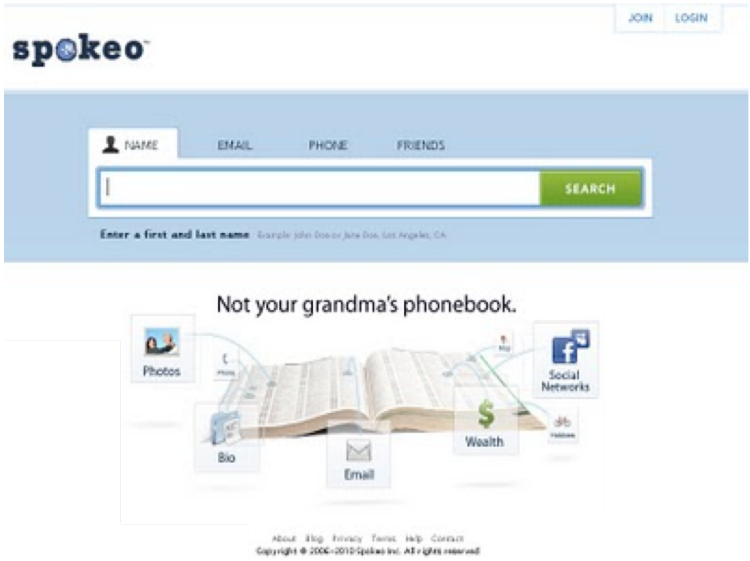

Exposure of inferred information **after aggregation** > **Exposure** of inferred information **before aggregation**

Take away 2: Exposure control extends access control and capture violations of contextual integrity which are not captured by access control

Outline

Access control is inadequate to enforce contextual integrity

Exposure: An extension of access control to better enforce contextual integrity

Discussion: How to **manage privacy via exposure**

Discussion: Managing privacy via exposure

Challenge 1:

How to estimate exposure for a content?

Challenge 2:

How to make users aware of the estimated exposure?

Challenge 3:

How to allow users more control over exposure?

Challenge 1: Estimating exposure

Situations where predicting exposure is very hard

Cross site prediction, exposure of inferred information

Situations where predicting exposure is possible

Predicting exposure of content in a site

Lots of research in content popularity growth

[Borghol et al] [Figueiredo et al.]
[Hong et al.] [Zaman et al]
[Bernstein et al.]

Challenge 1: Who can best estimate exposure

Platform operators are in the **best position to predict** exposure accurately with the data they collect

- They log who is accessing what content
- They collect historical data for content access

Platform operators can also **control** exposure

- They decide which content to show other users

Challenge 2: How to make users aware of the exposure?

Prediction can be shown to users at different granularity

- **List** of predicted people for a content
- **Number** of predicted people for a content
- Showing the prediction for a certain **time period**
- Showing the prediction with **error bounds**
- Showing how a **specific dissemination mechanism** changes the prediction
  - e.g., 200 more people are likely to see your content due to newsfeed

Challenge 3: How to allow users more control over exposure?

Different "knobs" can be provided to the user

- Change access control to a more restrictive setting
- Disabling particular dissemination mechanisms, e.g. search
- Enabling tripwires
  - Take content offline if more than 50 people view
  - Take content offline after two months
  - …

Enabling tripwires to control exposure

Limiting third party crawlers like Spokeo to control exposure

- Built **Genie**, a credit network based system
- Protect against crawling data from online social media

Understanding retrospective privacy management preferences for online social media site users

- Is contextual integrity violated over time due to static privacy settings?

Take away 3: There are lots of open challenges and substantial research opportunities in how to design and deploy exposure based systems

Conclusion

Take away 1: Access control is inadequate to capture user intention and enforce contextual integrity

Take away 2: Exposure based privacy model extends access control and better enforce contextual integrity

Take away 3: Lots of open challenges to design systems which can manage privacy by controlling exposure

Thank you!

Limitation of exposure control in enforcing contextual exposure

- Exposure control leverages user's past behavior as a proxy for user's mental model
  - Can not capture contextual norms rooted in social processes
  - E.g., keeping an acquaintance in the expected expected set who turns out to be a stalker