You are who you know: Inferring user profiles in online social networks

Alan Mislove[†‡§] Bimal Viswanath[†] Krishna P. Gummadi[†] Peter Druschel[†]

[§] *Northeastern University* [†]*MPI-SWS* [‡]*Rice University*

February 5th, 2010, WSDM

Facebook and personal data

Users upload information to sites like Facebook
- Profile information
- Status updates
- Photos, videos

Privacy model for data
- Choose what to reveal
- And what to keep private

When reasoning about privacy
- Don't often consider *implicit* data
- What our friends reveal about ourselves

facebook

+

What is implicit data?

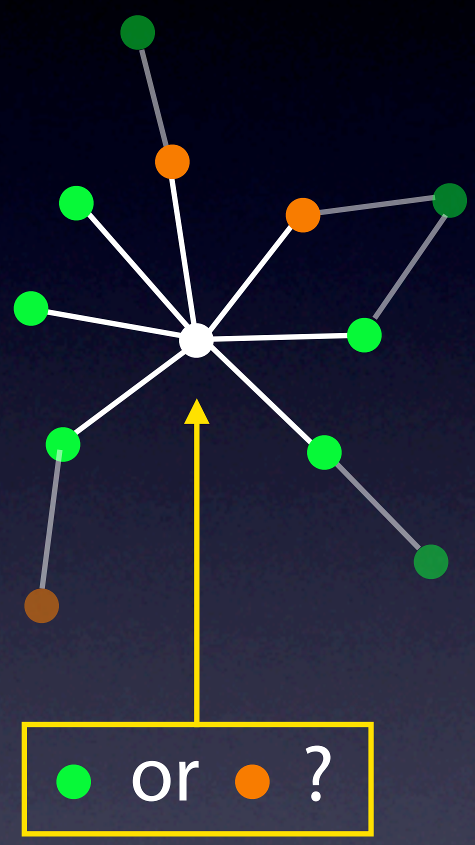

Example: MIT's Project Gaydar

Predict sexual orientation based on friends

Exploiting *homophily*

People associate with others like them

What about other attributes?

Using friends-of-friends?

This talk

Explore *how much implicit data exists on online social networks?*

- Or, how much information can be inferred?
- How much data is needed to be able to infer?

Focus on one source: social network

Develop methodology to infer user attributes

- Test on real-world network data

Roadmap

1. Idea: Use communities to infer attributes
2. Collect fine-grained community data
3. Do attribute-based communities exist?
4. How well can we infer user attributes?

Idea: Use communities

Project Gaydar used 1-hop friends

Using >1 hop friends is challenging

- Exponential growth in size
- Unclear relationship to source

Look for groupings of users

- Called *communities*
- Potentially share attributes

Leverage literature in community detection

What do we mean by *communities*?

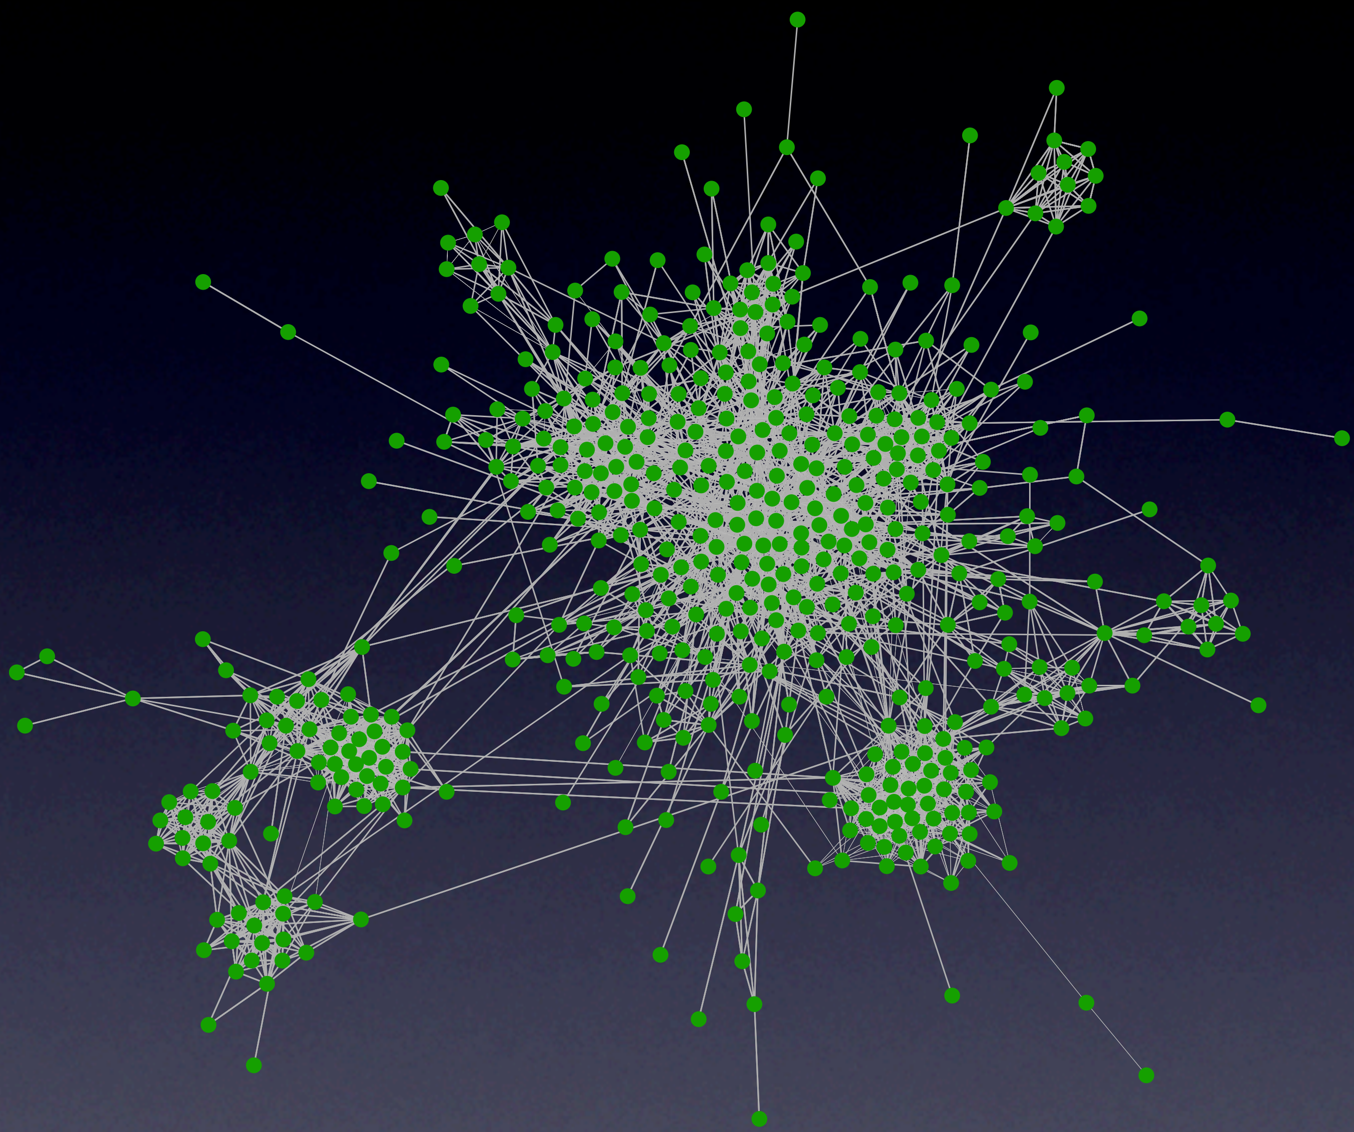

Group: Users who share a common attribute

Community: Users more densely connected than overall graph

What do we mean by *communities*?

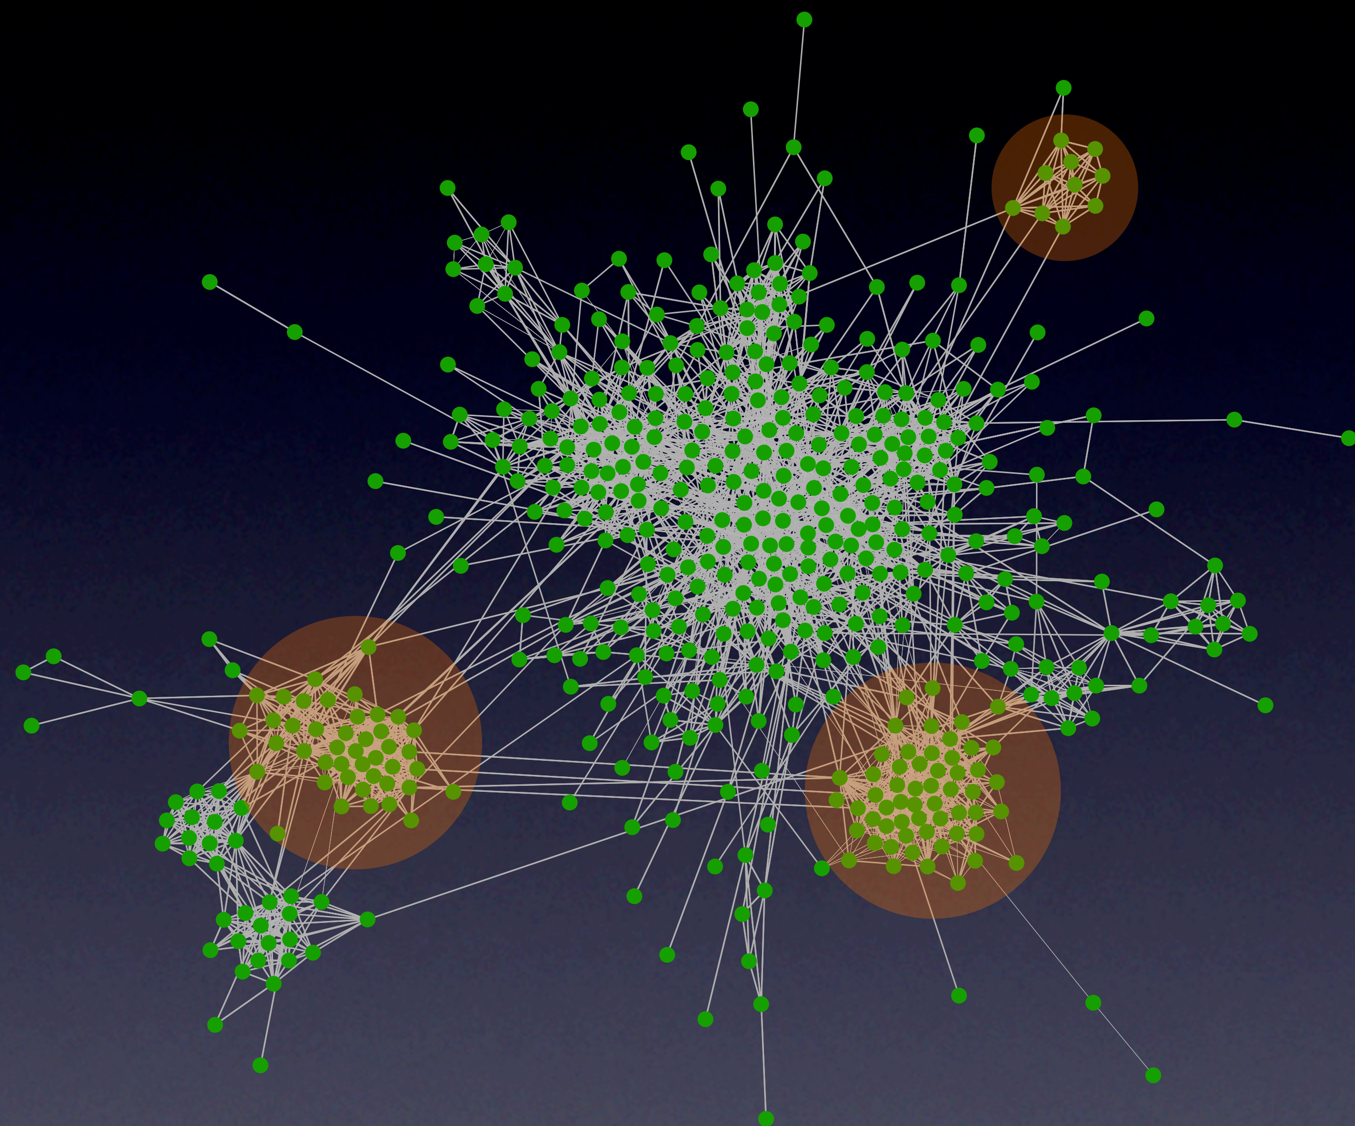

Group: Users who share a common attribute

Community: Users more densely connected than overall graph

Roadmap

1. ~~Idea: Use communities to infer attributes~~
2. Collect fine-grained community data
3. Do attribute-based communities exist?
4. How well can we infer user attributes?

Social network data

Crawled two Facebook networks

- Rice University (university)
- New Orleans (regional)

	Users	Avg. Degree
Rice ugrad		
Rice grad		
New Orleans		

Picked known seed user

- Crawled all of his friends, added new users to list
- Effectively performed a BFS of graph

Social network data

Crawled two Facebook networks

- Rice University (university)
- New Orleans (regional)

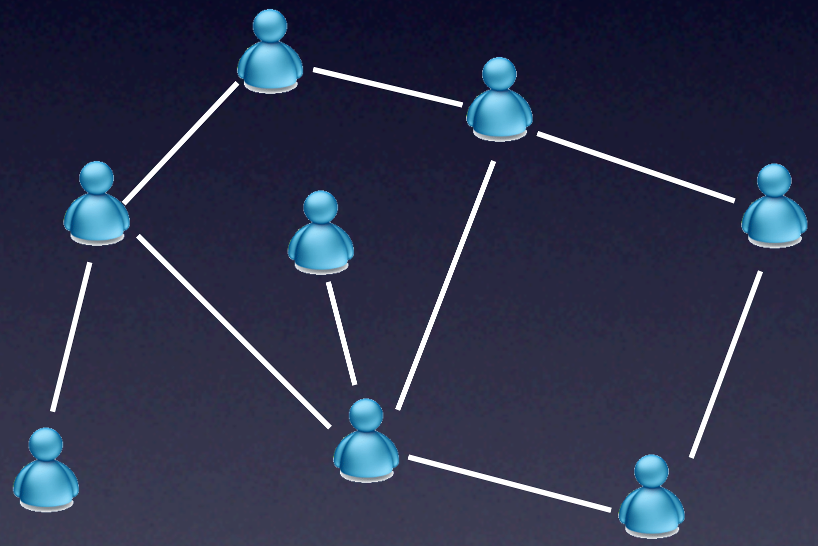

	Users	Avg. Degree
Rice ugrad		
Rice grad		
New Orleans		

Picked known seed user

- Crawled all of his friends, added new users to list
- Effectively performed a BFS of graph

Social network data

Crawled two Facebook networks

- Rice University (university)
- New Orleans (regional)

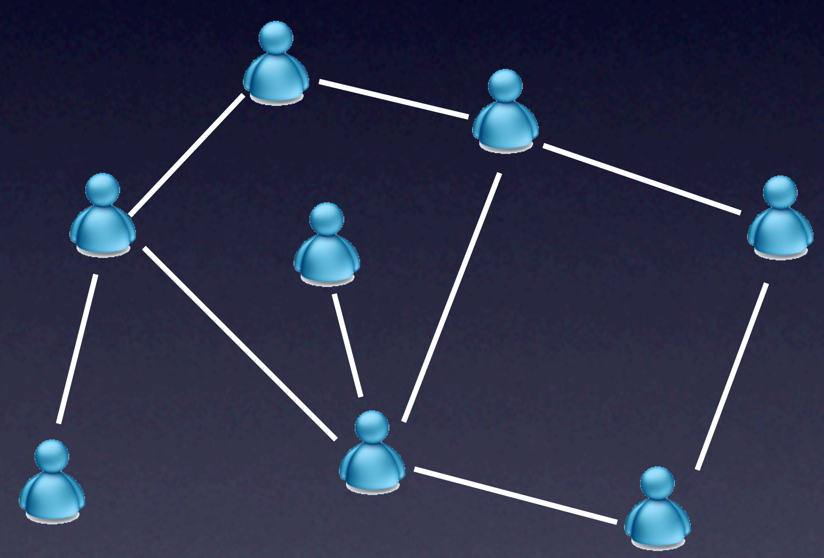

	Users	Avg. Degree
Rice ugrad	1,220	35.4
Rice grad	501	6.5
New Orleans	63,731	24.2

Picked known seed user

- Crawled all of his friends, added new users to list
- Effectively performed a BFS of graph

Collecting attributes

Obtained authoritative information

- Queried student directory
- College (dormitory), major(s), year

Could not collect Facebook profiles

Collected Facebook profiles

Extracted all attributes

- E.g., high school, groups, gender
- Attributes are freeform text

Roadmap

1. ~~Idea: Use communities to infer attributes~~
2. ~~Collect fine-grained community data~~
3. Do attribute-based communities exist?
4. How well can we infer user attributes?

Do attributes define communities?

Put users into groups based on attributes

- Determine if these are communities

Need metric to rate communities

Modularity rates community strength

- Range [-1,1]
- 0 represents expected in random graph
- ≥0.25 represents community structure

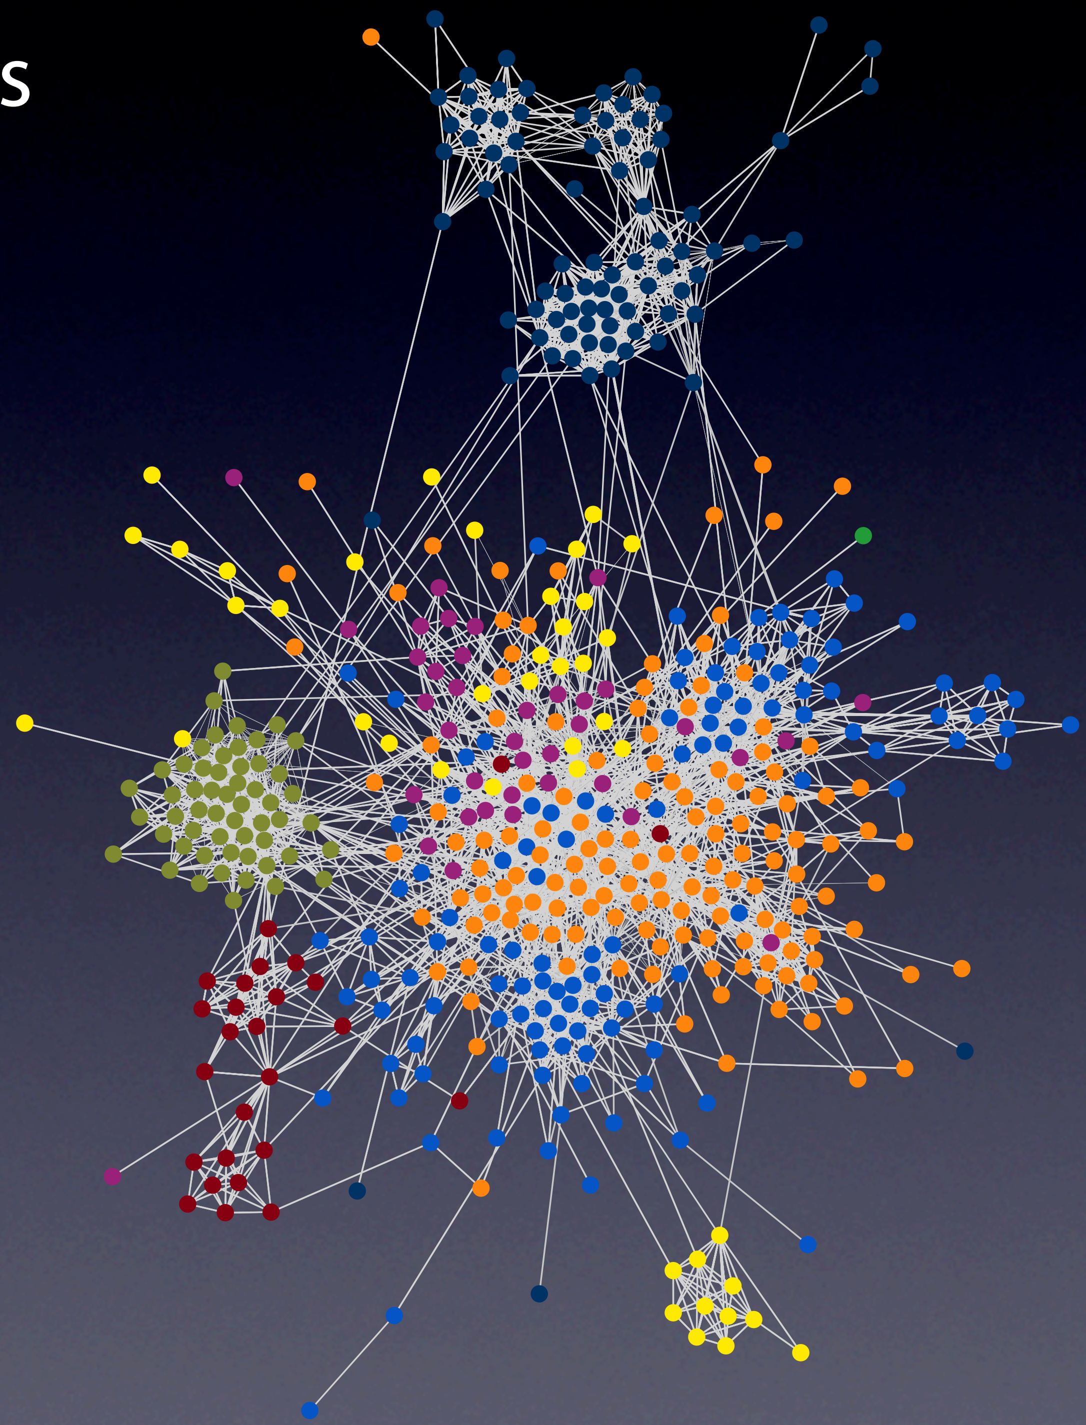

Attribute communities for Rice ugrads

	Communities	Community Size			Modularity
		Min	Avg	Max	
major	65	1	23	105	0.004
matriculation year	4	95	305	398	0.259
residential college	9	130	135	142	0.385

Communities based on shared college or year

- Multiple, overlapping community structures

Roadmap

1. ~~Idea: Use communities to infer attributes~~
2. ~~Collect fine-grained community data~~
3. ~~Do attribute-based communities exist?~~
4. How well can we infer user attributes?

Using communities to infer attributes

Can we detect a single attribute community?

- Given a few users in the community

Previous approaches proposed (local community detection)

- Not designed for social networks
- Never evaluated on a large-scale social network

Propose a new algorithm to detect a specific community

- Problem: How to evaluate community *strength*?

Normalized Conductance

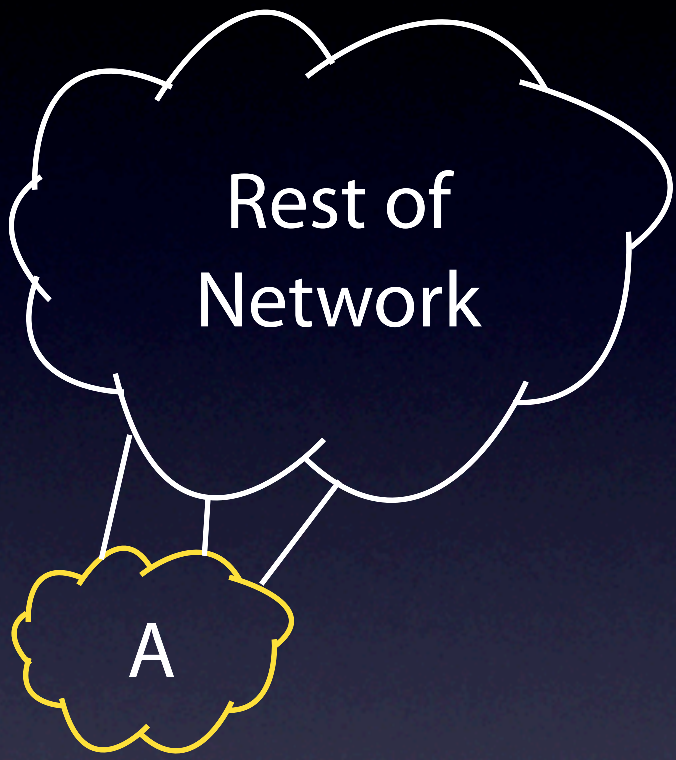

$$C = \frac{e_{AA}}{e_{AA} + e_{AB}} - \frac{e_A e_A}{e_A e_A + e_A e_B}$$

How *strong* is a particular community *A*?

Conductance previously proposed
- But, biased towards large communities

Metric: Normalized conductance C
- Fraction of *A*'s links within *A*
- Relative to a random graph

Range is [-1,1]
- 0 represents no stronger than random

Algorithm

Given seed users, find a community by

- Adding users
- Stopping at some point

At each step, add user who increases normalized conductance by the most

Stop when no user increases normalized conductance

How to evaluate?

Evaluate performance using *precision* and *recall*

Algorithm takes in fraction sharing attribute

recall = fraction of remaining attribute-sharing users identified

precision = fraction of identified users that share attribute

Ideally want a precision and recall of 1.0

Can we infer Rice undergrad classes?

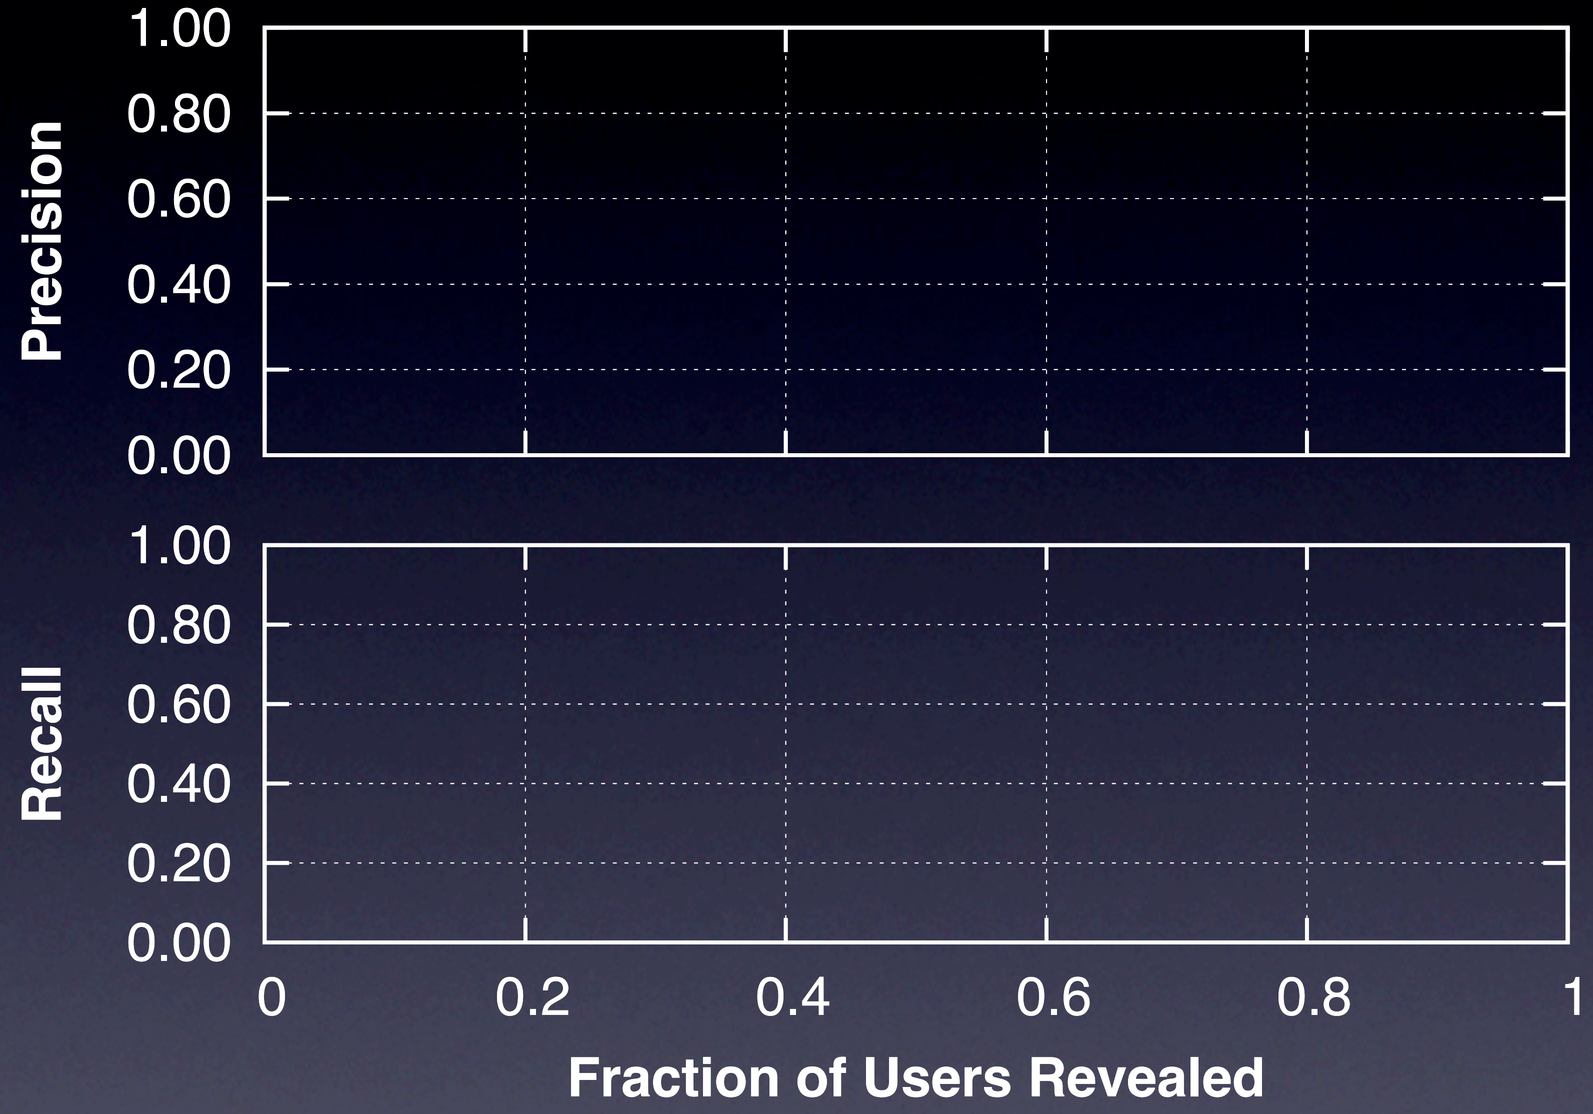

Yes; different communities show different characteristics

- In next graphs, average across all groups

Can we infer Rice undergrad classes?

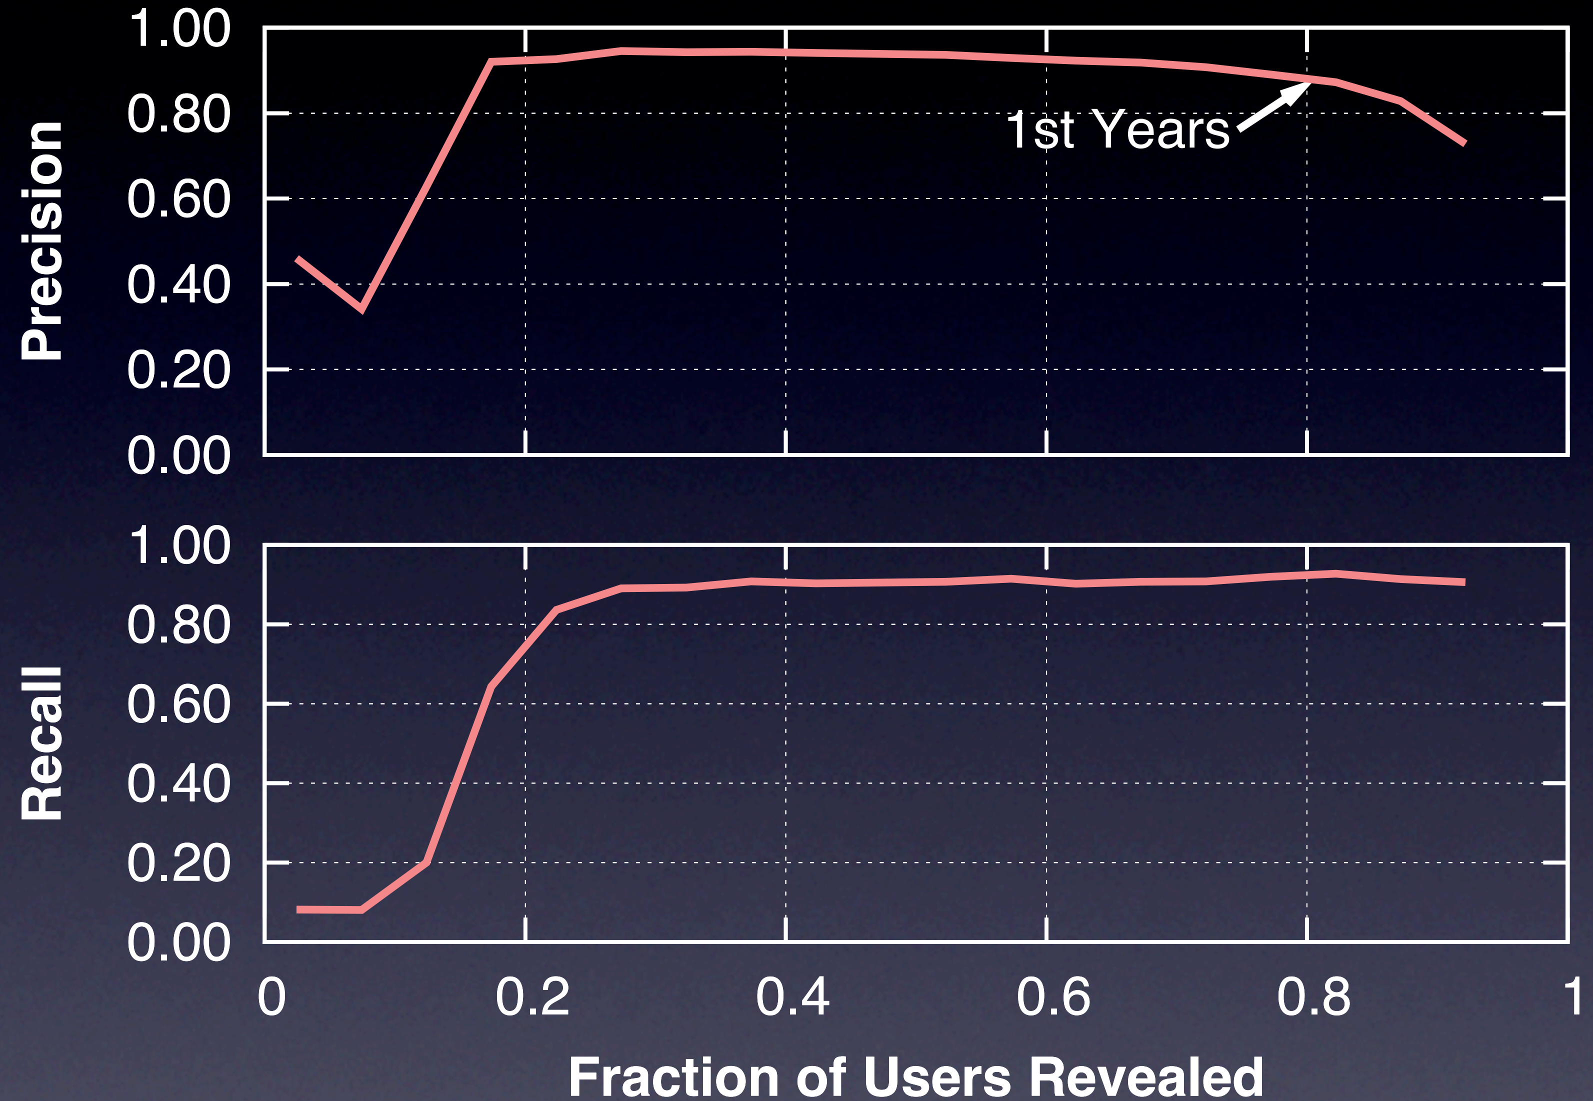

Yes; different communities show different characteristics

In next graphs, average across all groups

Can we infer Rice undergrad classes?

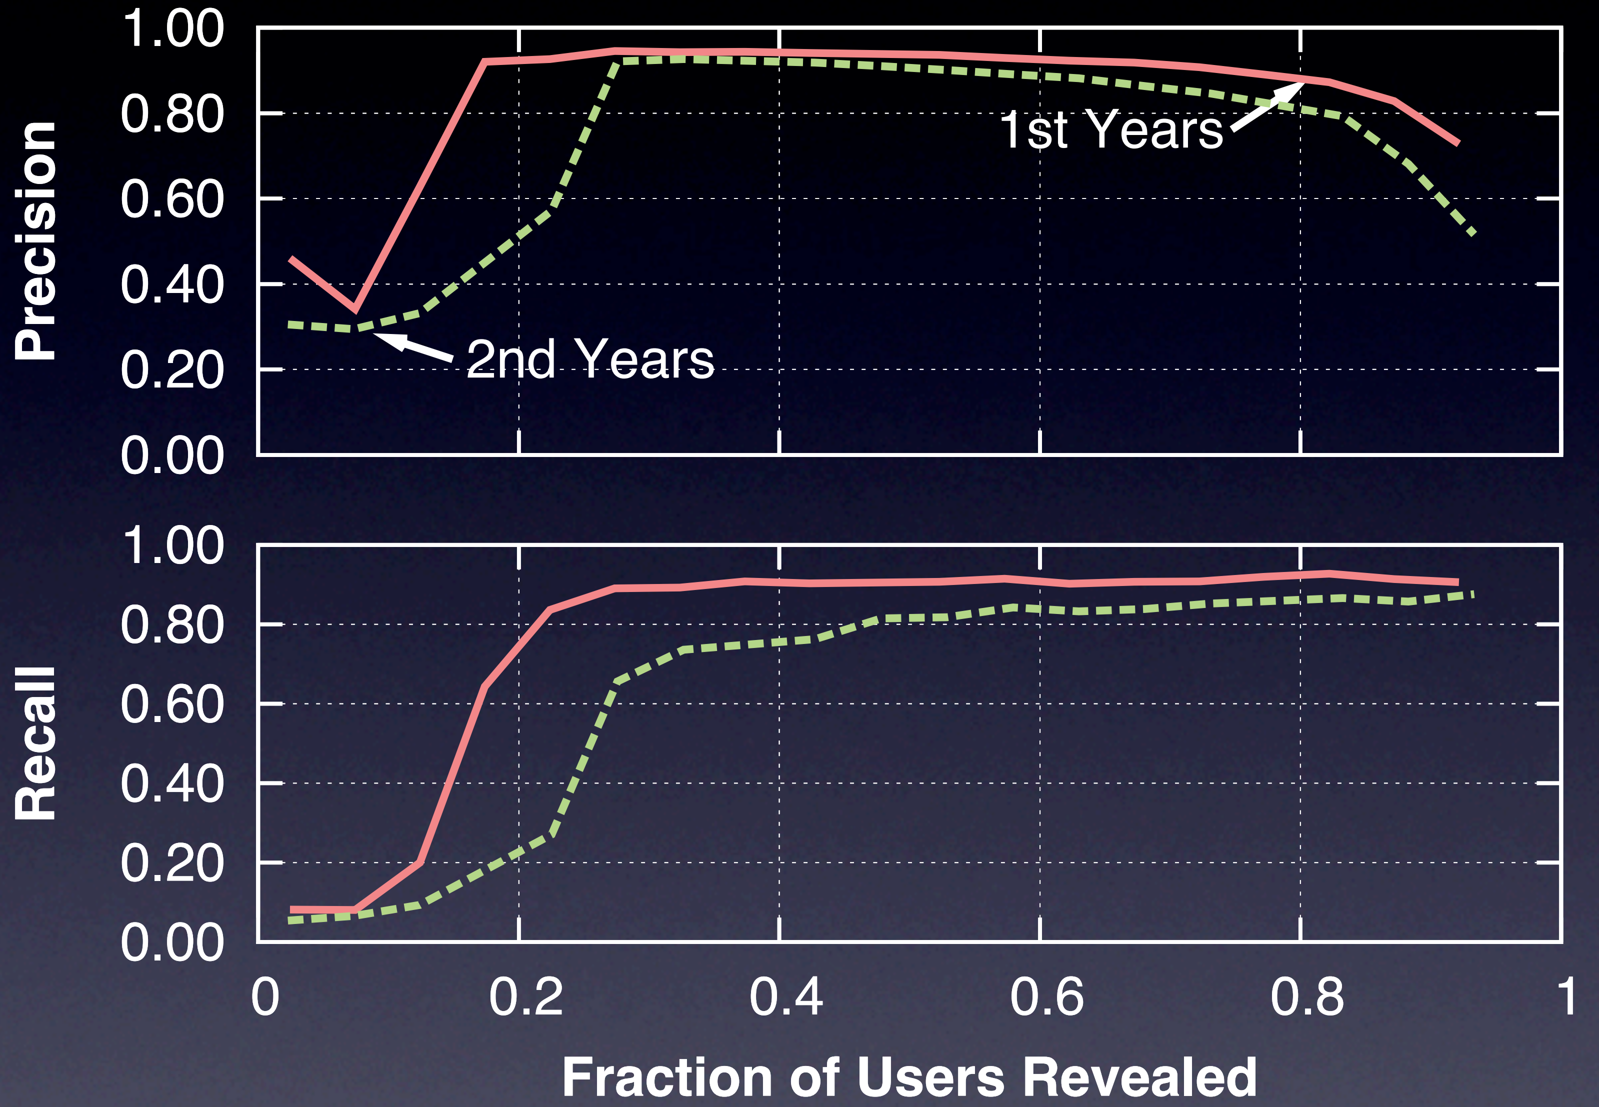

Yes; different communities show different characteristics

- In next graphs, average across all groups

Can we infer Rice undergrad classes?

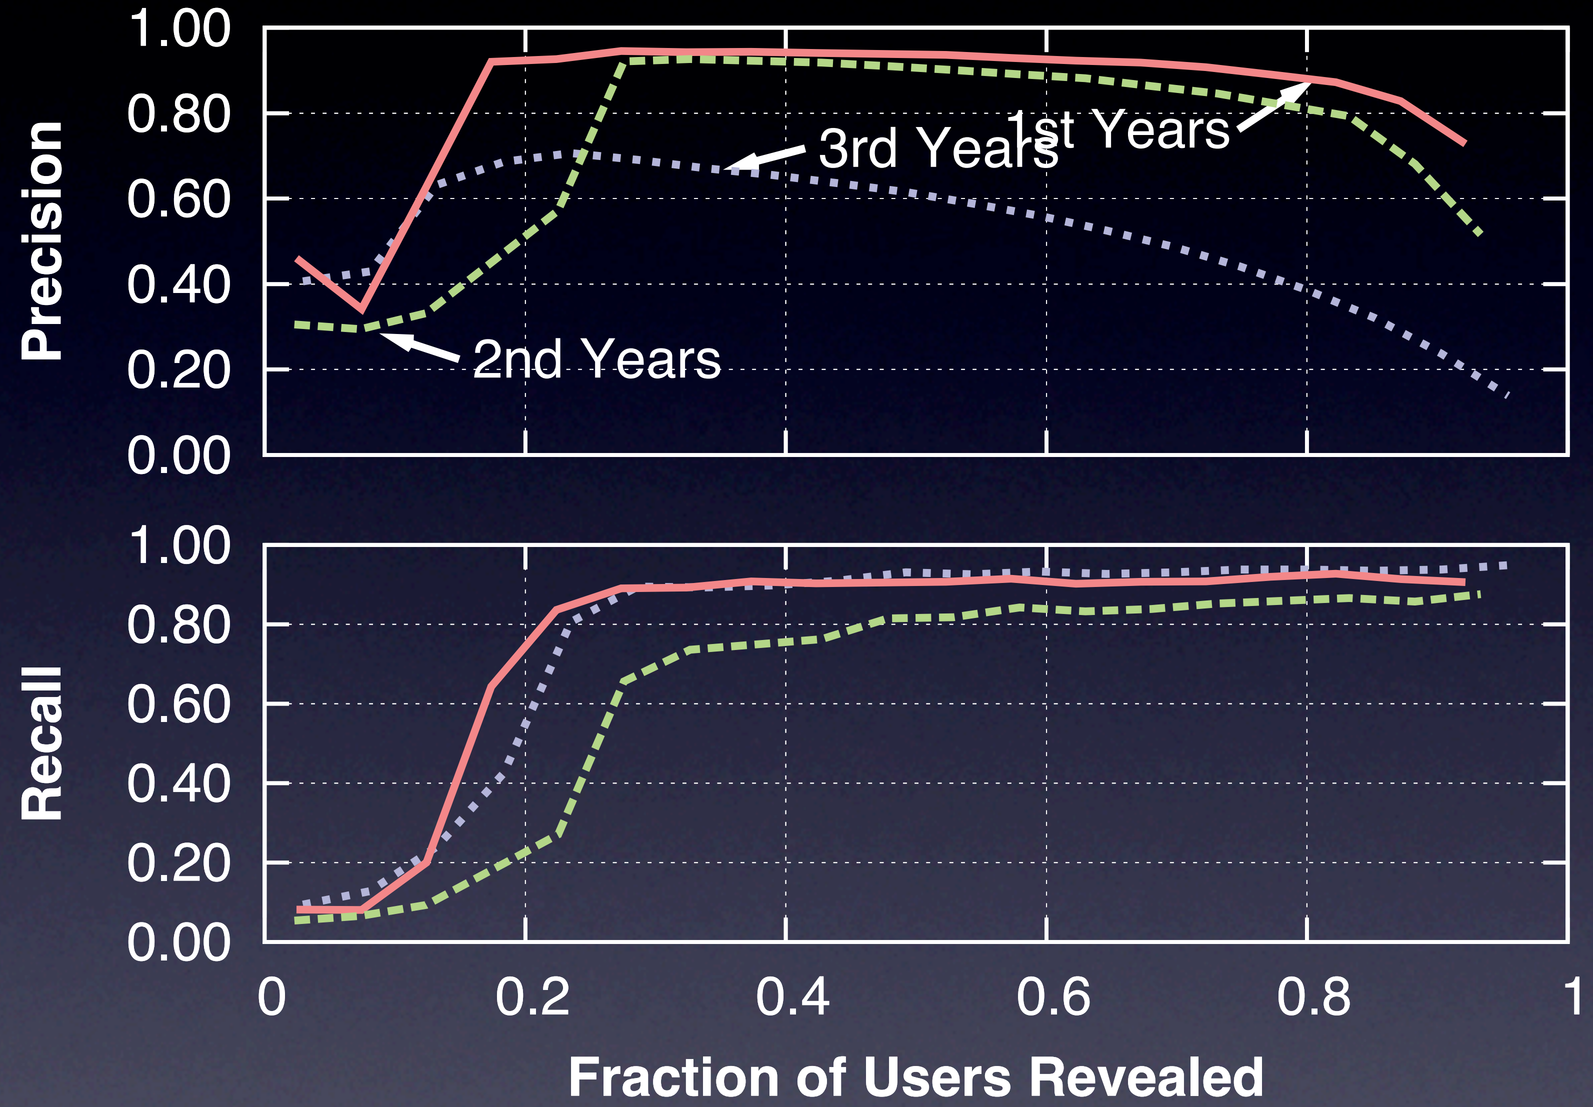

Yes; different communities show different characteristics

In next graphs, average across all groups

Can we infer Rice undergrad classes?

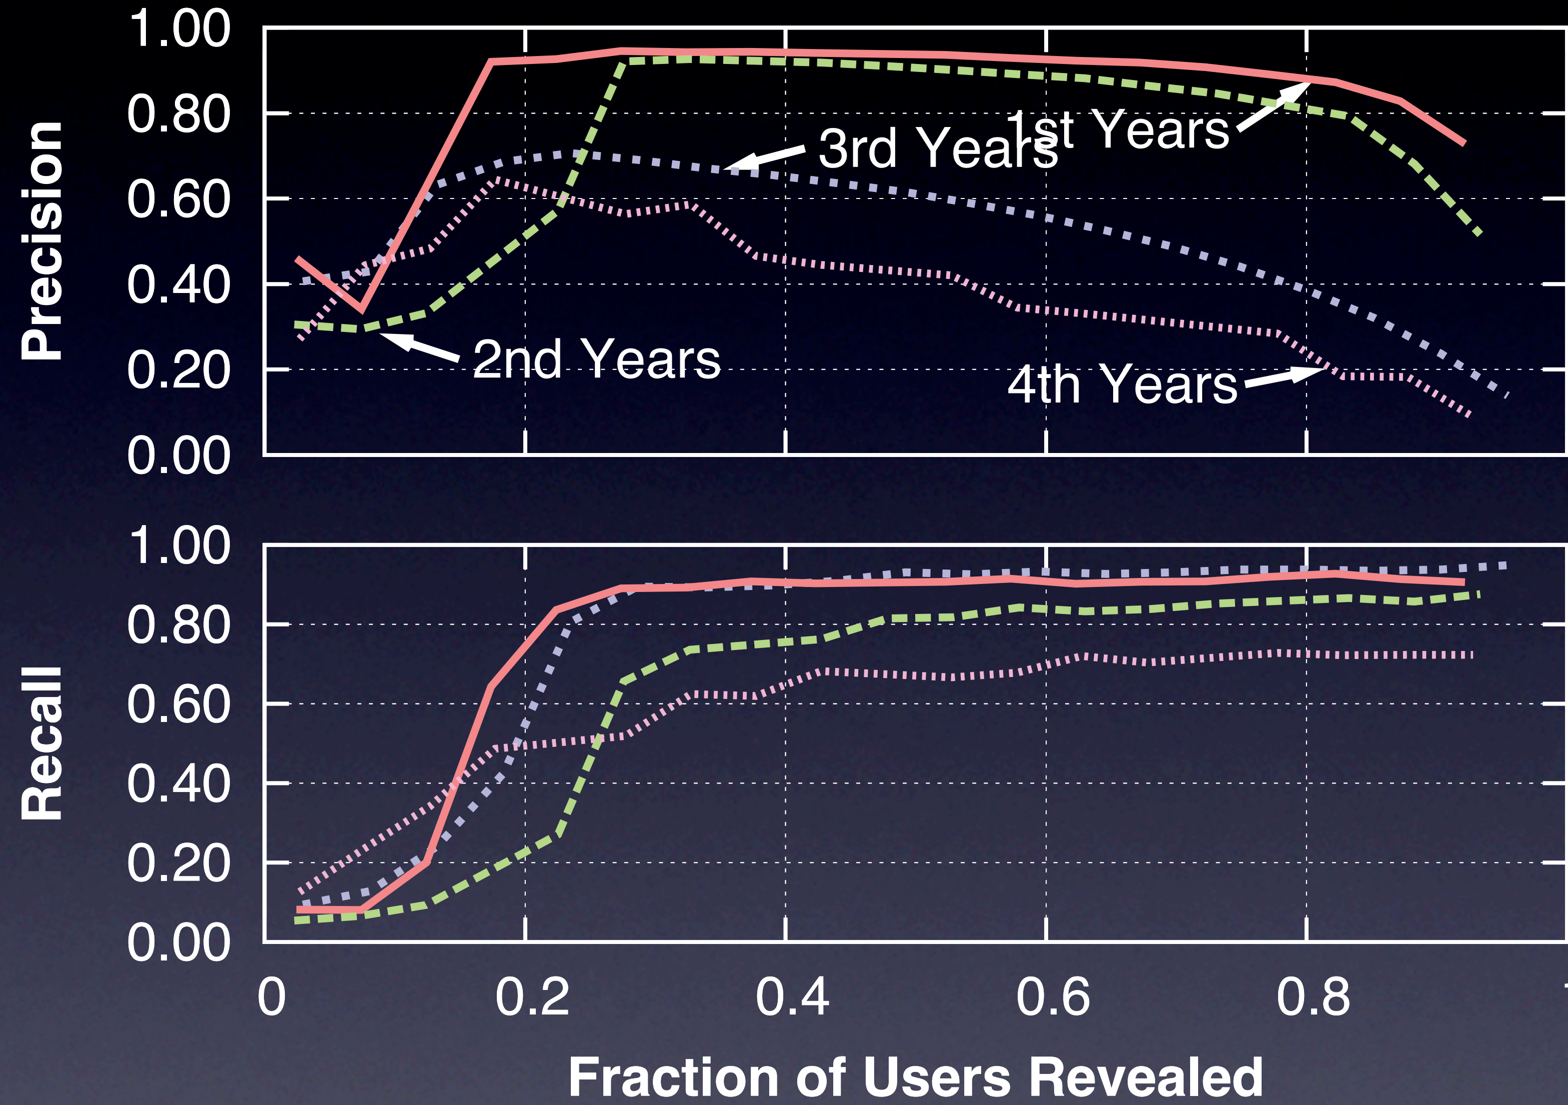

Yes; different communities show different characteristics

In next graphs, average across all groups

Inferring other attributes

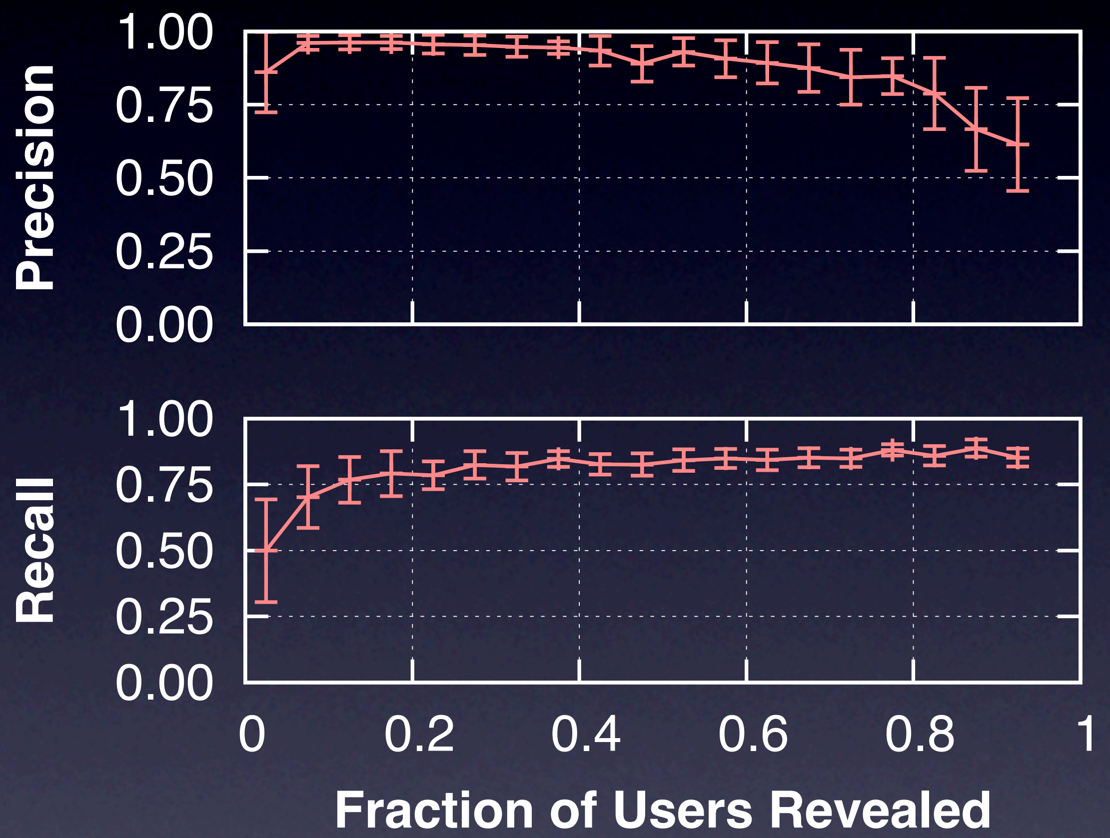

Dormitory

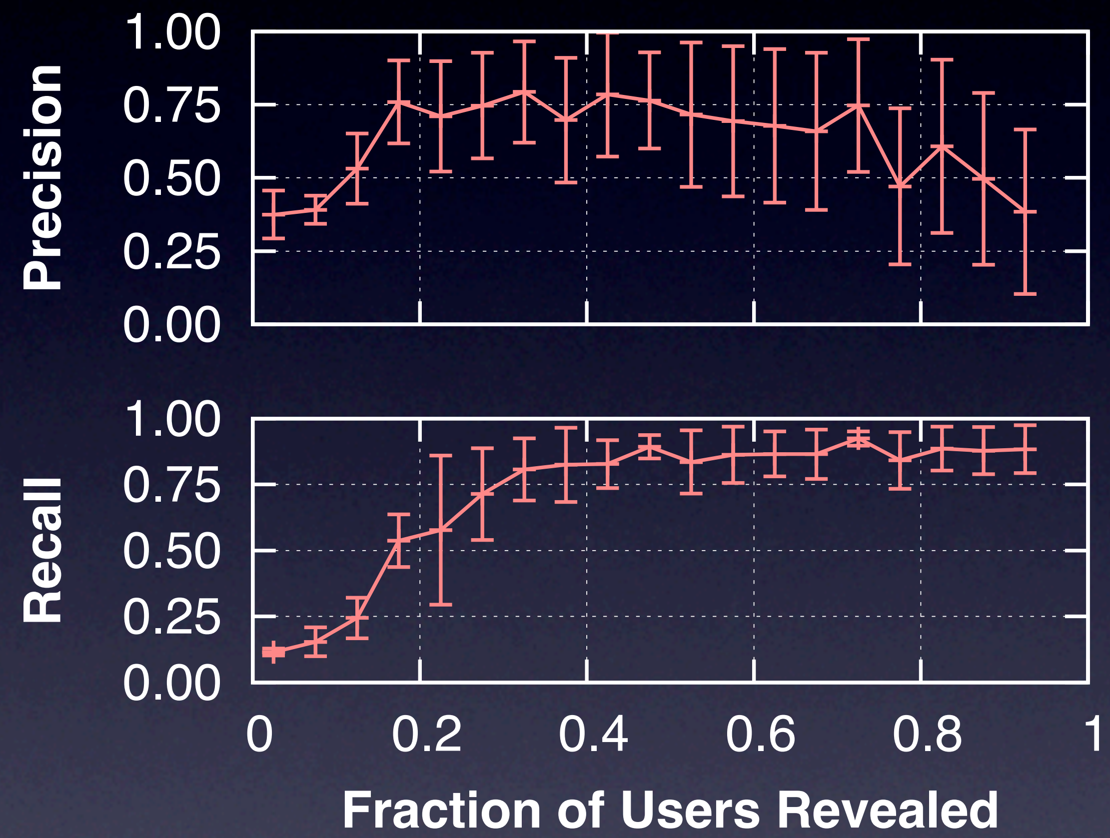

Matriculation year

Inferring other attributes

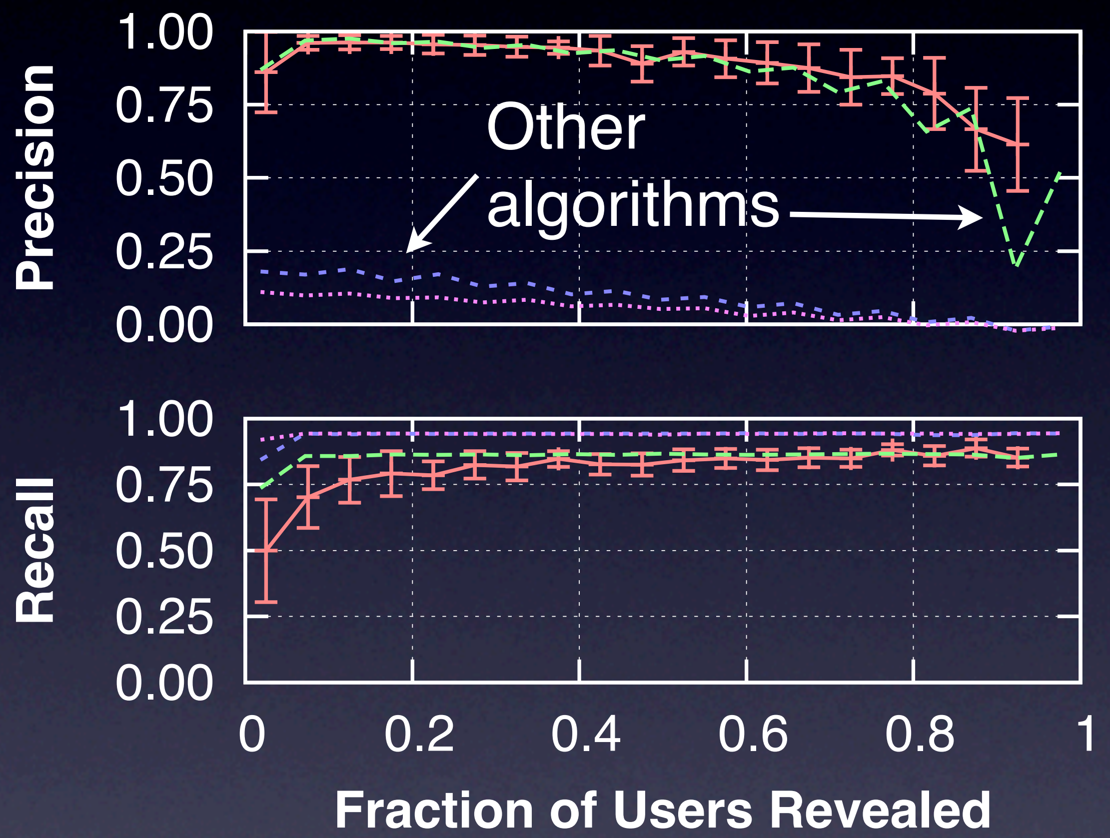

Dormitory

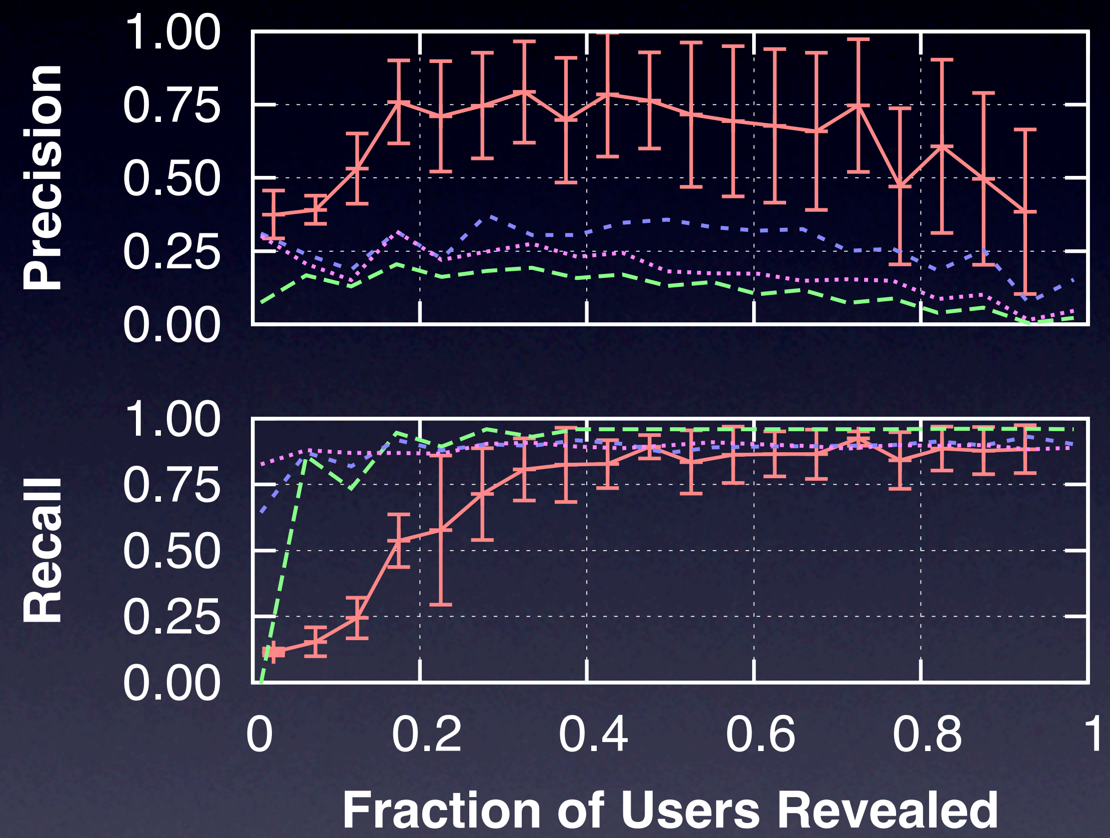

Matriculation year

Can we infer user-provided attributes?

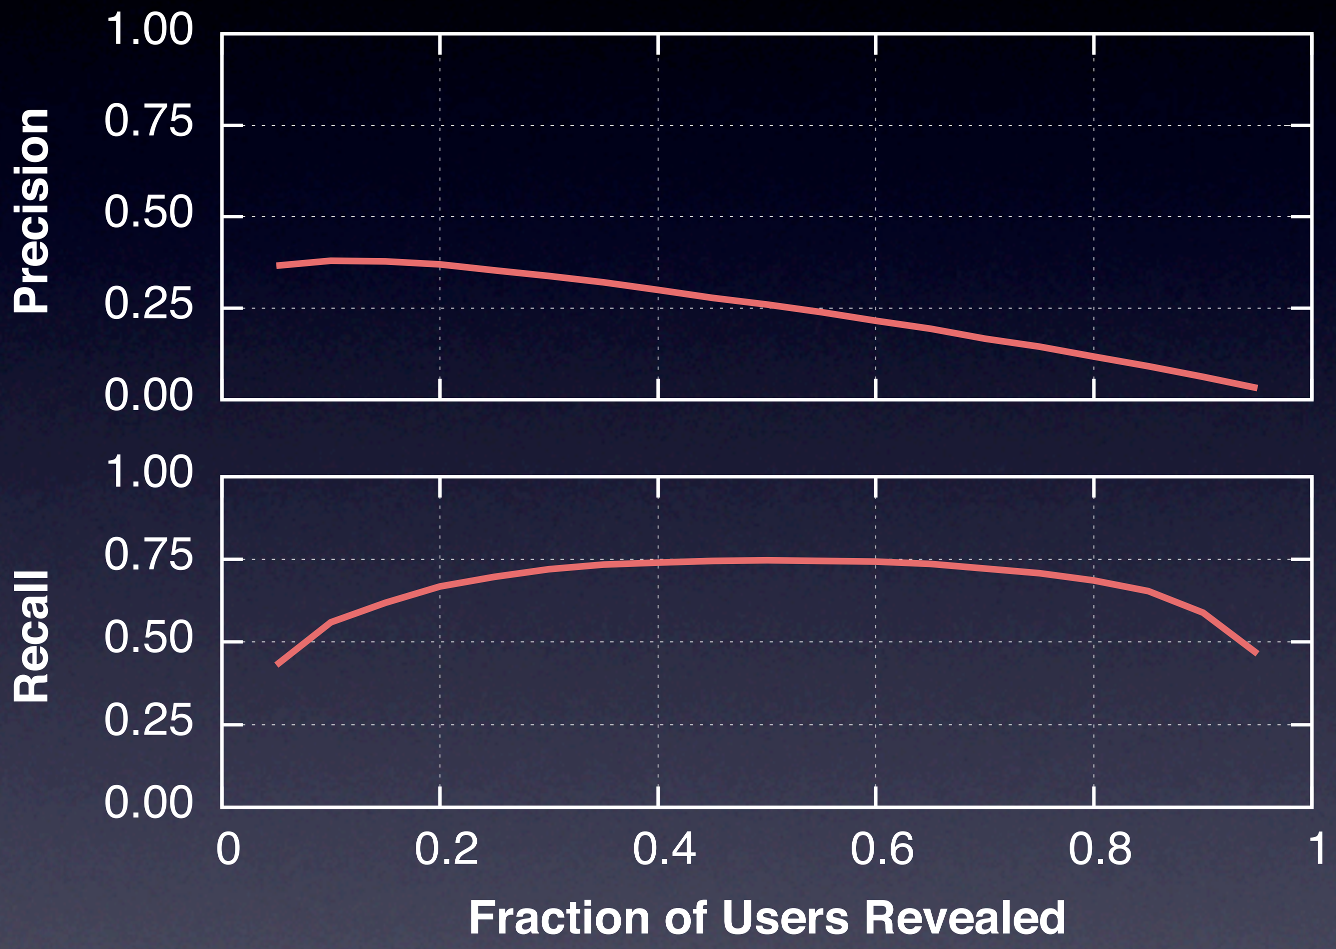

Use New Orleans data

Much more challenging
- Freeform text
- Non-authoritative attributes
- Missing data
- Most not communities
  - (gender, birthday, etc)

Results for 92 groups
- With conductance > 0.2

Summary

Ongoing online social network privacy debate

- Focuses mainly on *explicitly provided* attributes

Demonstrated that many attributes can be inferred

- Even if user didn't provide them

Good interpretation: Can reduce burden on users

- Don't have to fill in entire profile

Bad interpretation: Can figure out attributes users don't reveal

- Privacy is a function of what friends reveal

Questions?

Backup slides

Facebook privacy debate

Debate over privacy model and defaults

- Who can see users' attributes, status, friends

Scale, intensity of debate illustrates importance

So far, focused on explicit data

- Things the user uploaded or provided

What about *implicit* data?

- Data users didn't *explicitly* reveal?

facebook

+

Obtaining authoritative information

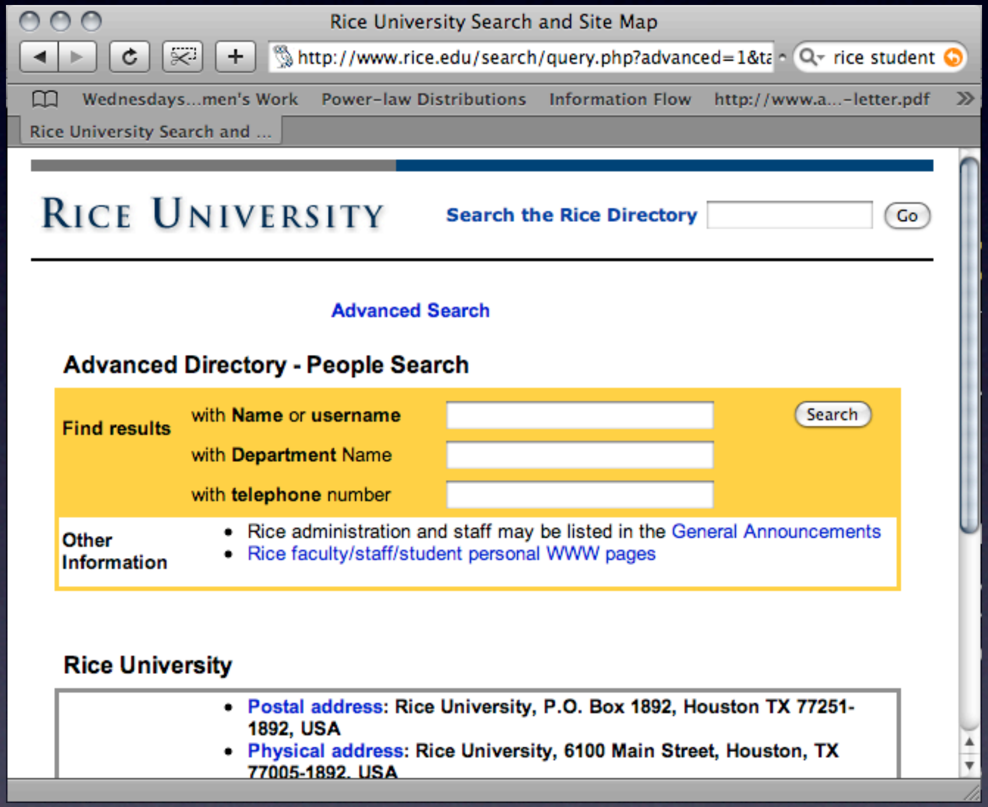

Additional information from student directory and alumni directory

Found matches for

- 1,233 (20.0%) undergraduates
- 548 (8.9%) graduate students
- 2,093 (33.9%) alumni

Focus on undergraduate network

- Obtained college, major(s), year

Similar results for others

Modularity

$$Q = \sum_i (e_{ii} - a_i^2)$$
$$= \mathrm{Tr}\,\mathbf{e} - ||\mathbf{e}^2||$$

How *good* is a community division?

Metric: Modularity *Q*

- Fraction of links within communities
- Relative to a random graph

Range is [-1,1]

- 0 represents no more community structure than random

Modularity > 0.25 indicates strong communities

Modularity

$$Q = \sum_i (e_{ii} - a_i^2)$$
$$= \mathrm{Tr}\ \mathbf{e} - ||\mathbf{e}^2||$$

How *good* is a community division?

Metric: Modularity *Q*

- Fraction of links within communities
- Relative to a random graph

Range is [-1,1]

- 0 represents no more community structure than random

Modularity > 0.25 indicates strong communities