# CSC 495.002 – Lecture 2
# Web/Social Networks Privacy: Inference

Dr. Özgür Kafalı

North Carolina State University
Department of Computer Science

Fall 2017



## What You Will Learn

- How additional information about individuals can be inferred from known data
- Sharing behaviors of users and how shared content propagates in the network
- Common violations and regret scenarios
- Methods for targeted advertising and how to mitigate those
- K-anonymity for ensuring privacy of datasets

## Inference (Logic)

- The act or process of deriving logical conclusions from premises known or assumed to be true
- Example
  1. Fact: All humans are mortal.
  2. Fact: All Greeks are humans.
  3. Inference: All Greeks are mortal.





## Inference Example

- Fact: The guy in the picture wears a ring on his ring finger
- Inference rule: People who wear a ring on their ring finger are married
- Outcome: The guy in the picture is married

## Inference (Privacy)

- Can user attributes and social network links be used to predict attributes of additional users?
- Facts: $\text{User}_1$.interests, $\text{User}_2$.interests, ..., $\text{User}_{n-1}$.interests, friend($\text{User}_1$, $\text{User}_2$), ..., friend($\text{User}_1$, $\text{User}_n$)
- Infer: $\text{User}_n$.interests
- Abductive reasoning: Some of the facts are assumed to be true (abducted) unless any evidence to the contrary





## Related Problems

- Opposite problem?
- Make use of user attributes to predict social links

- Preparing datasets for research purposes (upcoming lectures)
- Inference attacks: Analyze data to gain knowledge about a subject or database (upcoming k-anonymity lecture)

## Machine Learning

- Probabilistic inference: Quite high accuracy with big amounts of data
- Ethics:
  - Invasiveness of emerging machine learning technologies on user privacy
  - Created with good-natured intent . . . until it goes wrong
  - Balance the predictive value of data models with fairness and ethical values
  - No accountability for wrong predictions
- Facial recognition: Identify protestors in rallies and demonstrations
- Criminal intent: Wrongly accusing people because of their demographics (race, residential area, . . . )





## Hiring

- New interview questions in addition to your CV
  - Based on social media profiles
  - Who your friends are on Facebook
  - Whom you follow on Twitter
  - What you like on Instagram
- “We’ll verify your CV via Facebook, Twitter, and Instagram, and get back to you”
- “Do you have any other people skills besides 500+ Facebook friends?”

https://www.pinterest.com/ppimidland/funny-interviews/

## You Are Who You Know: Inferring User Profiles in Online Social Networks

### You Are Who You Know: Inferring User Profiles in Online Social Networks

Alan Mislove[†‡§], Bimal Viswanath[†], Krishna P. Gummadi[†], Peter Druschel[†]
[†]MPI-SWS, Saarbrücken, Germany
[‡]Rice University, Houston, TX, USA
[§]Northeastern University, Boston, MA, USA
amislove@ccs.neu.edu, {bviswana, gummadi, druschel}@mpi-sws.org

**ABSTRACT**

Online social networks are now a popular way for users to connect, express themselves, and share content. Users in today's online social networks often post a profile, consisting of attributes like geographic location, interests, and schools attended. Such profile information is used on the sites as a basis for grouping users, for sharing content, and for suggesting users who may benefit from interaction. However, in practice, not all users provide these attributes.

In this paper, we ask the question: given attributes for some fraction of the users in an online social network, can we *infer* the attributes of the remaining users? In other words, can the attributes of users, in combination with the social network graph, be used to predict the attributes of another user in the network? To answer this question, we gather fine-grained data from two social networks and try to infer user profile attributes. We find that users with common attributes are more likely to be friends and often form dense communities, and we propose a method of inferring user attributes that is inspired by previous approaches to detecting communities in social networks. Our results show that certain user attributes can be inferred with high accuracy when given information on as little as 20% of the users.

ple, MySpace (over 275 million users)[1], Facebook (over 300 million users), Orkut (over 67 million users), and LinkedIn (over 50 million "professionals") are examples of wildly popular networks used to find and organize contacts. Some networks such as Flickr, YouTube, and Picasa are used to share multimedia content, and others like LiveJournal and BlogSpot are popular networks for sharing blogs.

Users often post profiles to today's online social networks, consisting of *attributes* like geographic location, interests, and schools attended. Such profile information is used as a basis for grouping users, for sharing content, and for recommending or introducing people who would likely benefit from direct interaction. Today's online social networks rely on users to manually input profile attributes, representing a significant burden on users, especially when users are members of multiple online social networks. Thus, in practice, not all users provide these attributes, thereby reducing the usefulness of the social networking applications.

In this paper, we ask the question: is it possible to *infer* the missing attributes of a user in an online social network from the attribute information provided by other users in the network? In other words, can the attributes of other users in the network, in combination with the social network graph, be used to predict those of a given user? In offline social

---

Mislove et al. You Are Who You Know: Inferring User Profiles in Online Social Networks. Conference on Web Search and Data Mining, pages 251–260, 2010





## Research Questions

- Is it possible to infer missing attributes of a user in an online social network (OSN) from other users' attributes and their relations with the user in subject?
- What user attributes and social links are necessary to infer another user's attributes?

## Motivation and Assumptions

- Attributes for grouping users:
  - Geographic location
  - Interests
  - Schools attended
- Homophily: People tend to befriend others who share similar traits
  - Support same sports clubs
  - Date similar people
  - Offline behavior observed in the online world too





## Implications

- A user's privacy no longer depends on what they reveal (who you're friends with also reveals information about you)
- Not all bad implications though
- Better recommender systems
- Connect people that might benefit from the interaction (e.g., job search)
- Reduce burden on users by avoiding manual data entry
- Data analysis: Filling in missing data helps produce better quality results from aggregated data

## Datasets

- Dataset 1: Rice University
- 4,000 students and alumni of Rice University collected from Facebook
- Attributes collected:
  - Major(s) of study
  - Year of matriculation
  - Dormitory
- Dataset 2: New Orleans
- 63,000 users in the New Orleans Facebook regional network
- Attributes collected from Facebook profile page





## Datasets: Rice University

- Crawl Rice University Facebook network (students and alumni)
- Name and list of friends collected
- Additional information collected from Rice University Student and Alumni Directories
- Matriculation year, graduation year, residential college, and major(s) collected
- Subsets used: Undergraduates (1,220 users with 43,208 links) and graduates (501 users with 3,255 links)
- Very few links between the two subsets

## Datasets: New Orleans

- Crawl New Orleans Facebook regional network
- Collect user profiles (some are private)
- 63,731 users





## Exercise: Revealed Attributes in New Orleans Network

- Employer, high school, interests, location, university

| Attribute | Fraction revealed |
|---|---|
| | 68.9% |
| | 58.3% |
| | 42.3% |
| | 35.5% |
| | 19.3% |

## Revealed Attributes in New Orleans Network

| Attribute | Fraction revealed |
|---|---|
| high school | 68.9% |
| university | 58.3% |
| employer | 42.3% |
| interests | 35.5% |
| location | 19.3% |





## Methodology: Social Network as a Graph

- G = (V, E)
- Users are nodes (or vertices V)
- Friend links are edges (E)

## Attribute Commonality and Affinity

- Fraction of links for which users share the same value of attribute a
- $S_a = \frac{|(i,j) \in E : a_i = a_j|}{|E|}$
- Divide that by $E_a$: expected if attributes were placed randomly
- Affinity = $S_a$ / $E_a$
- Values greater than 1 indicate that links are positively correlated with attributes



NC STATE UNIVERSITY
TECHNIQUES & STUDIES

## Exercise: Affinity Values

| Users | Attribute | Affinity |
|---|---|---|
| Rice undergrads | college<br>major<br>year | |
| Rice grads | department<br>school<br>year | |
| New Orleans | high school<br>hometown<br>political views | |

## Affinity Values

| Users | Attribute | Affinity |
|---|---|---|
| Rice undergrads | college | 4.49 |
| | major | 2.33 |
| | year | 1.97 |
| Rice grads | department | 9.71 |
| | school | 4.02 |
| | year | 1.79 |
| New Orleans | high school | 53.2 |
| | hometown | 2.87 |
| | political views | 1.86 |





## Communities

- Users who share common attributes form communities (dense clusters in the network)
- Quantify the strength of communities using modularity
- Modularity: Positive values indicate strong community structure

## Modularity

- Rice undergraduates

| Attributes | Communities | Modularity |
|---|---|---|
| college, major, year | 582 | 0.023 |
| college, major | 317 | 0.029 |
| year, major | 147 | 0.045 |
| major | 52 | 0.055 |
| college, year | 44 | 0.248 |
| year | 7 | 0.259 |
| college | 9 | 0.384 |





## Methodology: Global Community detection

- Global detection: Assume knowledge of the entire network
- Remove edges until network graph is partitioned
- Various partitions are communities
- Betweenness centrality: Tells which edge to remove
  - Bridging links between communities
  - Once removed, underlying community structure emerges
- Limitations: Computationally expensive, hard to obtain complete network structure

# Community Building from Rice Undergraduates

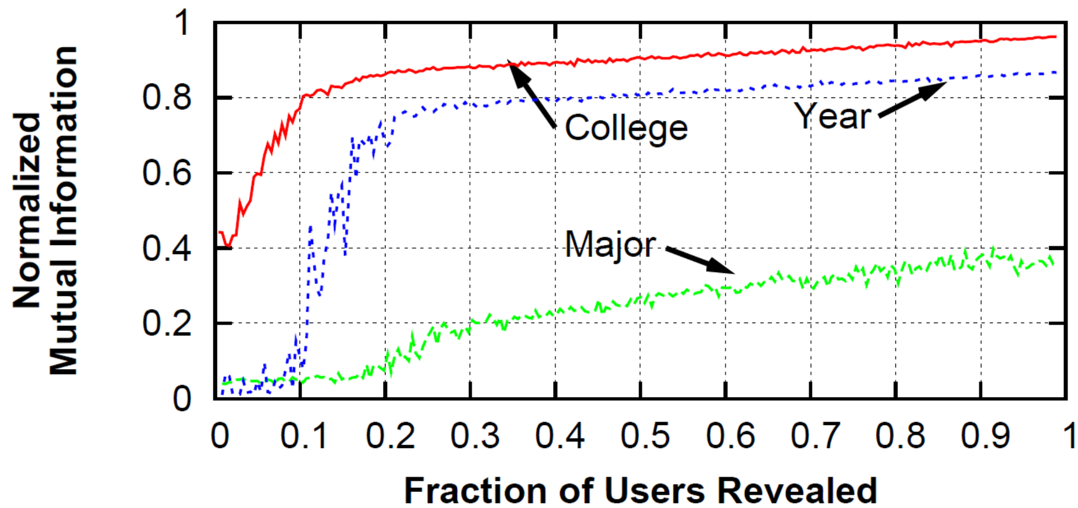






# Pareto Principle

- Named after Italian economist Vilfredo Pareto
- For many events, roughly 80% of the effects come from 20% of the causes
- 80% of the land in Italy was owned by 20% of the population
- 20% of the peapods in his garden contained 80% of the peas
- Software: 80% of errors eliminated by fixing 20% of bugs
- Can be applied to anything . . .

## Methodology: Local Community detection

- Local detection: Assume knowledge of a local region
- Scalable, applicable to larger networks
- Start with a set of seed nodes
- Add neighbors until sufficiently strong community is found
- Outwardness: Neighbors outside community - neighbors within community
- Add the node with lowest outwardness





## Inferring Attributes: Conductance

- Measure the quality of a community
- Let A, B subsets of V (e.g., users sharing common attributes)
- $B = V \setminus A$
- $e_{AB}$: Number of edges between A and B
- $e_{AA}$: Number of edges within A
- Traditional conductance of A: $e_{AB}$ / $e_{AA}$
- Small value means stronger community
- Biased towards large communities

## Normalized Conductance

- K = $\frac{e_{AA}}{e_{AA}+e_{AB}}$ (value close to 1 indicates good community structure)
- Also biased towards very large communities

- Normalized conductance: K - expected value of K for a random network
  - Strongly positive values indicate good community structure
  - Zero indicates random graph
  - Negative values indicate less structure than random graph





## Community Detection Algorithm

- Input: A subset (S) of users in a community
- Output: Other members of the community
- Greedy approach: Maximize C (normalized conductance)
- Divide network into A and B, initially A = S
- At each step, select user v $\in$ B
- Adding v to A yields the highest increase in C for A
- Stop when no more such users are found

# Evaluation Metrics

- H: Users in community
- S: Initial subset of users in H (before algorithm)
- R: Additional users (believed to be in H) returned by algorithm

- Precision: Fraction of returned users who are actually in the community
  $| R \cap H | / | R |$

- Recall: Fraction of remaining community members returned
  $| R \cap H | / | H \setminus S |$





# Results

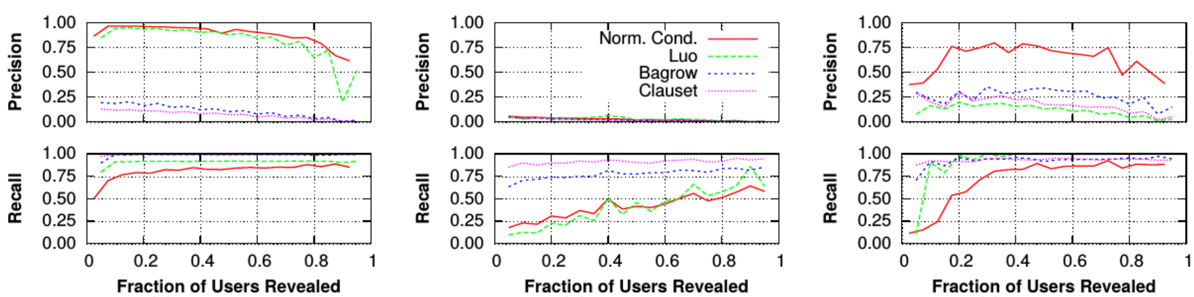


(a) College (b) Major (c) Year