CS145: INTRODUCTION TO DATA MINING

18: Bias, Privacy, and Ethics Issues on Graph Data

Instructor: Yizhou Sun

yzsun@cs.ucla.edu

December 1, 2021

Content

- Bias
- Privacy
- Social influence

slido

Do you think your friends in average have more friends than you?

ⓘ Start presenting to display the poll results on this slide.

Friendship Paradox

- Your friends have more friends than you?

https://arxiv.org/pdf/1605.04470.pdf

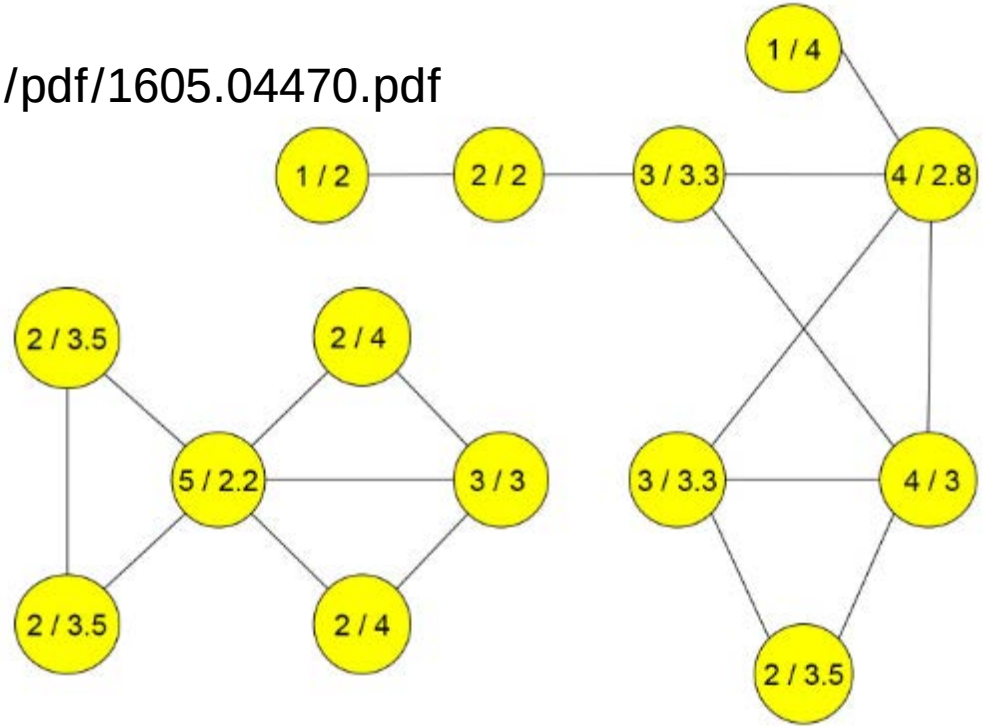

Figure 1: Data from James Coleman's (1961) study of high school friendships. Nodes are girls and links are mutual friendships. The first number listed for each girl is how many friends the girl has and the second number is the average number of friends that the girl's friends have. For instance, the girl in the lower left-hand corner has 2 friends, and those friends have 2 and 5 friends, for an average of 3.5.

Some Python Code

- https://towardsdatascience.com/observe-the-friend-paradox-in-facebook-data-using-python-314c23fd49e4

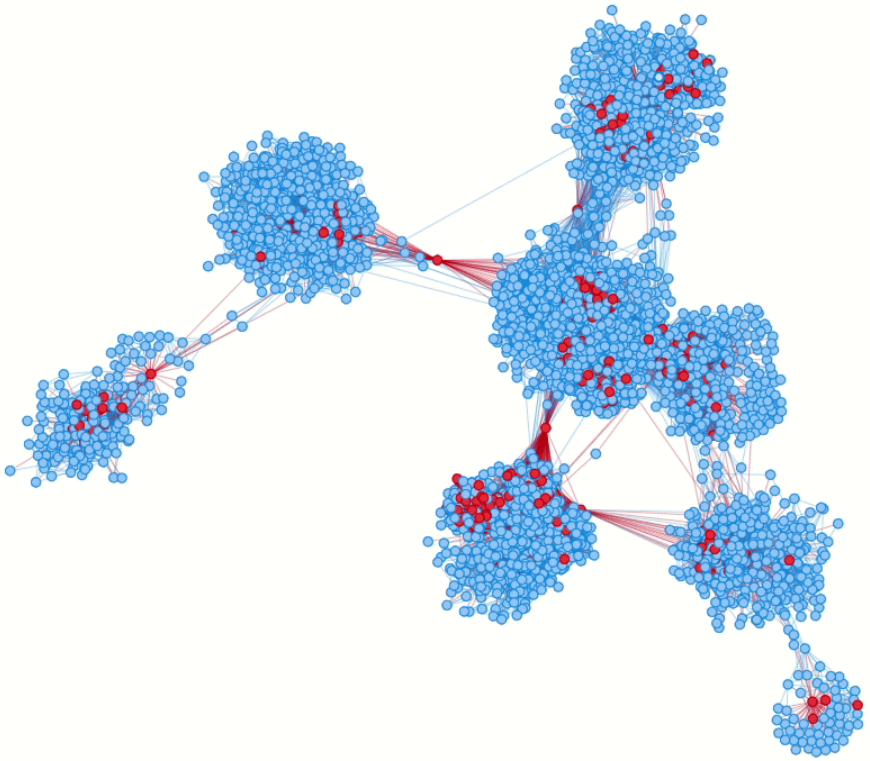

Game Time

- Given a social network, with n=30, each one with color red or blue

- Now everyone is assigned with a note
  - Red, {5 red, 6 blue} means your color is red, your neighbors have 5 reds and 6 blues

slido

Do you see more reds or blues in your local neighborhood including yourself and friends?

ⓘ Start presenting to display the poll results on this slide.

slido

Based on previous result, which color do you think is more popular?

ⓘ Start presenting to display the poll results on this slide.

slido

Tell us the color assigned to you

ⓘ Start presenting to display the poll results on this slide.

Friendship Paradox

- Your opinion is different from “majority”?

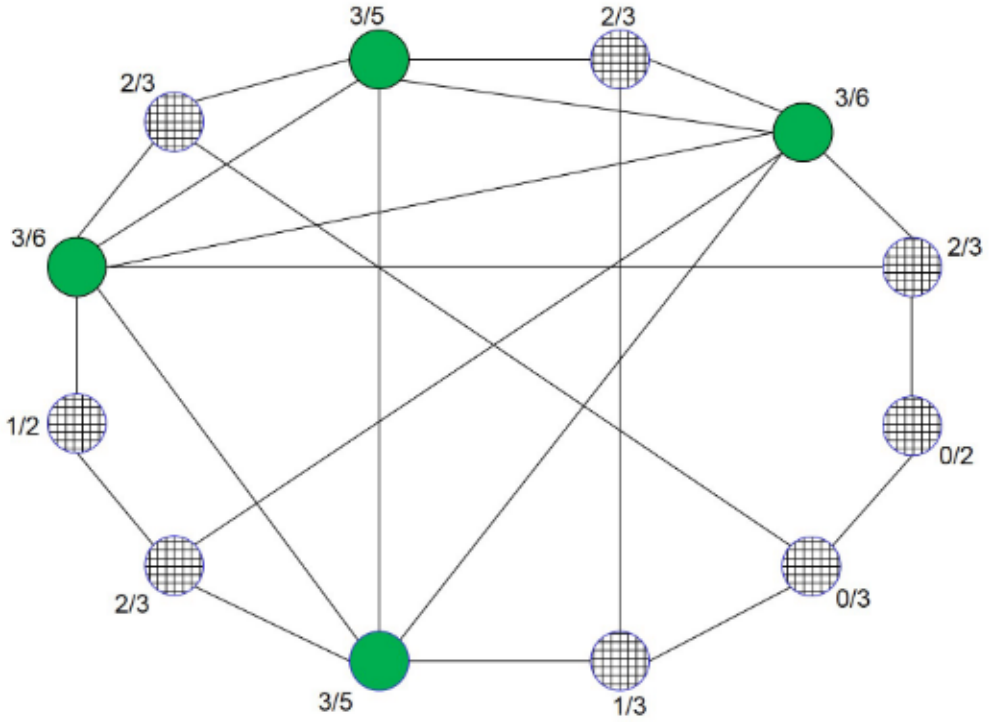

Figure 3: The Friendship paradox at work. The four most connected agents have a base preference for solid and the eight others prefer plaid. The fractions next to the agents are their perceptions of the preferences are for solids over plaids, based on what they see among their friends in the first period. Most of them perceive a majority preference for solid, with only the few agents in the lower right perceiving a majority for plaids.

Content

- Bias
- Privacy 
- Social influence

Facebook MyPersonality App

- https://www.youtube.com/watch?v=GOZArvMMHKs

PNAS

Private traits and attributes are predictable from digital records of human behavior

Michal Kosinski[a,1], David Stillwell[a], and Thore Graepel[b]

[a]Free School Lane, The Psychometrics Centre, University of Cambridge, Cambridge CB2 3RQ United Kingdom; and [b]Microsoft Research, Cambridge CB1 2FB, United Kingdom

Edited by Kenneth Wachter, University of California, Berkeley, CA, and approved February 12, 2013 (received for review October 29, 2012)

We show that easily accessible digital records of behavior, Facebook Likes, can be used to automatically and accurately predict a range of highly sensitive personal attributes including: sexual orientation, ethnicity, religious and political views, personality traits, intelligence, happiness, use of addictive substances, parental separation, age, and gender. The analysis presented is based on a dataset of over 58,000 volunteers who provided their Facebook Likes, detailed demographic profiles, and the results of several psychometric tests. The proposed model uses dimensionality reduction for preprocessing the Likes data, which are then entered into logistic/linear regression to predict individual psychodemographic profiles from Likes. The model correctly discriminates between homosexual and heterosexual men in 88% of cases, African Americans and Caucasian Americans in 95% of cases, and between Democrat and Republican in 85% of cases. For the personality trait "Openness," prediction accuracy is close to the test–retest accuracy of a standard personality test. We give examples of associations between attributes and Likes and discuss implications for online personalization

browsing logs (11–15). Similarly, it has been shown that personality can be predicted based on the contents of personal Web sites (16), music collections (17), properties of Facebook or Twitter profiles such as the number of friends or the density of friendship networks (18–21), or language used by their users (22). Furthermore, location within a friendship network at Facebook was shown to be predictive of sexual orientation (23).

This study demonstrates the degree to which relatively basic digital records of human behavior can be used to automatically and accurately estimate a wide range of personal attributes that people would typically assume to be private. The study is based on Facebook Likes, a mechanism used by Facebook users to express their positive association with (or "Like") online content, such as photos, friends' status updates, Facebook pages of products, sports, musicians, books, restaurants, or popular Web sites. Likes represent a very generic class of digital records, similar to Web search queries, Web browsing histories, and credit card purchases. For example, observing users' Likes related to music


Like

How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did

impending bundle of joy. Target assigns every customer a Guest ID number, tied to their credit card, name, or email address that becomes a bucket that stores a history of everything they've bought and any demographic information Target has collected from them or bought from other sources. Using that, Pole looked at historical buying data for all the ladies who had signed up for Target baby registries in the past. From the NYT:

https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/?sh=60853d6a6668

TARGET

> [Pole] ran test after test, analyzing the data, and before long some useful patterns emerged. Lotions, for example. Lots of people buy lotion, but one of Pole's colleagues noticed that women on the baby registry were buying larger quantities of unscented lotion around the beginning of their second trimester. Another analyst noted that sometime in the first 20 weeks, pregnant women loaded up on supplements like calcium, magnesium and zinc. Many shoppers purchase soap and cotton balls, but when someone suddenly starts buying lots of scent-free soap and extra-big bags of cotton balls, in addition to hand sanitizers and washcloths, it signals they could be getting close to their delivery date.

> One Target employee I spoke to provided a hypothetical example. Take a fictional Target shopper named Jenny Ward, who is 23, lives in Atlanta and in March bought cocoa-butter lotion, a purse large enough to double as a diaper bag, zinc and magnesium supplements and a bright blue rug. There's, say, an 87 percent chance that she's pregnant and that her delivery date is sometime in late August.

made up -- that conveys how eerily accurate the targeting is. An angry man went into a Target outside of Minneapolis, demanding to talk to a manager:

> "My daughter got this in the mail!" he said. "She's still in high school, and you're sending her coupons for baby clothes and cribs? Are you trying to encourage her to get pregnant?"

Target knows before it shows.

> The manager didn't have any idea what the man was talking about. He looked at the mailer. Sure enough, it was addressed to the man's daughter and contained advertisements for maternity clothing, nursery furniture and pictures of smiling infants. The manager apologized and then called a few days later to apologize again.

> On the phone, though, the father was somewhat abashed. "I had a talk with my daughter," he said. "It turns out there's been some activities in my house I haven't been completely aware of. She's due in August. I owe you an apology."

slido

Which privacy policy do you agree (multiple choices)?

ⓘ Start presenting to display the poll results on this slide.

Content

- Bias
- Privacy
- Social influence

PNAS
Proceedings of the
National Academy of Sciences
of the United States of America

RESEARCH ARTICLE

Experimental evidence of massive-scale emotional contagion through social networks

Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock

+ See all authors and affiliations

PNAS June 17, 2014 111 (24) 8788-8790; first published June 2, 2014; https://doi.org/10.1073/pnas.1320040111

Edited by Susan T. Fiske, Princeton University, Princeton, NJ, and approved March 25, 2014 (received for review October 23, 2013)

- https://www.pnas.org/content/111/24/8788

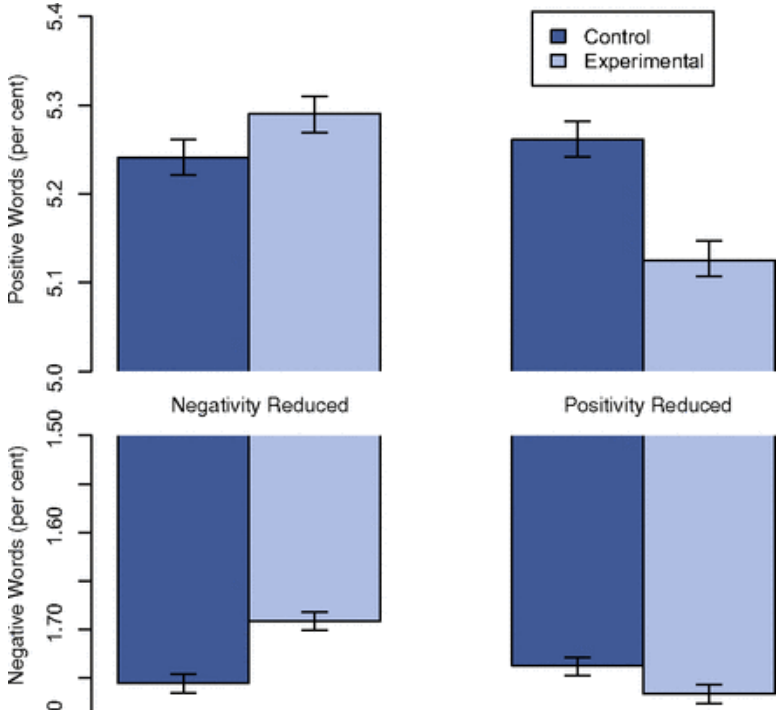

Scientists published a paper revealing that in 2012, Facebook researchers conducted a study into "emotional contagion." The social media company altered the news feeds (the main page users land on for a stream of updates from friends) of nearly 700,000 users. Feeds were changed to reflect more "positive" or "negative" content, to determine if seeing more sad messages makes a person sadder.

Facebook Manipulates Our Moods For Science And Commerce: A Roundup

Laurie Penny explains that the study's findings are not the point — that Facebook did this is the point — and argues *the potential* for more is why the research feels so wrong.

> "I am not convinced that the Facebook team knows what it's doing. It does, however, know what it *can* do — what a platform with access to the personal information and intimate interactions of 1.25 billion users can do. ...
>
> "What the company does now will influence how the corporate powers of the future understand and monetise human emotion. Dr Adam Kramer, the man behind the study and a longtime member of the company's research team, commented in an excited Q & A that 'Facebook data constitutes the largest field study in the history of the world.' The ethics of this situation have yet to be unpacked."

https://www.npr.org/sections/alltechconsidered/2014/06/30/326929138/facebook-manipulates-our-moods-for-science-and-commerce-a-roundup

Into Political Domain

Social media manipulation by political actors an industrial scale problem - Oxford report

SOCIAL SCIENCES | RESEARCH

Social media manipulation of public opinion is a growing threat to democracies around the world, according to the 2020 media manipulation survey from the Oxford Internet Institute, which found evidence in every one of the 80+ countries surveyed.

https://www.ox.ac.uk/news/2021-01-13-social-media-manipulation-political-actors-industrial-scale-problem-oxford-report

Key findings the OII researchers identified include:

- Private ‘strategic communications’ firms are playing an increasing role in spreading computational propaganda, with researchers identifying state actors working with such firms in 48 countries.
- Almost $60 million has been spent on firms who use bots and other amplification strategies to create the impression of trending political messaging.
- Social media has become a major battleground, with firms such as Facebook and Twitter taking steps to combat ‘cyber troops’, with some $10 million has been spent on social media political advertisements. The platforms removed more than 317,000 accounts and pages from ‘cyber troops’ actors between January 2019 and November 2020.

Professor Philip Howard, Director of the Oxford Internet Institute, and the report's co-author says, 'Our report shows misinformation has become more professionalised and is now produced on an industrial scale. Now, more than ever, the public needs to be able to rely on trustworthy information about government policy and activity. Social media companies need to raise their game by increasing their efforts to flag misinformation and close fake accounts without the need for government intervention, so the public has access to high-quality information.'

> "Social media manipulation of public opinion is a growing threat to democracies around the world"

slido

Audience Q&A Session

ⓘ Start presenting to display the audience questions on this slide.