Tutorial at the ACM SIGKDD conference, 2011
http://snap.stanford.edu/proj/socmedia-kdd

Social Media Analytics : Part 2: Rich Interactions

Jure Leskovec
Stanford University

Social Media: Interactions

- In Social Media users interact with one another and the content they both crate and consume
- Traditional social network analysis only distinguishes between pairs of people that are linked vs. not-linked
- But, user interactions in social media are much richer

Part 2 of the Tutorial: Outline

- How to learn to recommend/ predict links in social networks?
- User interactions in social media:
  - Strength: strong vs. weak ties
  - Friends vs. Foes
  - Trust vs. Distrust
- How people evaluate one another and the content that is being produced by others?

Tutorial Outline

- Part 1: Information flow in networks
- **Part 2: Rich interactions**
  - 2.1: Recommending links in networks
  - 2.2: Predicting tie strength
  - 2.3: Predicting friends vs. foes
  - 2.4: How do people evaluate others?

Tutorial Outline

- Part 1: Information flow in networks
- **Part 2: Rich interactions**
  - **2.1: Recommending links in networks**
  - 2.2: Predicting tie strength
  - 2.3: Predicting friends vs. foes
  - 2.4: How do people evaluate others?

[LibenNowell-Kleinberg '03]

Link prediction in networks

- Link prediction task:
  - Given $G[t_0,t_0']$ a graph on edges up to time t_0' output a ranked list L of links (not in $G[t_0,t_0']$) that are predicted to appear in $G[t_1,t_1']$
- Evaluation:
  - $n=|E_{new}|$: # new edges that appear during the test period $[t_1,t_1']$
  - Take top n elements of L and count correct edges

Link prediction via node distance

- Predict links evolving collaboration network

	training period			Core		
	authors	papers	collaborations[1]	authors	$\lvert E_{old} \rvert$	$\lvert E_{new} \rvert$
astro-ph	5343	5816	41852	1561	6178	5751
cond-mat	5469	6700	19881	1253	1899	1150
gr-qc	2122	3287	5724	486	519	400
hep-ph	5414	10254	47806	1790	6654	3294
hep-th	5241	9498	15842	1438	2311	1576

- Core: Since network data is very sparse
  - Consider only nodes with in-degree and out-degree of at least 3

[LibenNowell-Kleinberg '03]

Link prediction via proximity

- For every pair of nodes *(x,y)* compute:

graph distance	(negated) length of shortest path between x and y
common neighbors	$\lvert\Gamma(x) \cap \Gamma(y)\rvert$
Jaccard's coefficient	$\frac{\lvert\Gamma(x)\cap\Gamma(y)\rvert}{\lvert\Gamma(x)\cup\Gamma(y)\rvert}$
Adamic/Adar	$\sum_{z\in\Gamma(x)\cap\Gamma(y)} \frac{1}{\log\lvert\Gamma(z)\rvert}$
preferential attachment	$\lvert\Gamma(x)\rvert \cdot \lvert\Gamma(y)\rvert$
Katz_β	$\sum_{\ell=1}^{\infty} \beta^\ell \cdot \lvert\mathsf{paths}_{x,y}^{\langle\ell\rangle}\rvert$

where $\mathsf{paths}_{x,y}^{\langle\ell\rangle} :=$ {paths of length exactly ℓ from x to y}
weighted: $\mathsf{paths}_{x,y}^{\langle 1\rangle} :=$ number of collaborations between x, y.
unweighted: $\mathsf{paths}_{x,y}^{\langle 1\rangle} := 1$ iff x and y collaborate.

- Sort the pairs by score and predict top *n* pairs as new links

$$E^*_{new} := E_{new} \cap (\text{Core} \times \text{Core})$$

Γ(x) … degree of node *x*

Results: Improvement over random

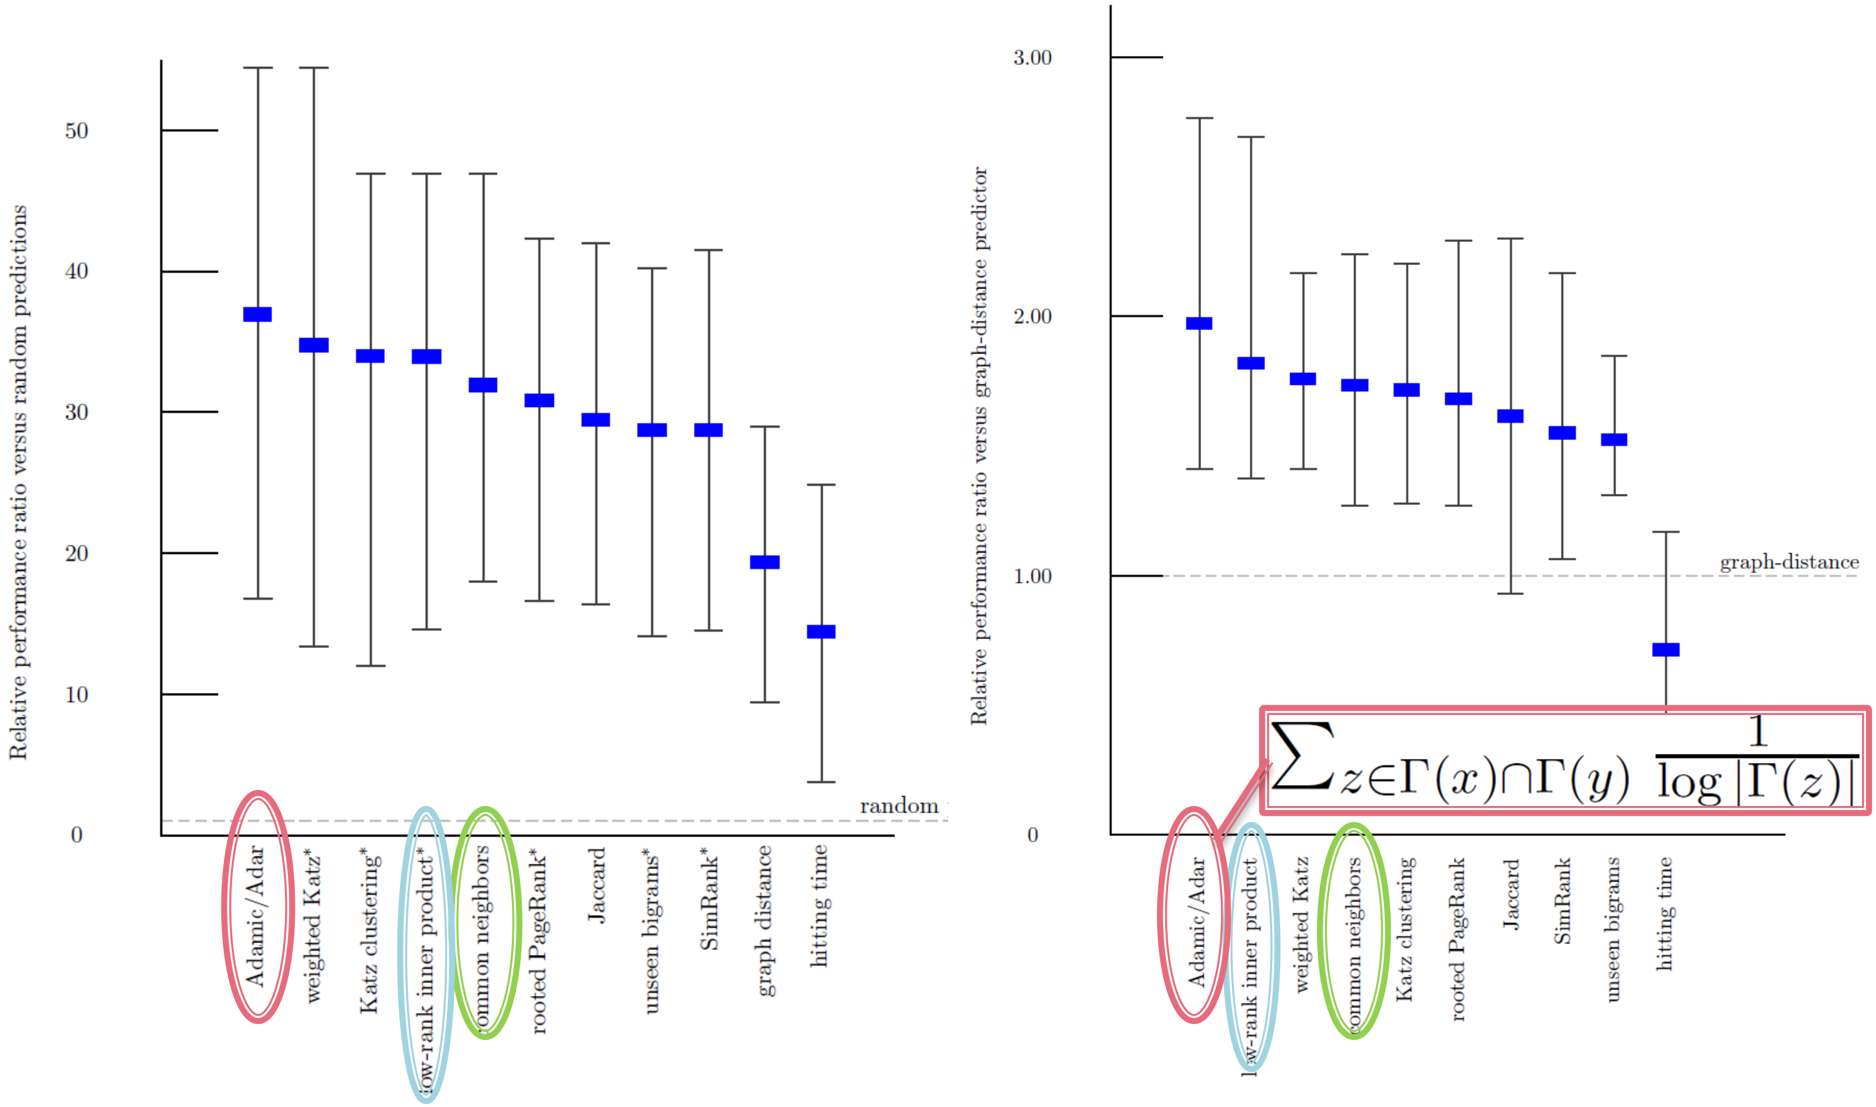

Supervised Link Prediction

- How to learn to predict new friends in networks?
  - Facebook's People You May Know
  - Let's look at the data:
    - 92% of new friendships on FB are friend-of-a-friend
    - More common friends helps

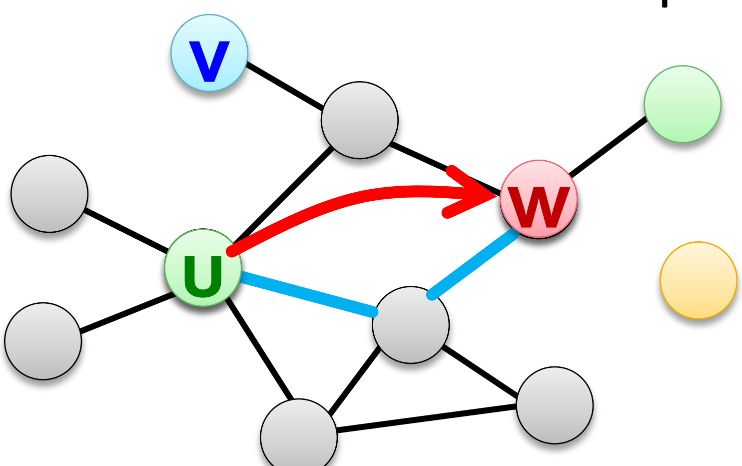

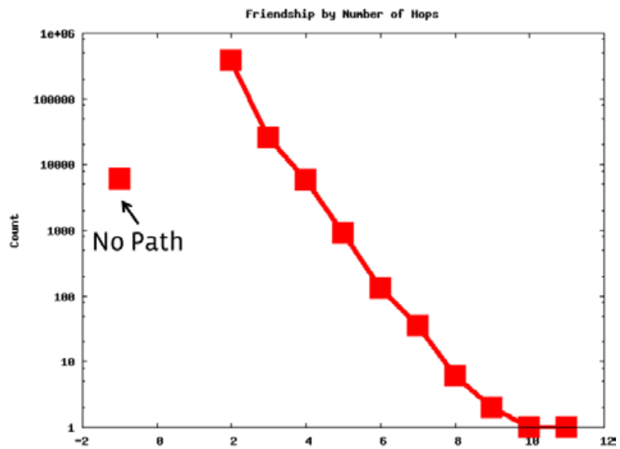

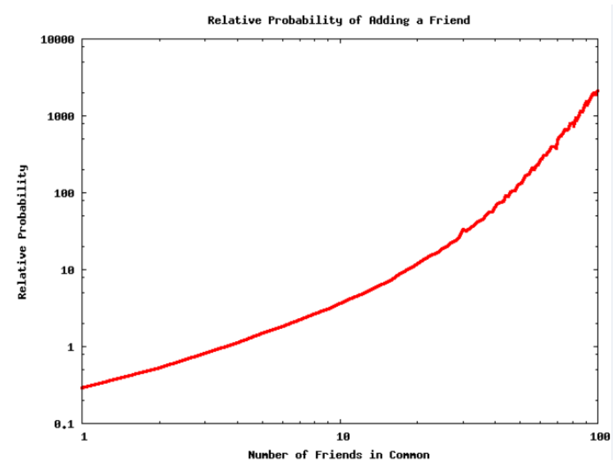

Supervised Link Prediction

- Recommend a list of possible friends
- Supervised machine learning setting:
  - Training example:
    - For every node s have a list of nodes she will create links to $\{v_1, \ldots, v_k\}$
      - *E.g.*, use FB network from May 2011 and $\{v_1, \ldots, v_k\}$ are the new friendships you created since then
  - Problem:
    - For a given node s learn to rank nodes $\{v_1, \ldots, v_k\}$ higher than other nodes in the network
- Supervised Random Walks based on word by Agarwal&Chakrabarti

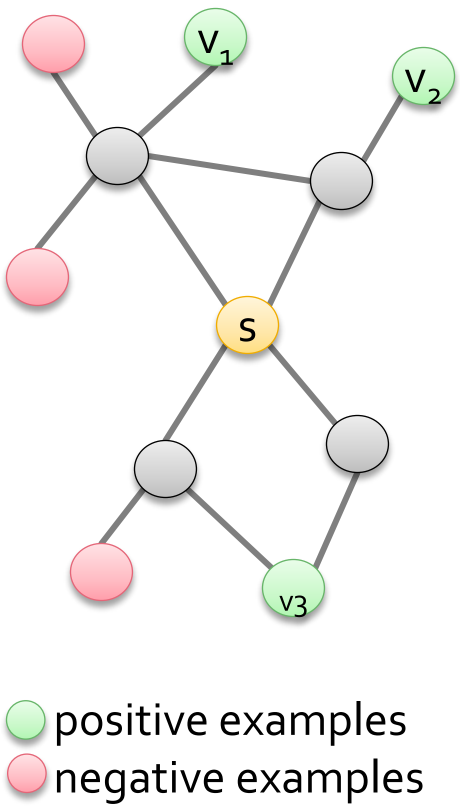

Supervised Link Prediction

- How to combine node/edge attributes and the network structure?
  - Learn a strength of each edge based on:
    - Profile of user u, profile of user v
    - Interaction history of u and v
  - Do a PageRank-like random walk from s to measure the "proximity" between s and other nodes
  - Rank nodes by their "proximity" (i.e., visiting prob.)

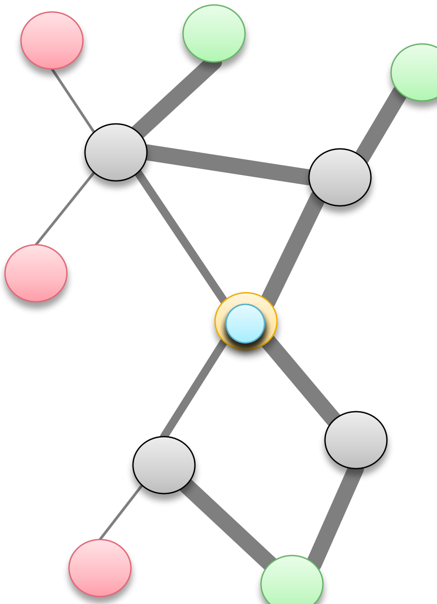

Supervised Random Walks

- Let s be the center node
- Let $f_w(u,v)$ be a function that assigns a strength to each edge:

  $a_{uv} = f_w(u,v) = exp(-w^T \Psi_{uv})$
  - Ψ_{uv} is a feature vector
    - Features of node u
    - Features of node v
    - Features of edge (u,v)
  - **w is the parameter vector we want to learn**
- Do Random Walk with Restarts from s where transitions are according to edge strengths
- How to learn $f_w(u,v)$?

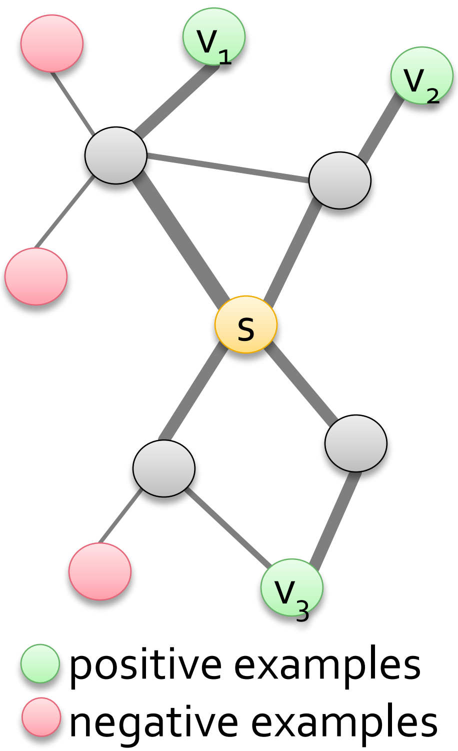

Personalized PageRank

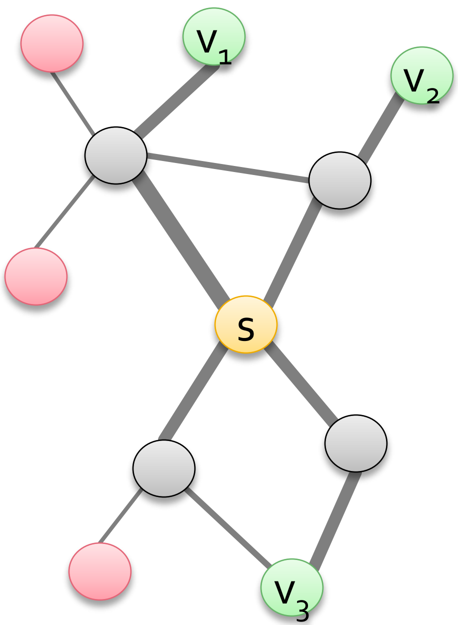

- Random walk transition matrix:

$$Q'_{uv} = \begin{cases} \frac{a_{uv}}{\sum_w a_{uw}} & \text{if } (u,v) \in E, \\ 0 & \text{otherwise} \end{cases}$$

- PageRank transition matrix:

$$Q_{ij} = (1-\alpha)Q'_{ij} + \alpha \mathbf{1}(j = s)$$

  - with prob. α jump back to s

- Compute PageRank vector: $p = p^T Q$
- Rank nodes by p_u

The Optimization Problem

- Each node u has a score p_u
- Destination nodes $D = \{v_1, \ldots, v_k\}$
- No-link nodes $L = \{the\ rest\}$

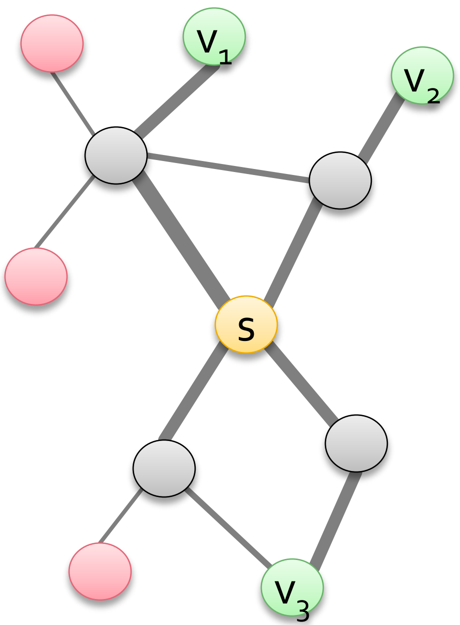

- What do we want?
  Want to find w such that $p_l < p_d$

$$\min_w F(w) = ||w||^2$$

such that

$$\forall\, d \in D, l \in L:\; p_l < p_d$$

  - Hard constraints, make them soft

Making constraints soft

- Want to minimize:

$$\min_{w} F(w) = ||w||^2 + \lambda \sum_{ld} h(p_l - p_d)$$

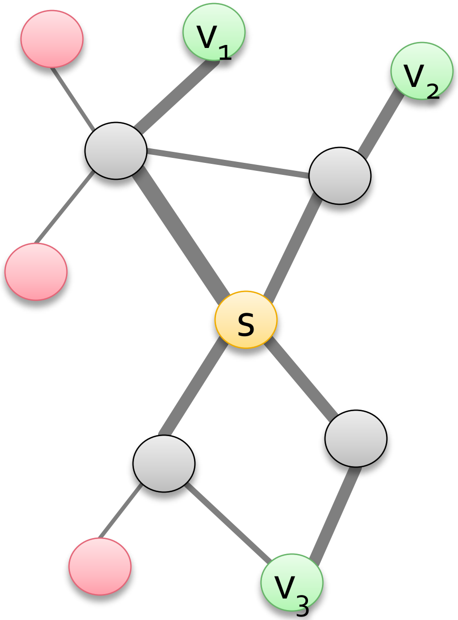

  - **Loss:** *h(x)=0* if *x<0*, x^2 else

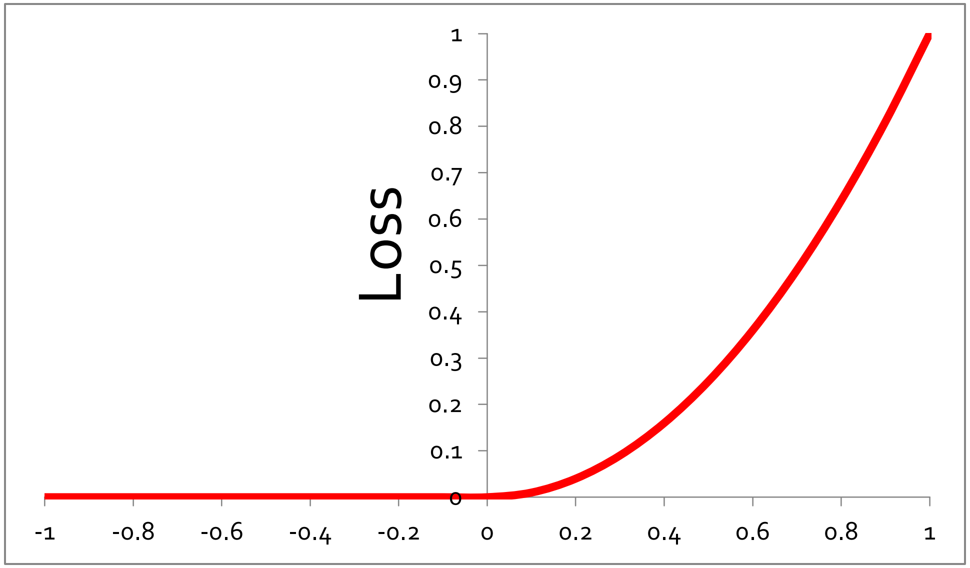

$p_l < p_d$ $p_l = p_d$ $p_l > p_d$

Solving the problem: Intuition

- How to minimize *F*?

$$\min_{w} F(w) = ||w||^2 + \lambda \sum_{ld} h(p_l - p_d)$$

- p_l and p_d depend on w
  - Given w assign edge weights $a_{uv}=f_w(u,v)$
  - Using transition matrix $Q=[a_{uv}]$ compute PageRank scores p_u
  - Rank nodes by the PageRank score
- Want to find w such that $p_l < p_d$

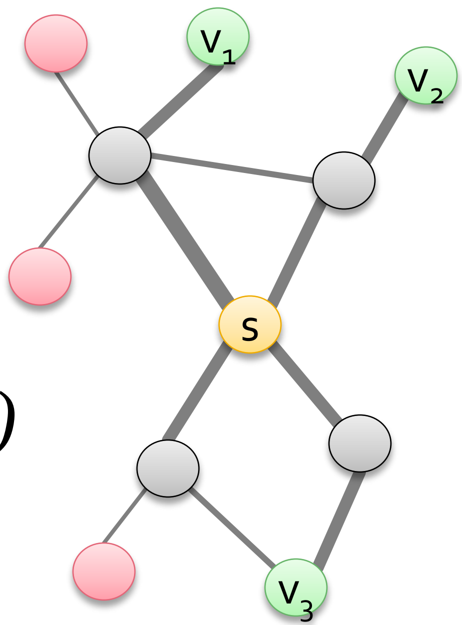

Gradient Descent

- How to minimize F?

$$\min_w F(w) = ||w||^2 + \lambda \sum_{ld} h(p_l - p_d)$$

  - Take the derivative!

$$\frac{\partial F}{\partial w} = 2w + \sum_{l,d} \frac{\partial h(p_l - p_d)}{\partial w}$$

$$= 2w + \sum_{l,d} \frac{\partial h(\delta_{ld})}{\partial \delta_{ld}} \left(\frac{\partial p_l}{\partial w} - \frac{\partial p_d}{\partial w}\right)$$

- We know:

$$p = p^T Q \text{ i.e. } p_u = \sum_j p_j Q_{ju}$$

- So:

$$\frac{\partial p_u}{\partial w} = \sum_j Q_{ju} \frac{\partial p_j}{\partial w} + p_j \frac{\partial Q_{ju}}{\partial w}$$

  - Solve by power iteration!

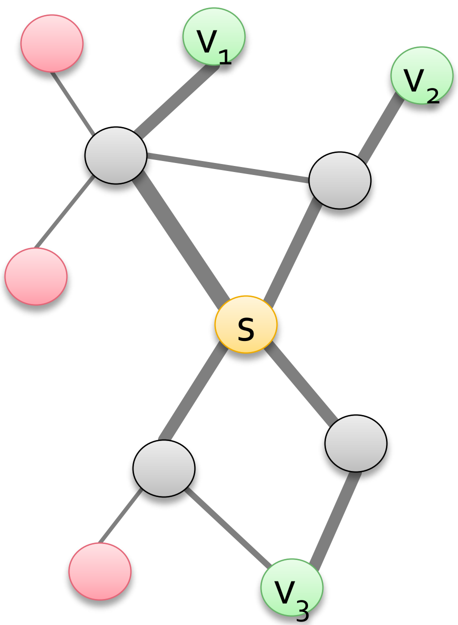

Optimizing F

- To optimize F, use gradient based method:
  - Pick a random starting point w_0
  - Compute the personalized PageRank vector p
  - Compute gradient with respect to weight vector w
  - Update w
    - Optimize using quasi-Newton method

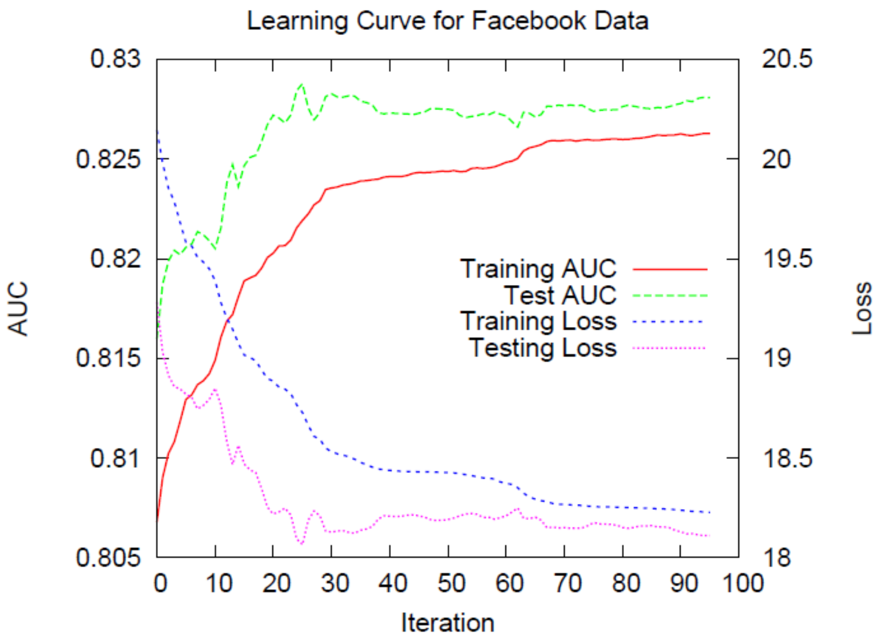

Data: Facebook

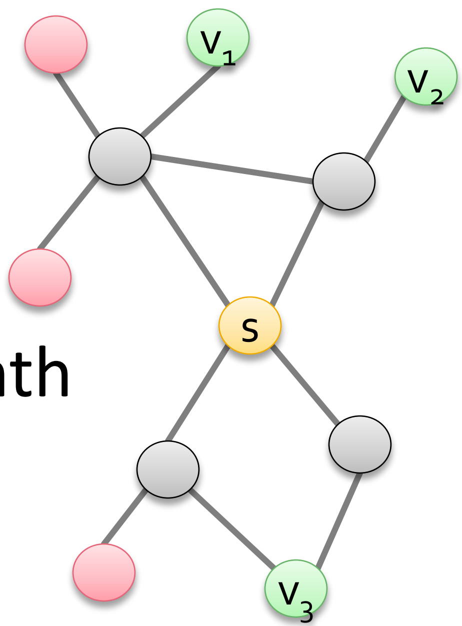

- Facebook Iceland network
  - 174,000 nodes (55% of population)
  - Avg. degree 168
  - Avg. person added 26 new friends/month
- For every node s:
  - Positive examples:
    - $D = \{$ new friendships of s created in Nov '09 $\}$
  - Negative examples:
    - $L = \{$ other nodes s did not create new links to $\}$
  - Limit to friends of friends
    - on avg. there are 20k FoFs (max 2M)!

Experimental setting

- Node and Edge features for learning:
  - Node:
    - Age
    - Gender
    - Degree
  - Edge:
    - Age of an edge
    - Communication,
    - Profile visits
    - Co-tagged photos
- Baselines:
  - Decision trees and logistic regression:
    - Above features + 10 network features (PageRank, common friends)
- Evaluation:
  - AUC and precision at Top20

Results: Facebook Iceland

- Facebook: predict future friends
  - Adamic-Adar already works great
  - Logistic regression also strong
  - SRW gives slight improvement

Learning Method	AUC	Prec@20
Random Walk with Restart	0.81725	6.80
Adamic-Adar	0.81586	7.35
Common Friends	0.80054	7.35
Degree	0.58535	3.25
DT: Node features	0.59248	2.38
DT: Network features	0.76979	5.38
DT: Node+Network	0.76217	5.86
DT: Path features	0.62836	2.46
DT: All features	0.72986	5.34
LR: Node features	0.54134	1.38
LR: Network features	0.80560	7.56
LR: Node+Network	0.80280	7.56
LR: Path features	0.51418	0.74
LR: All features	0.81681	7.52
SRW: one edge type	**0.82502**	6.87
SRW: multiple edge types	**0.82799**	**7.57**

Tutorial Outline

- Part 1: Information flow in networks
- **Part 2: Rich interactions**
  - 2.1: Recommending links in networks
  - **2.2: Predicting tie strength**
  - 2.3: Predicting friends vs. foes
  - 2.4: How do people evaluate others?

Tie strength

- *"The strength of a tie is a combination of the amount of TIME, the emotional INTENSITY, the INTIMACY, and the reciprocal SERVICES which characterize the tie."* [Grannovetter]

- Gilbert & Karahalios surveyed 35 Facebook users to label 2,184 friendships (links)

- Describe each link by 70+ features
- Train a regression model to predict tie strength

[Gilbert-Karahalios, '09]

Five aspects of tie strength

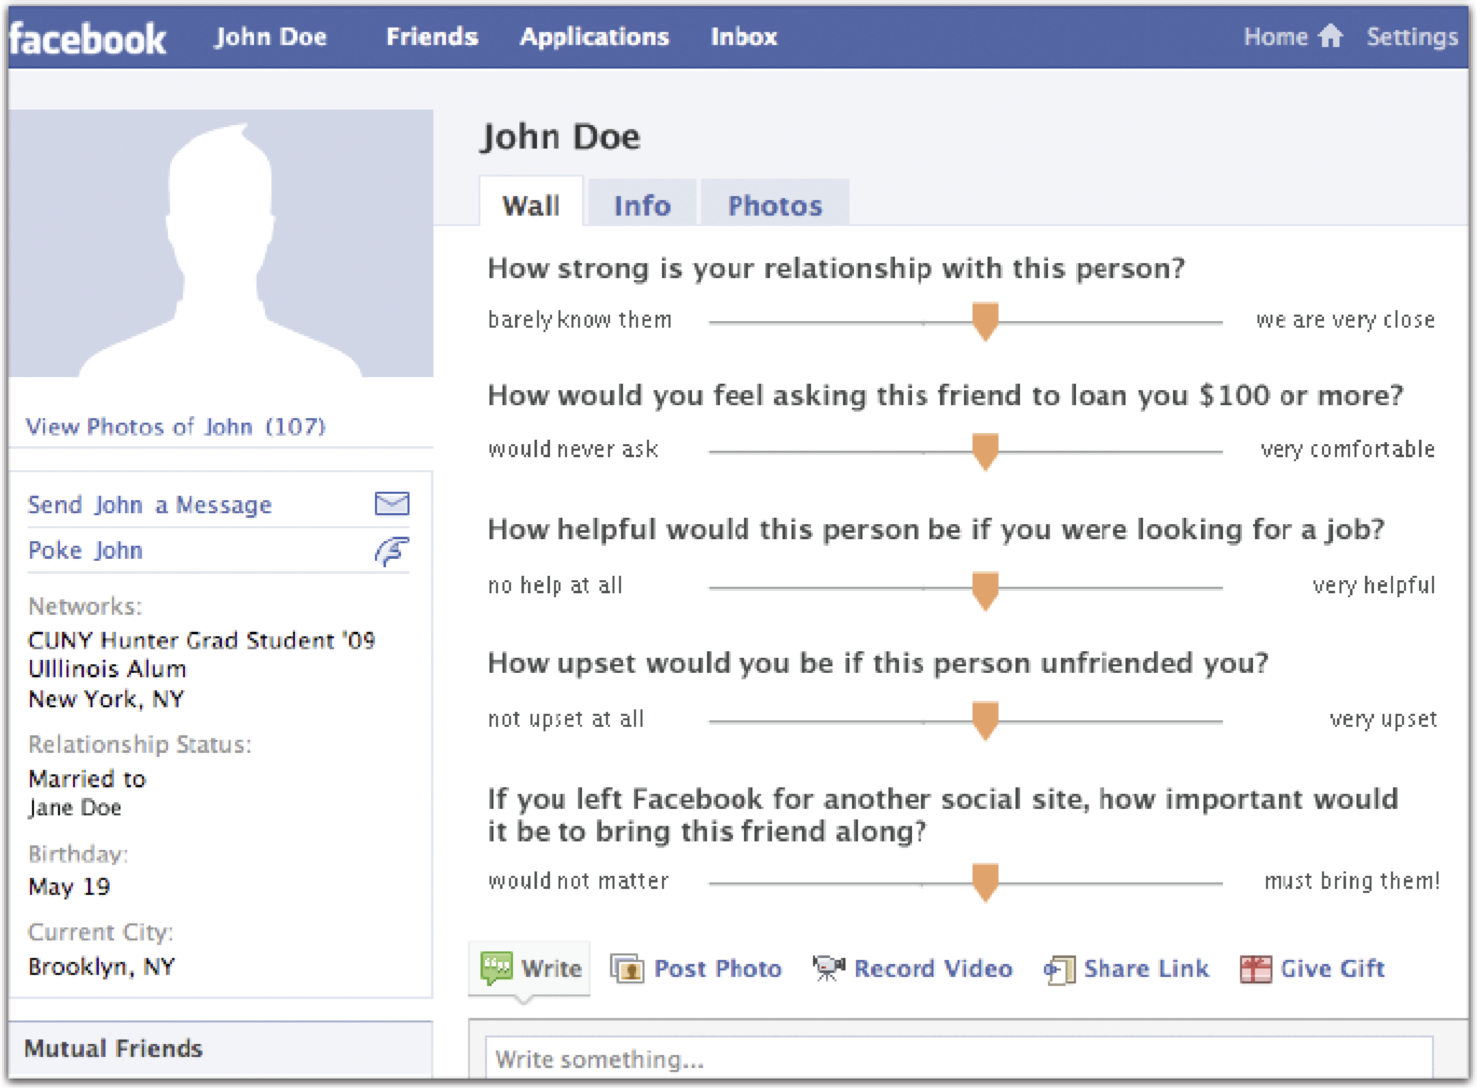

[Gilbert-Karahalios, '09]

Attributes of the friendship (1)

- Features that are used for learning
- Intensity
  - wall words exchanged
  - friend-initiated wall posts
  - part.-initiated wall posts
  - inbox messages together
  - inbox thread depth
  - part.'s status updates
  - friend's status updates
- Intimacy
  - participant's friends
  - friend's friends
  - days since last comm.
  - wall intimacy words
  - inbox intimacy words
  - together in photo
  - miles between hometowns

[Gilbert-Karahalios, '09]

Attributes of the friendship (2)

- Features that are used for learning
- Social Distance
  - age difference
  - # occupations difference
  - educational difference
  - political difference
- Reciprocal services
  - Links exchanged by wall
  - Applications in common
- Structural
  - mutual friends
  - groups in common
  - Cosine similarity of interests
- Emotional support
  - Positive emotion words
  - Negative emotion words

Results: How strong is relationship

- Train a linear (regression) model
- Results for the "How strong is your relationship?"

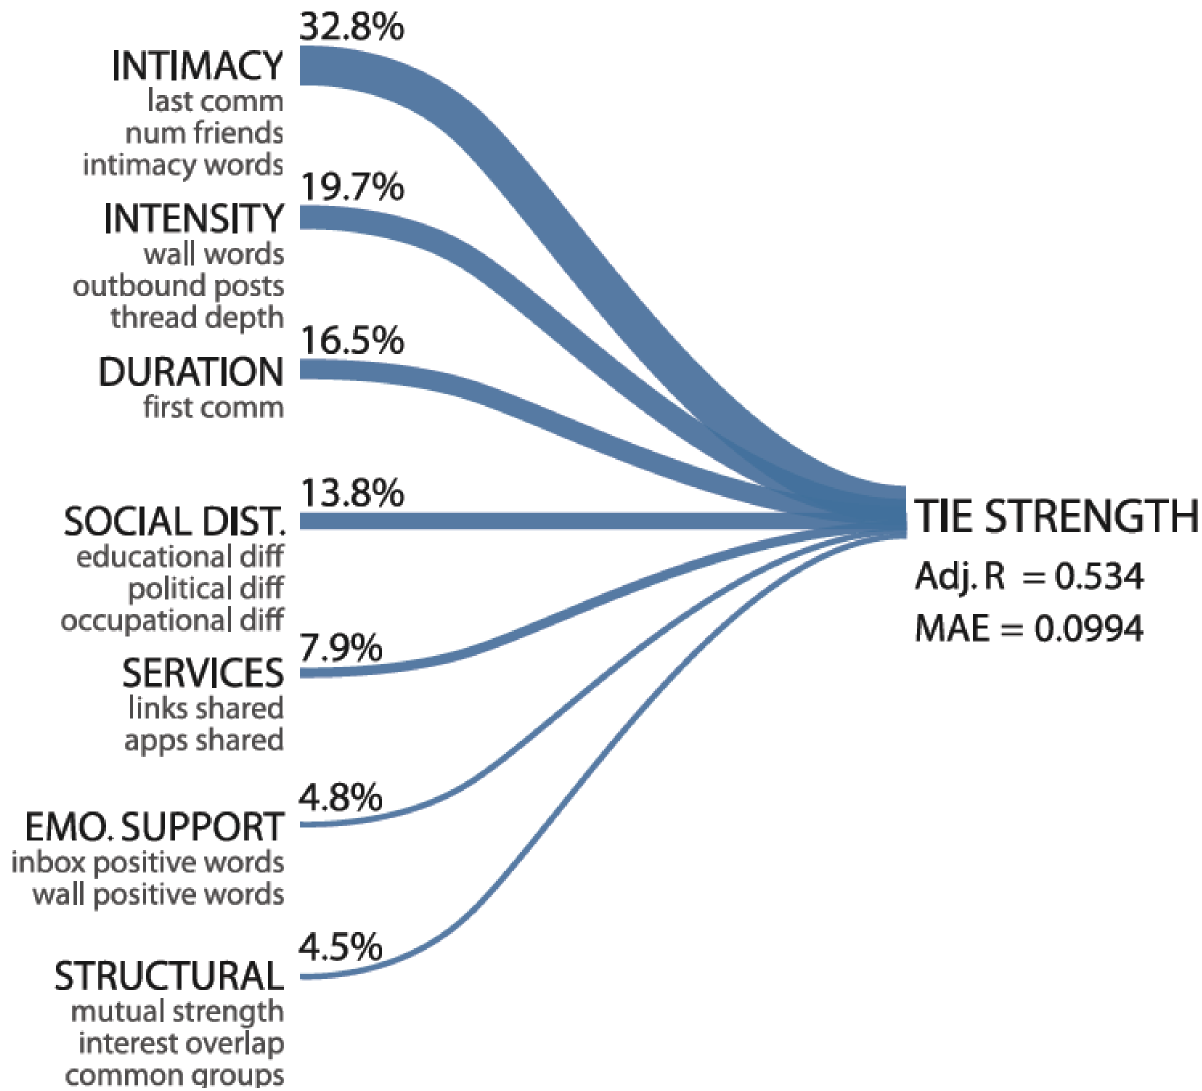

[Gilbert-Karahalios, '09]

Results: Most predictive features

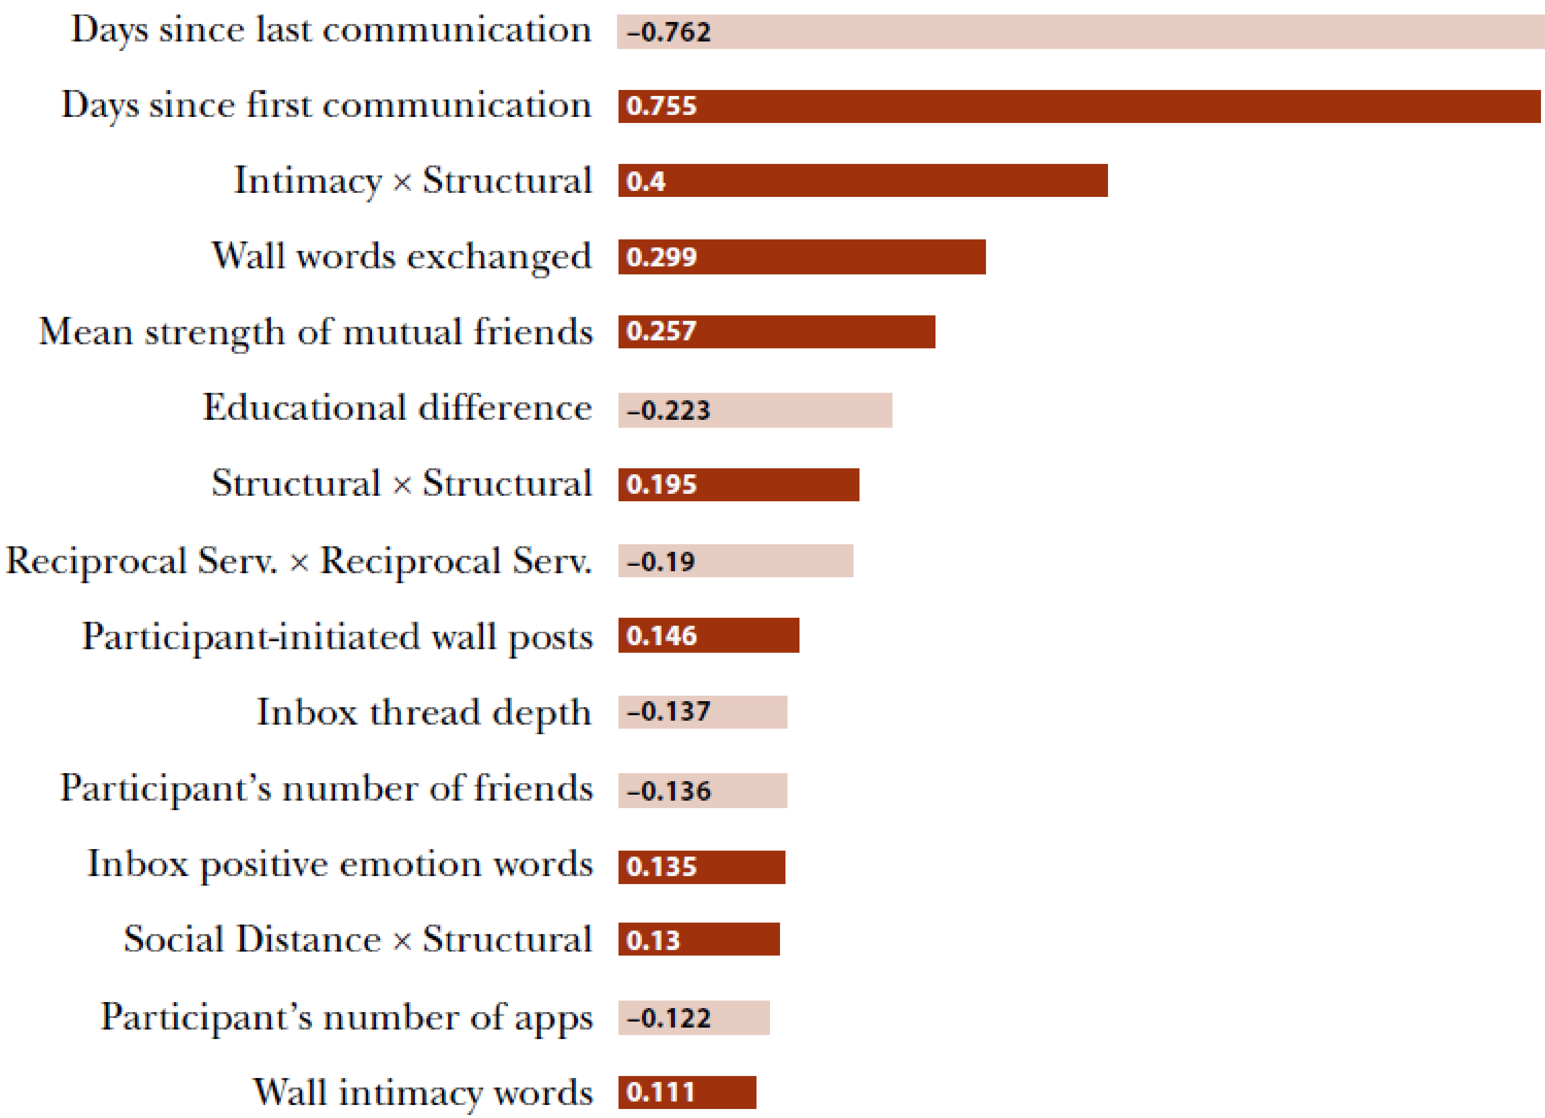

Tutorial Outline

- Part 1: Information flow in networks
- **Part 2: Rich interactions**
  - 2.1: Recommending links in networks
  - 2.2: Predicting tie strength
  - **2.3: Predicting friends vs. foes**
  - 2.4: How do people evaluate others?

Friends vs. Foes

- So far we viewed links as positive but links can also be negative
- Question:
  - How do edge signs and network interact?
  - How to model and predict edge signs?
- Applications:
  - Friend recommendation
    - Not just whether you know someone but what do you think of them

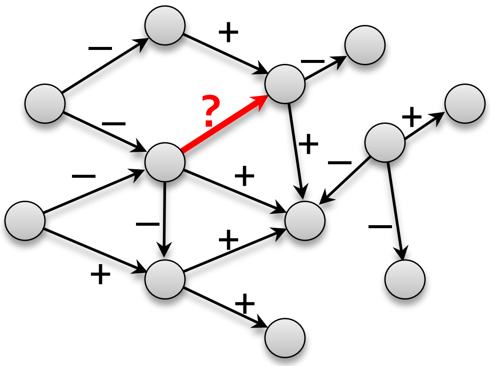

Networks with Explicit Signs

- Each link A→B is **explicitly** tagged with a sign:
  - Epinions: Trust/Distrust
    - Does A trust B's product reviews?
      (only positive links are visible)
  - Wikipedia: Support/Oppose
    - Does A support B to become Wikipedia administrator?
  - Slashdot: Friend/Foe
    - Does A like B's comments?

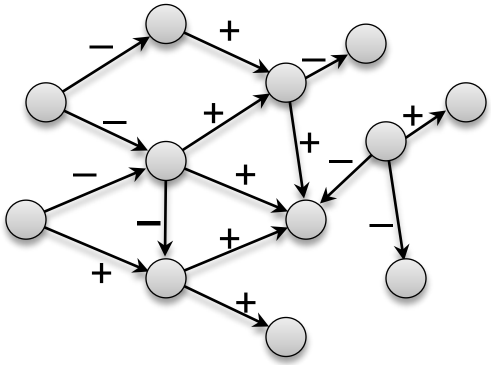

	Epinions	Slashdot	Wikipedia
Nodes	119,217	82,144	7,118
Edges	841,200	549,202	103,747
+ edges	85.0%	77.4%	78.7%
− edges	15.0%	22.6%	21.2%

Theory of Structural Balance

Consider edges as undirected

- Start with intuition [Heider '46]:
  - Friend of my friend is my friend
  - Enemy of enemy is my friend
  - Enemy of friend is my enemy
- Look at connected triples of nodes:

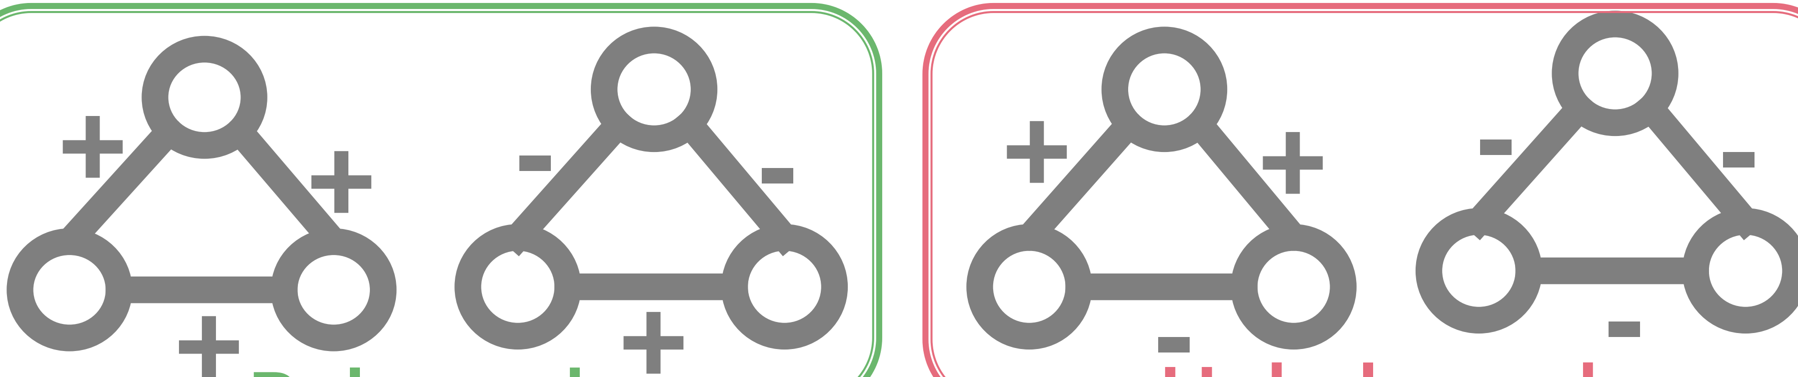

Balanced

Consistent with "friend of a friend" or "enemy of the enemy" intuition

Unbalanced

Inconsistent with the "friend of a friend" or "enemy of the enemy" intuition

Theory of Status

- Status theory [Davis-Leinhardt '68, Guha et al. '04, Leskovec et al. '10]
  - Link A $\xrightarrow{+}$ B means: B has higher status than A
  - Link A $\xrightarrow{-}$ B means: B has lower status than A
  - Based on signs/directions of links from/to node X make a prediction
- Status and balance can make different predictions:

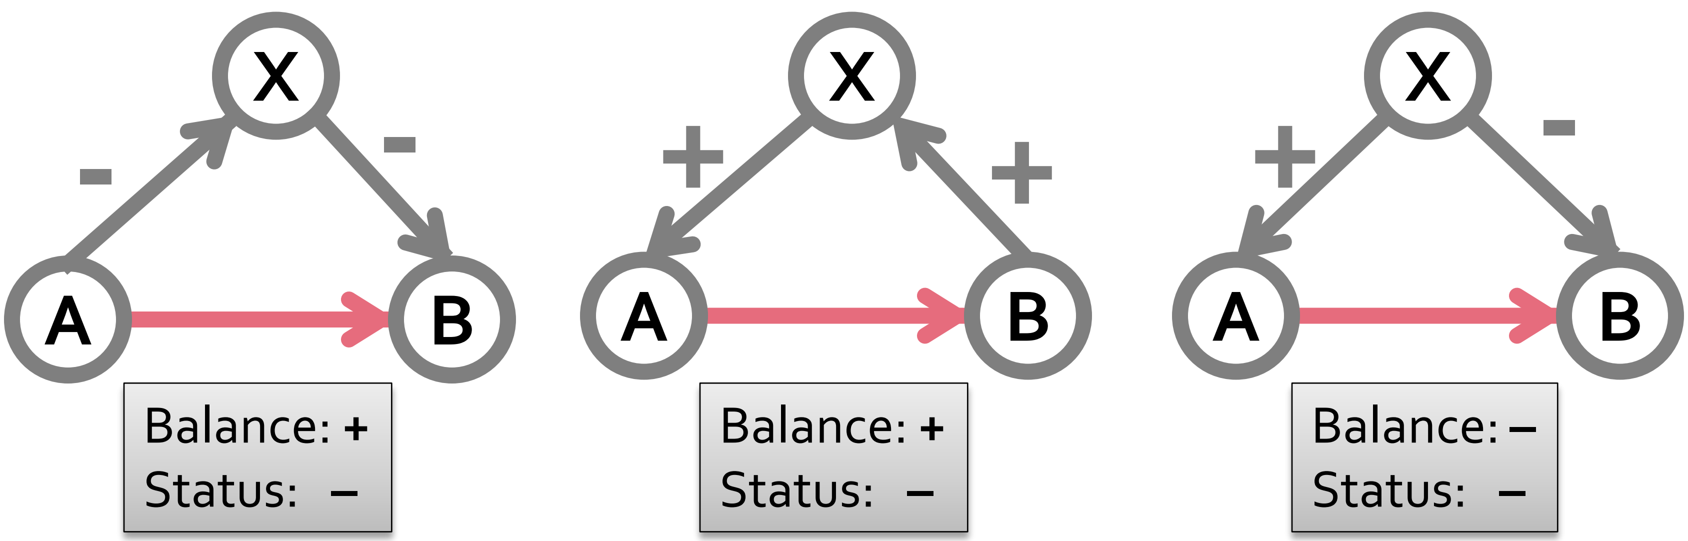

The Plan

- How do these two theories align with ways people create links:
  - Not just "which is right" but how are aspects of each reflected in the data
  - Provide insights into how these linking systems are being used
- Outline:
  - Study links as undirected: Balance theory
  - Study links as directed and evolving: Status theory
  - Predicting signs of edges

Evolving directed networks

- Our networks are really directed
  - trust, opinion (, friendship)
- How many △ are now explained by balance?
  - Half (8 out of 16)
- Is there a better explanantion?
  - Yes. **Theory of Status.**

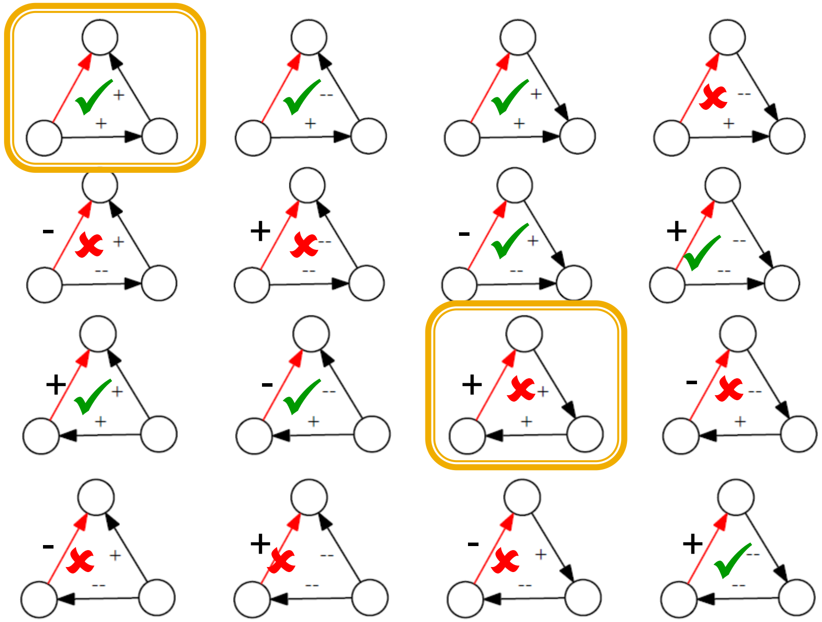

16 signed directed triads

Alternate theory: Status

- Links are directed and created over time
- Status theory [Davis-Leinhardt '68, Guha et al. '04, Leskovec et al. '10]
  - Link $A \xrightarrow{+} B$ means: B has higher status than A
  - Link $A \xrightarrow{-} B$ means: B has lower status than A
- Status and balance can give different predictions:

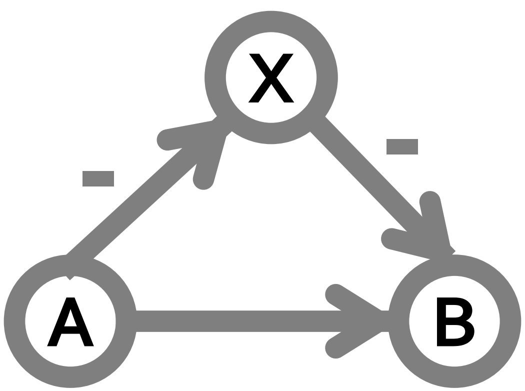

Balance: **+**
Status: –

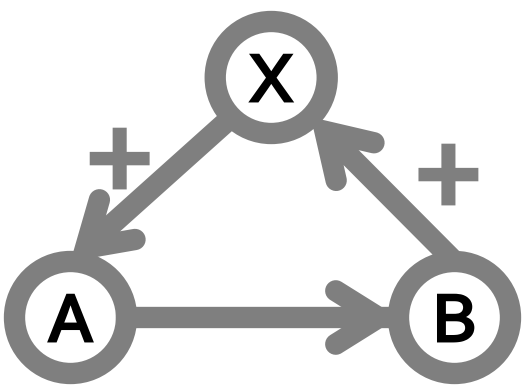

Balance: +
Status: –

Evolving Directed Networks

- Links are directed
- Links are created over time
  - X has links to/from A and B
  - Now, A links to B
- To compare balance and status we need to **formalize** :
  - Links are embedded in triads – provides context for signs
  - Users are heterogeneous in their linking behavior

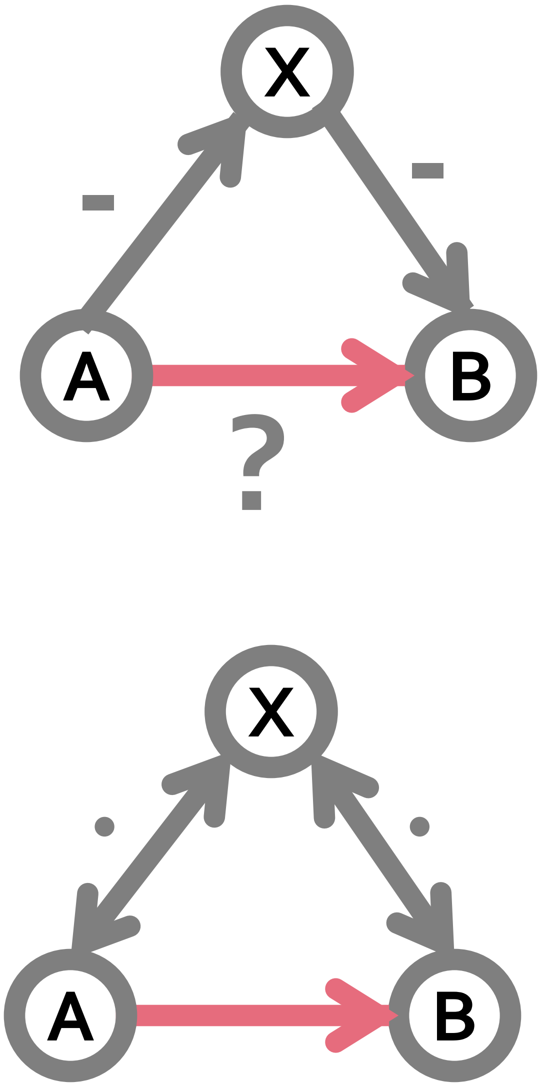

16 types of contextualized links

- **Link contexts:**
  - A contextualized link is a triple (A,B;X) such that directed A-B link forms after there is a two-step semi-path A-X-B
  - A-X and B-X links can have either direction and either sign: 16 possible types

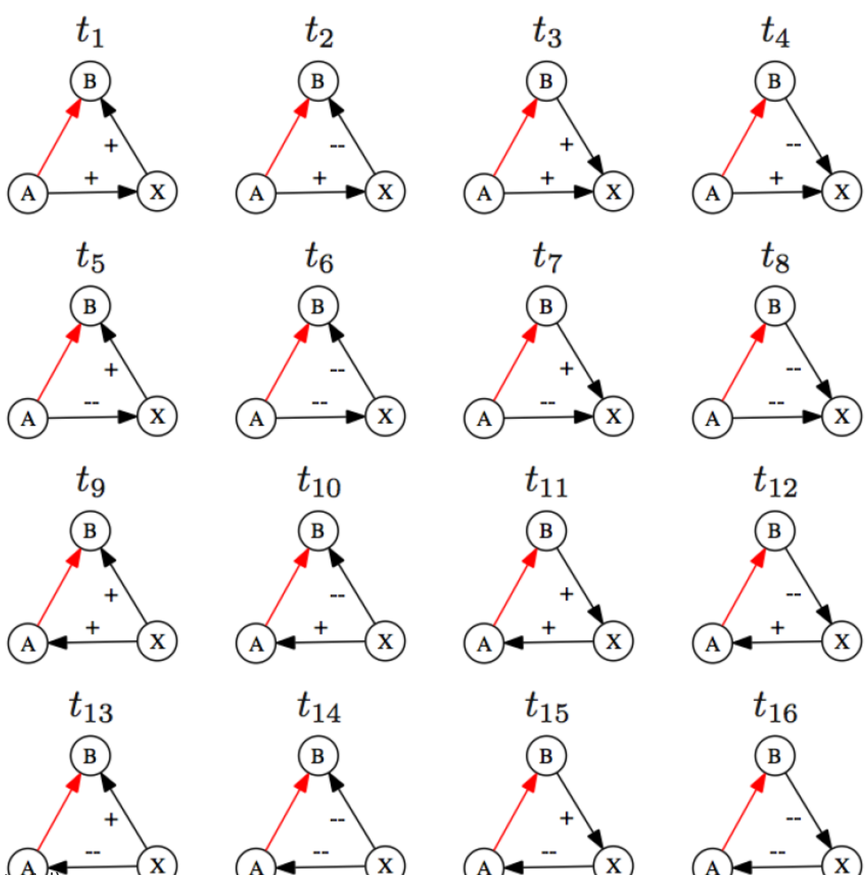

Heterogeneity in linking behavior

- Different users make signs differently:
  - Generative baseline (frac. of + given by A)
  - Receptive baseline (frac. of + received by B)
- How do different link contexts cause users to deviate from baselines?
- **Surprise**: How much behavior of A/B deviates from **baseline** when they are in **context**

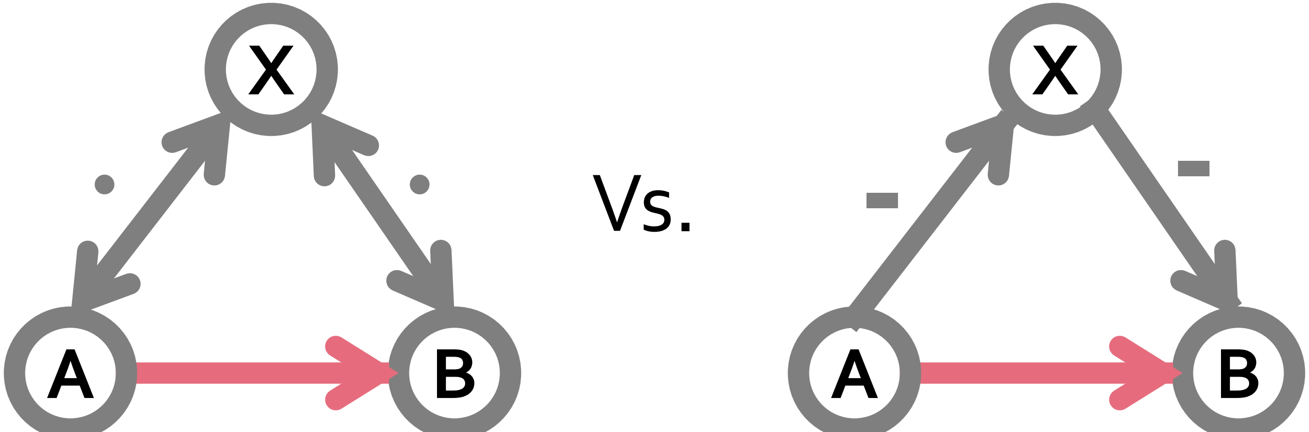

Status: Two Examples

- Two basic examples:

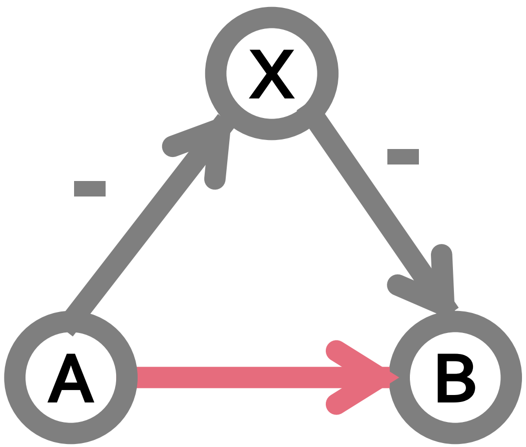

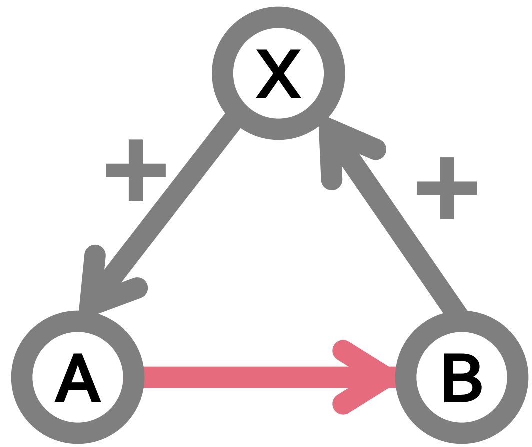

More **negative** than gen. baseline of A
More **negative** than rec. baseline of B

More **negative** than gen. baseline of A
More **negative** than rec. baseline of B

Consistency with Status

- Determine node status:
  - Assign X status 0
  - Based on signs and directions of edges set status of A and B

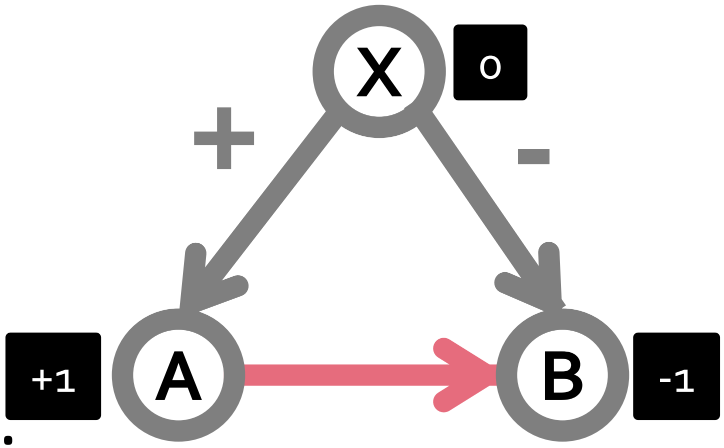

Status-consistent if:
Gen. surprise < 0
Rec. surprise < 0

- Surprise is **status**-consistent, if:
  - Gen. surprise is status-consistent if it has **same** sign as status of B
  - Rec. surprise is status-consistent if it has the **opposite** sign from the status of A
- Surprise is **balance**-consistent, if:
  - If it completes a balanced triad

Status vs. Balance (Epinions)

- Results for Epinions: Trust vs. Distrust

t_i	count	$P(+)$	s_{out}	s_{in}	B_{out}	B_{in}	S_{out}	S_{in}
t_1	178,051	0.97	95.9	197.8	✓	✓	✓	✓
t_2	45,797	0.54	-151.3	-229.9	✓	✓	✓	○
t_3	246,371	0.94	89.9	195.9	✓	✓	○	✓
t_4	25,384	0.89	1.8	44.9	○	○	✓	✓
t_5	45,925	0.30	18.1	-333.7	○	✓	✓	✓
t_6	11,215	0.23	-15.5	-193.6	○	○	✓	✓
t_7	36,184	0.14	-53.1	-357.3	✓	✓	✓	✓
t_8	61,519	0.63	124.1	-225.6	✓	○	✓	✓
t_9	338,238	0.82	207.0	-239.5	✓	○	✓	✓
t_{10}	27,089	0.20	-110.7	-449.6	✓	✓	✓	✓
t_{11}	35,093	0.53	-7.4	-260.1	○	○	✓	✓
t_{12}	20,933	0.71	17.2	-113.4	○	✓	✓	✓
t_{13}	14,305	0.79	23.5	24.0	○	○	✓	✓
t_{14}	30,235	0.69	-12.8	-53.6	○	○	✓	○
t_{15}	17,189	0.76	6.4	24.0	○	○	○	✓
t_{16}	4,133	0.77	11.9	-2.6	✓	○	✓	○
				Number of correct predictions	8	7	14	13

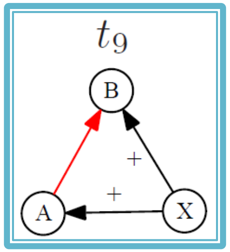

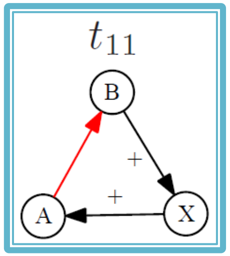

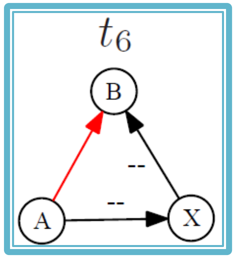

Global Structure of Signed Nets

- Intuitive picture of social network in terms of densely linked clusters
- How does structure interact with links?
- Embeddedness of link (A,B): number of hared neighbors

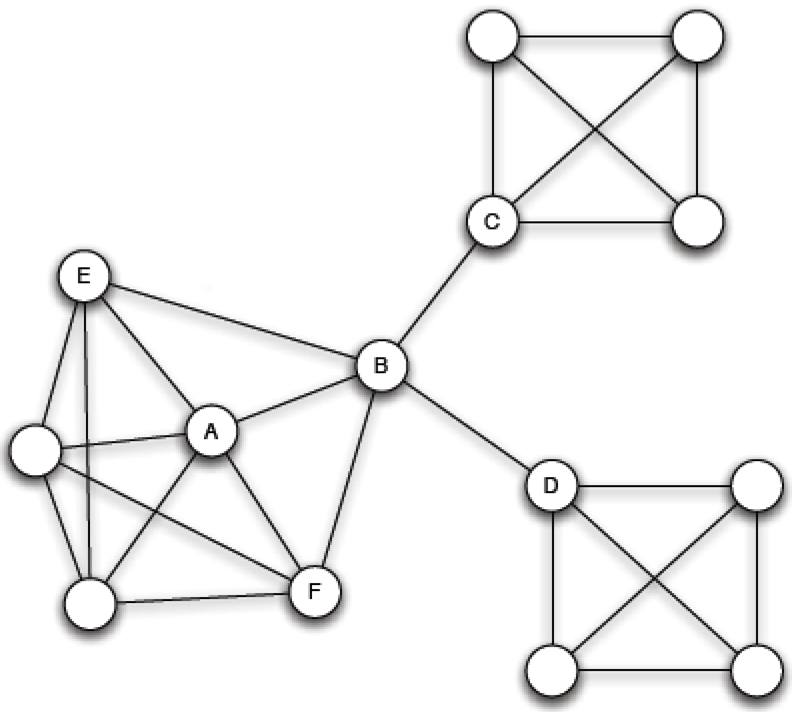

Global structure: Embeddedness

- Embeddedness of ties:
  - Embedded ties tend to be more positive
- A natural connection to closure based social capital [Coleman '88]
- Public display of signs (votes) in Wikipedia further strengthens this

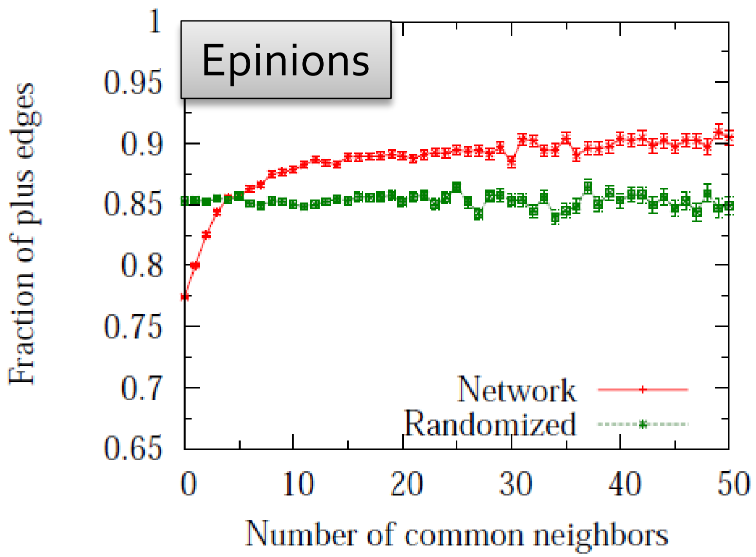

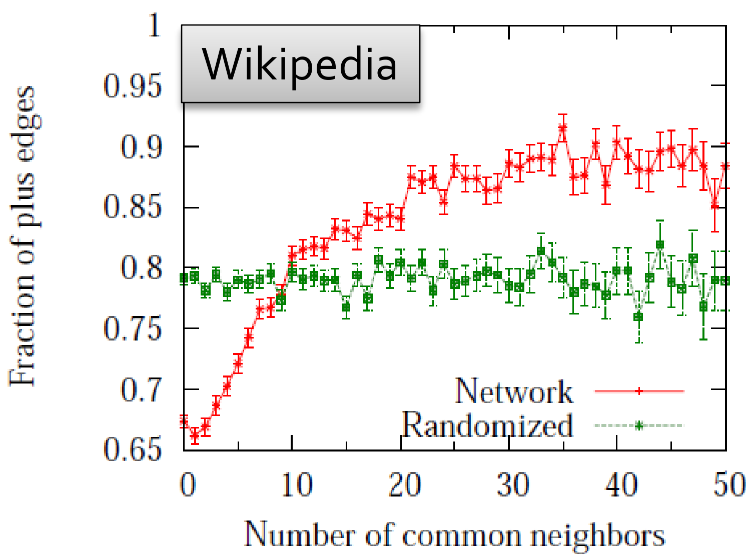

Predicting edge signs

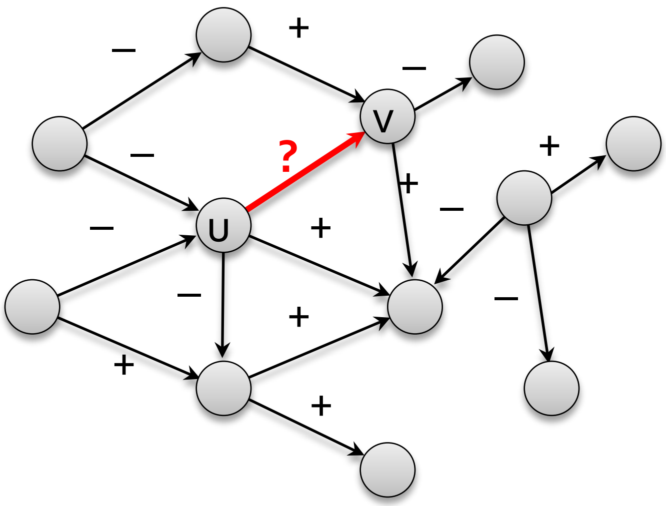

Edge sign prediction problem

- Given a network and signs on all but one edge, predict the missing sign

Machine Learning formulation:

- Predict sign of edge (u,v)
- Class label:
  - +1: positive edge
  - -1: negative edge
- Learning method:
  - Logistic regression

$$P(+|x) = \frac{1}{1 + e^{-(b_0 + \sum_i^n b_i x_i)}}$$

- Dataset:
  - Original: 80% +edges
  - Balanced: 50% +edges
- Evaluation:
  - Accuracy and ROC curves
- Features for learning:
  - Next slide

Features for learning

For each edge (u,v) create features:

- Triad counts (16):
  - Counts of signed triads edge $u \rightarrow v$ takes part in
- Degree (7 features):
  - Signed degree:
    - $d^{+}_{out}(u)$, $d^{-}_{out}(u)$, $d^{+}_{in}(v)$, $d^{-}_{in}(v)$
  - Total degree:
    - $d_{out}(u)$, $d_{in}(v)$
  - Embeddedness of edge (u,v)

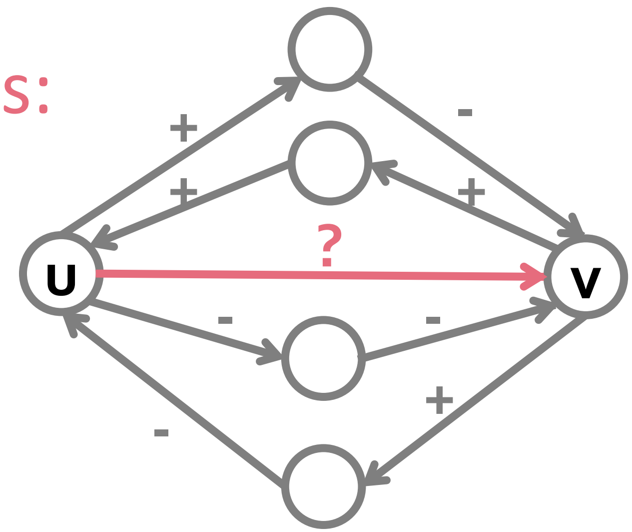

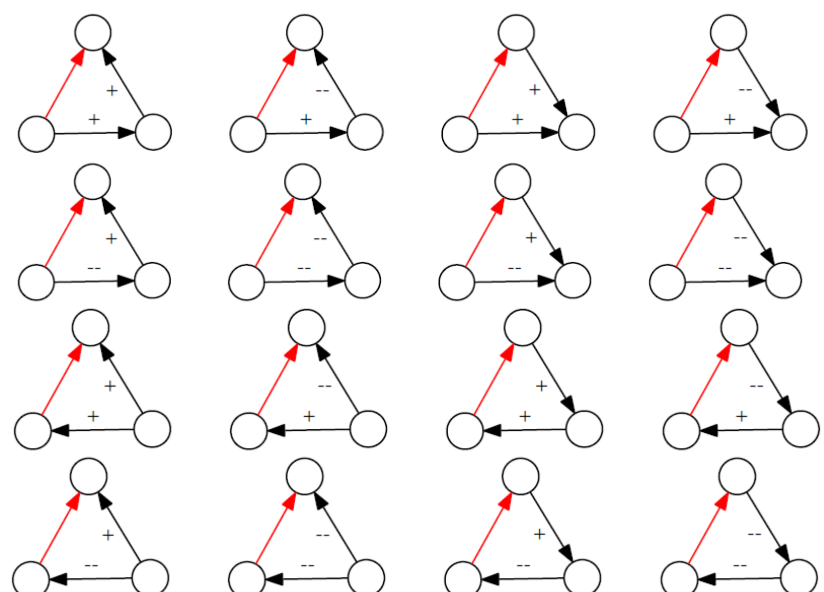

Edge sign prediction

- Error rates:
  - Epinions: 6.5%
  - Slashdot: 6.6%
  - Wikipedia: 19%
- Signs can be modeled from local network structure alone
  - Trust propagation model of [Guha et al. '04] has 14% error on Epinions
- Triad features perform less well for less embedded edges
- Wikipedia is harder to model:
  - Votes are publicly visible

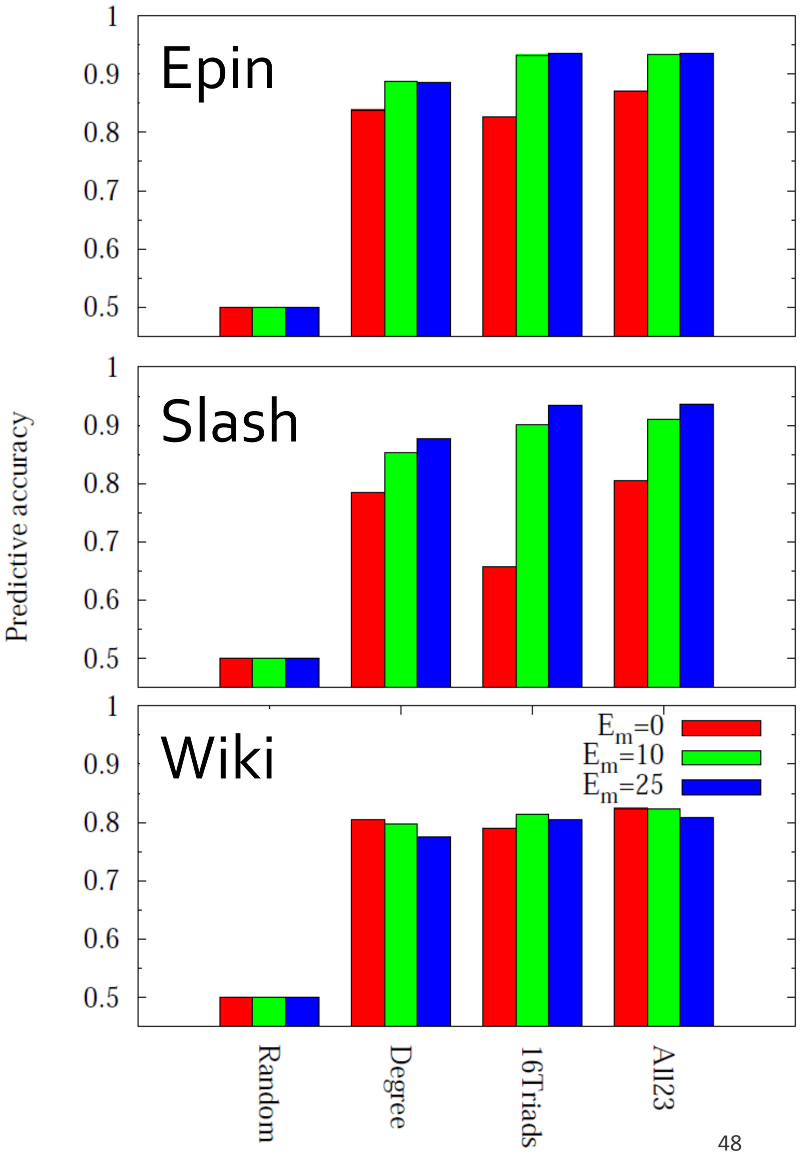

Generalization

- Do people use these very different linking systems by obeying the same principles?
  - How generalizable are the results across the datasets?
    - Train on row "dataset", predict on "column"

All23	Epinions	Slashdot	Wikipedia
Epinions	0.9342	0.9289	0.7722
Slashdot	0.9249	0.9351	0.7717
Wikipedia	0.9272	0.9260	0.8021

- Almost perfect generalization of the models even though networks come from very different applications

Balance and Status: Complete model

Feature	Bal	Stat	Epin	Slashd	Wikip
const			-0.2	0.02	-0.2
● +→ ● +→ ●	1	1	**0.5**	**0.9**	**0.3**
● +→ ● -→ ●	-1	0	**-0.5**	**-0.9**	**-0.4**
● -→ ● +→ ●	-1	0	**-0.4**	**-1.1**	**-0.3**
● -→ ● -→ ●	1	-1	**-0.7**	**-0.6**	**-0.8**
○ +→ ○ ←+ ○	1	0	0.3	0.4	0.05
○ +→ ○ ←- ○	-1	1	-0.01	-0.1	-0.01
● -→ ● ←+ ●	-1	-1	**-0.9**	**-1.2**	**-0.2**
○ -→ ○ ←- ○	1	0	0.04	-0.07	-0.03
○ ←+ ○ +→ ○	1	0	0.08	0.4	0.1
● ←+ ● -→ ●	-1	-1	**-1.3**	**-1.1**	**-0.4**
○ ←- ○ +→ ○	-1	1	-0.1	-0.2	0.05
○ ←- ○ -→ ○	1	0	0.08	-0.02	-0.1
○ ←+ ○ ←+ ○	1	-1	-0.09	-0.09	-0.01
○ ←+ ○ ←- ○	-1	0	-0.05	-0.3	-0.02
○ ←- ○ ←+ ○	-1	0	-0.04	-0.3	0.05
○ ←- ○ ←- ○	1	1	-0.02	0.2	-0.2

Balance and Status: Observations

- Both theories agree well with learned models

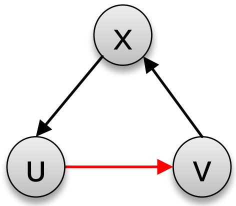

- Further observations:
  - Backward-backward triads have smaller weights than forward and mixed direction triads
  - Balance is in better agreement with Epinions and Slashdot while Status is with Wikipedia
  - Balance consistently disagrees with "enemy of my enemy is my friend"

Tutorial Outline

- Part 1: Information flow in networks
- **Part 2: Rich interactions**
  - 2.1: Recommending links in networks
  - 2.2: Predicting tie strength
  - 2.3: Predicting friends vs. foes
  - **2.4: How do people evaluate others?**

People have Opinions

People express positive and negative attitudes/opinions:

- Through actions:
  - Rating a product
  - Pressing “like” button
- Through text: Sentiment analysis [Pang-Lee ‘08]
  - Writing a comment, a review

People Express Opinions

- About items:
  - Movie and product reviews

IMDb Earth's Biggest Movie Database™ amazon.com

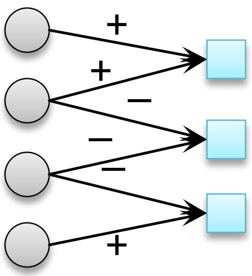

- About other users:
  - Online communities

Epinions.com WIKIPEDIA

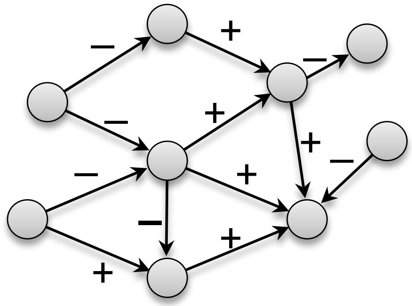

- About items created by others:
  - Q&A websites

stackoverflow YAHOO! ANSWERS

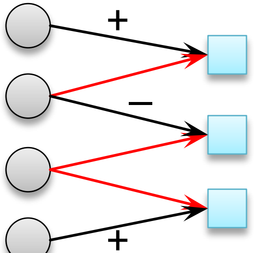

This talk: Users evaluating others

- Any user A can evaluate any user B:

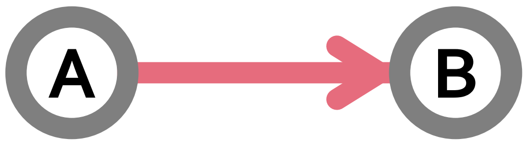

  - Positive (+) vs. negative (–) evaluation

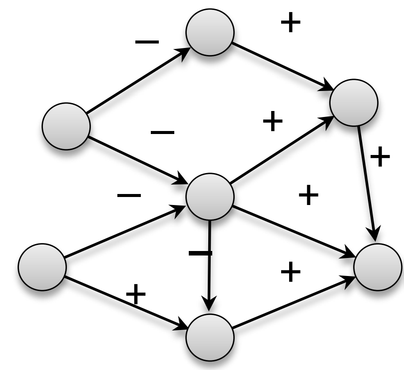

- In what (online) settings does this process naturally occur at large scale?
  - Epinions: Trust/Distrust (1M evals)
    - Does A trust B's product reviews?
  - Wikipedia: Support/Oppose (150k votes)
    - Does A support B to become Wiki admin?
  - Stackoverflow: Up/down vote (6M votes)
    - Does A think B contributed a good answer?

Relative vs. Absolute Assessment

- How do properties of **evaluator A** and **target B** affect A's vote?

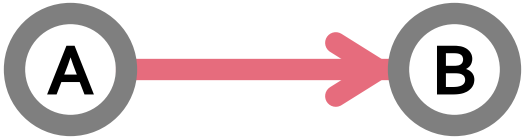

- Two natural (but competing) hypotheses:
  - (1) Prob. that B receives a positive evaluation depends primarily on the characteristics of B
    - There is some objective criteria for a user to receive a positive evaluation

Relative vs. Absolute Assessment

- How do properties of **evaluator A** and **target B** affect A's vote?

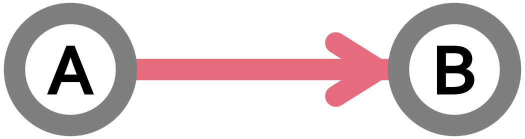

- Two natural (but competing) hypotheses:
  - (2) Prob. that B receives a positive evaluation depends on relationship between characteristics of A and B
    - **Similarity:** Prior interaction between A and B
    - **Status:** A compares status of B to her own status

Status (level of contribution)

Ways to quantify status (seniority, merit) of a user:

- Total number of edits of a user:
  - The more edits the user made the higher status she has
- Total number of answers of a user:
  - The more answers given by the user the higher status she has

Status: How to model?

- How does the prob. of A evaluating positively depend on the status of A and status of B?

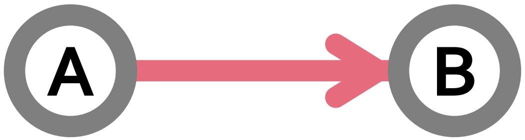

  - Model it as a function of status S_A of A and S_B of B separately?
  - Model as the status difference S_A-S_B?
  - Model as the status ratio S_A/S_B?

Status: Relative Assessment (1)

- How does status of B affect A's evaluation?
  - Each curve is fixed status difference: $\Delta = S_A - S_B$
- Observations:
  - Flat curves: Prob. of positive evaluation doesn't depend on B's status
  - Different levels: Different values of Δ result in different behavior

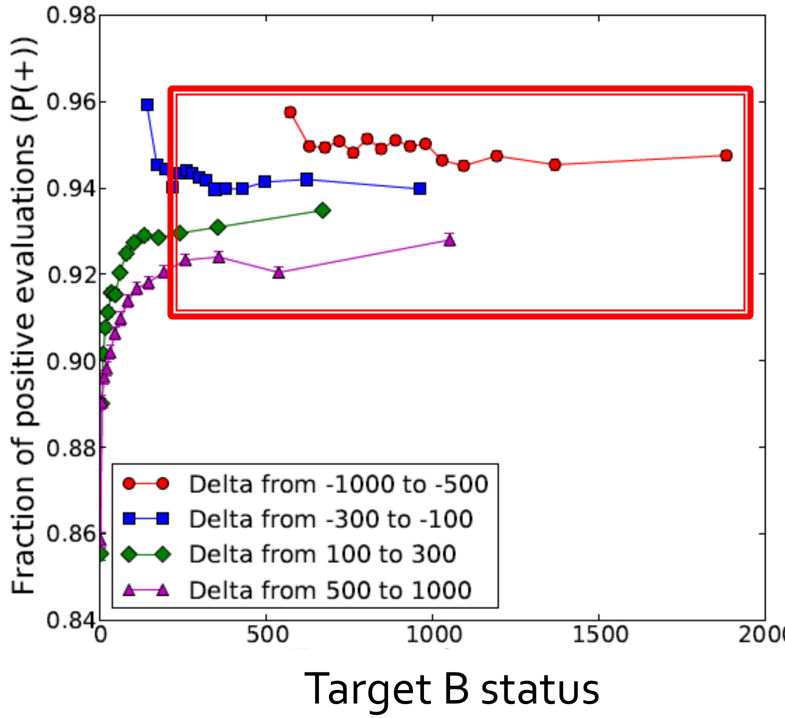

Status difference remains salient even as A and B acquire more status

Status: Relative Assessment (2)

- How does status of B affect A's evaluation?
  - Each curve is fixed status difference: $\Delta = S_A - S_B$
- Observations:
  - Below some threshold targets are judged based on their absolute status
    - And independently of evaluator's status

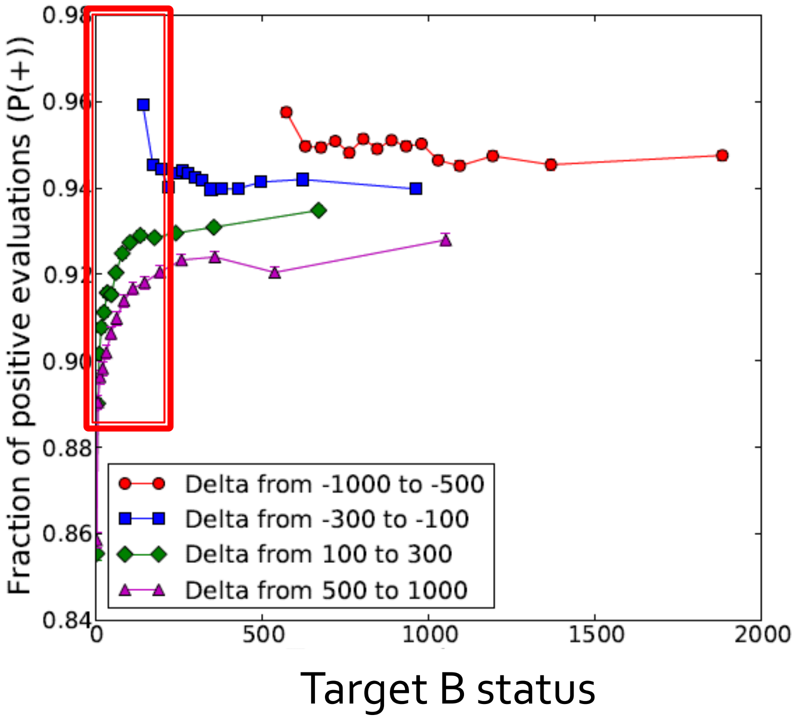

Low-status targets are evaluated based on absolute status

Effects of Similarity

- How does prior interaction shape evaluations?
  - (1) Evaluators are more supportive of targets in their area
  - (2) More familiar evaluators know weaknesses and are more harsh
- Observation:
  - Prior interaction/similarity increases prob. of a positive evaluation

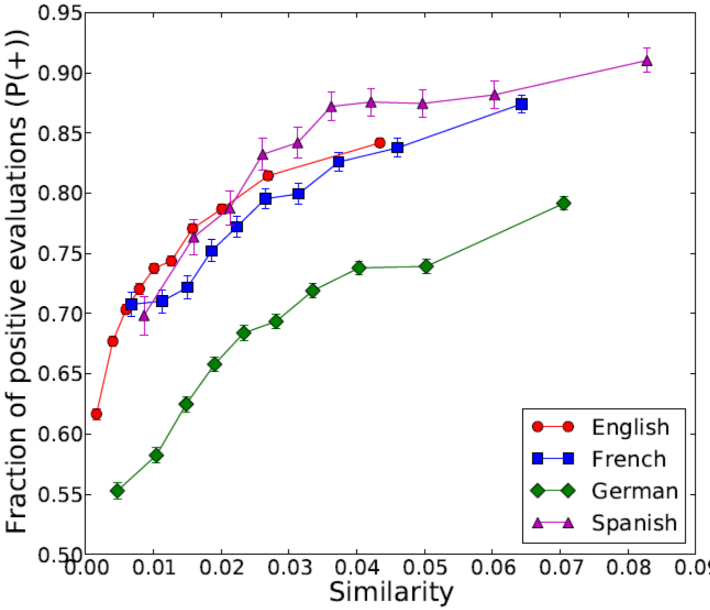

Prior interaction/ similarity boosts positive evaluations

Relating Status and Similarity (1)

- Observation:
  - Evaluation depends less on status when evaluator A is more informed
- Consequence:
  - Evaluators use status as proxy for quality in the absence of direct knowledge of B

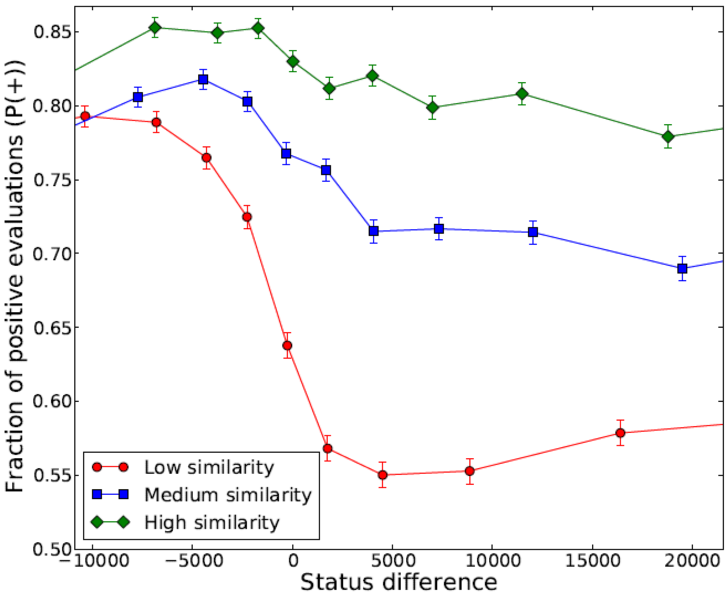

Status is a proxy for quality when evaluator does not know the target

Relating Status and Similarity (2)

- Observation:
  - Evaluators with higher status than the target are more similar to the target

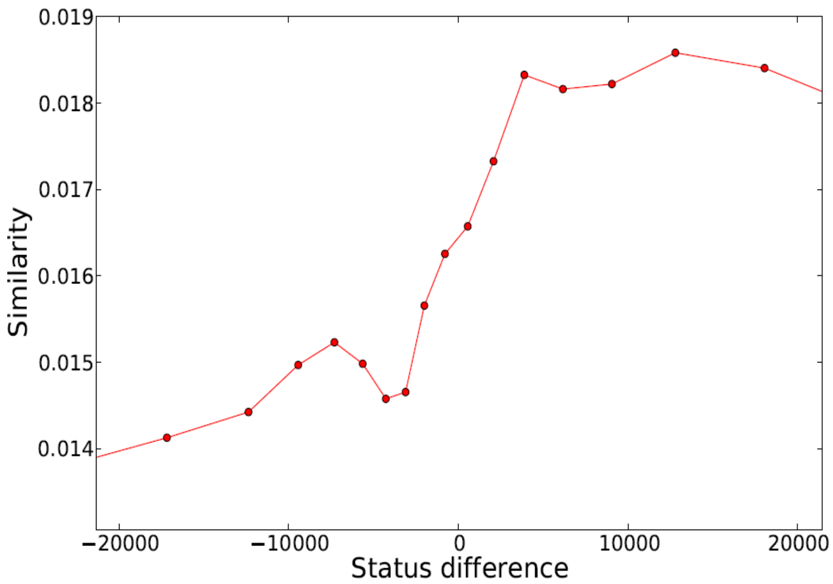

- Selection bias:
  - High-status evaluators are more similar to the target

Elite evaluators vote on targets in their area of expertise

Puzzle: Status

- Evaluator A evaluates target B
- Prob. of positive evaluation of A as a function of status difference: $\Delta = S_A - S_B$
  - Hypothesis: Monotonically decreases

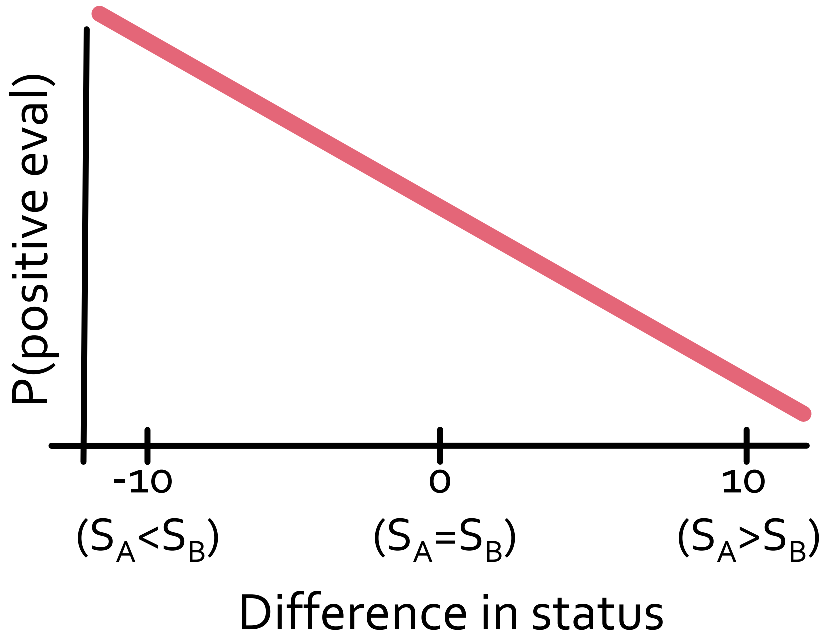

Puzzle: Status

- Prob. of positive evaluation of B as a function of status difference: $\Delta = S_A - S_B$
- Observations:
  - A is especially negative when status equals: $S_A=S_B$
  - "Mercy bounce" for $S_A>S_B$

How to explain the bounce?

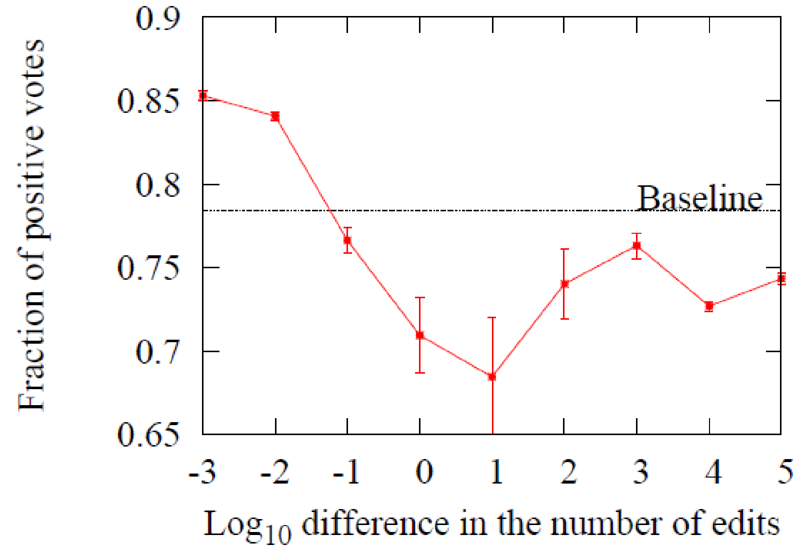

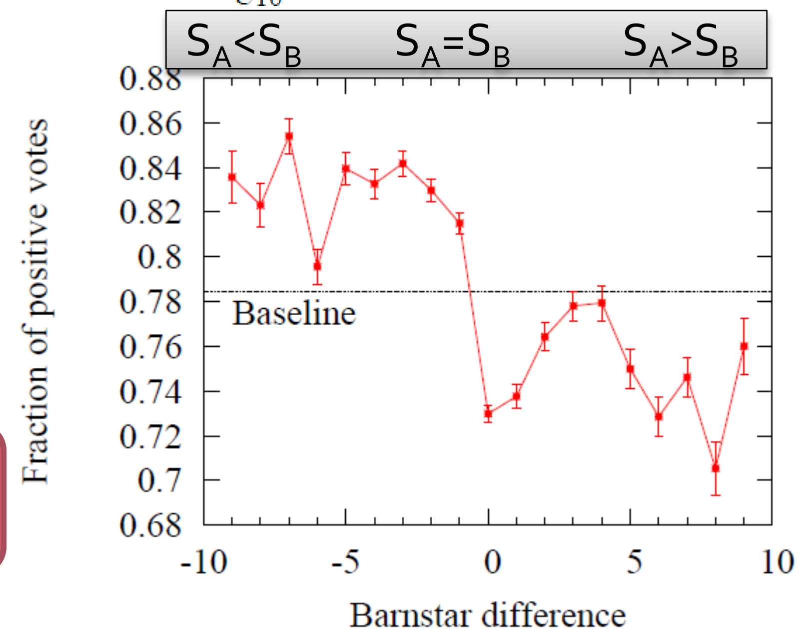

Why most harsh at zero difference?

How to explain low aggregate evaluations given by users to others of same status?

- Not due to users being tough on each other
  - Similarity increases the positivity of evaluations

Possible (but wrong) explanation:

- Most targets have low status (small $\Delta > 0$)
- Low-status targets are judged on abs. status
  - The rebound persists even for high-status targets

Important points

- Social media sites are governed by (often implicit) user evaluations
- Wikipedia voting process has an explicit, public and recorded process of evaluation
- Main characteristics:
  - Importance of relative assessment: **Status**
  - Importance of prior interaction: **Similarity**
  - Diversity of individuals' response functions

- Application: Ballot-blind prediction

Ballot-blind prediction

- Predict Wikipedia elections without seeing the votes
  - Observe identities of the first k(=5) people voting (but *not* how they voted)
  - Want to predict the election outcome (promotion/no promotion)

  - Why is it hard?
    - Don't see the votes (just voters)
    - Only see first 5 voters (10% of the election)

Ballot-blind: the Model

- Idea: Split the status-similarity space (s,Δ) in to 4 quadrants
- Model deviation in voter's behavior when they evaluate a candidate from a particular quadrant:
  - d(s,Δ) … avg. deviation in fraction of positive votes
    - When voters evaluate a candidate C from a particular (s,Δ) quadrant, how does this change their behavior

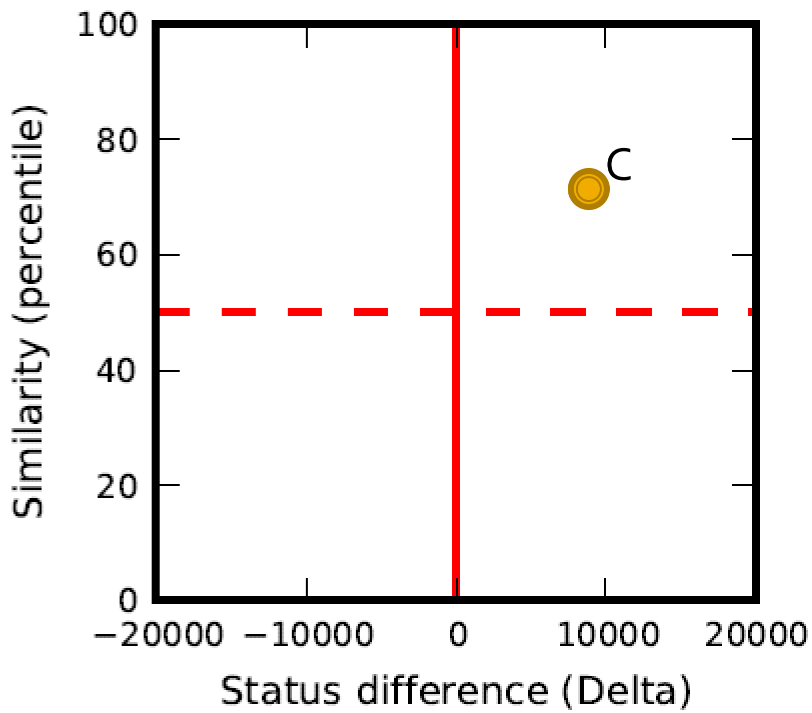

Ballot-blind: the Model

- d(s,Δ) … signed deviation in the fraction of positive votes when *E* evaluates *C* of similarity s and status difference Δ

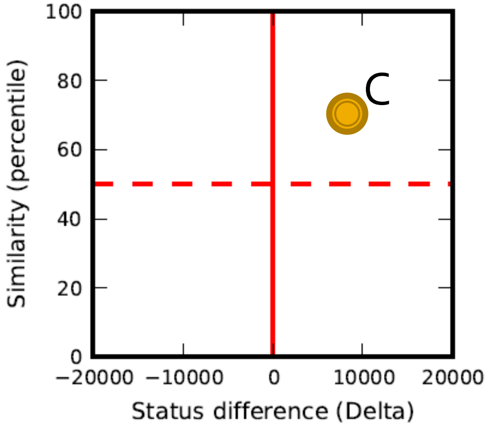

  - $P(E_i=1)$ … prob. evaluator *E* votes + in election *i*
- The models:
  - Global **M1**: $P(E_i = 1) = P_i + d(\Delta_i, s_i)$
  - Personal **M2**:

$$P(E_i = 1) = \alpha \cdot P_i(\Delta_i, s_i) + (1 - \alpha) \cdot d(\Delta_i, s_i)$$

where P_i is empirical frac. of + votes of E

Results: Wikipedia

- Predictive accuracy of baselines:
  - Guessing: 52%
  - If we know votes: 85%
  - Bag-of-features **B1**: 69%
- Model based on status and similarity:
    - Does not see votes
    - Sees only first 5 votes (10% of the lection)
  - Global model **M1**: 76%
  - Personal model **M2**: 75%

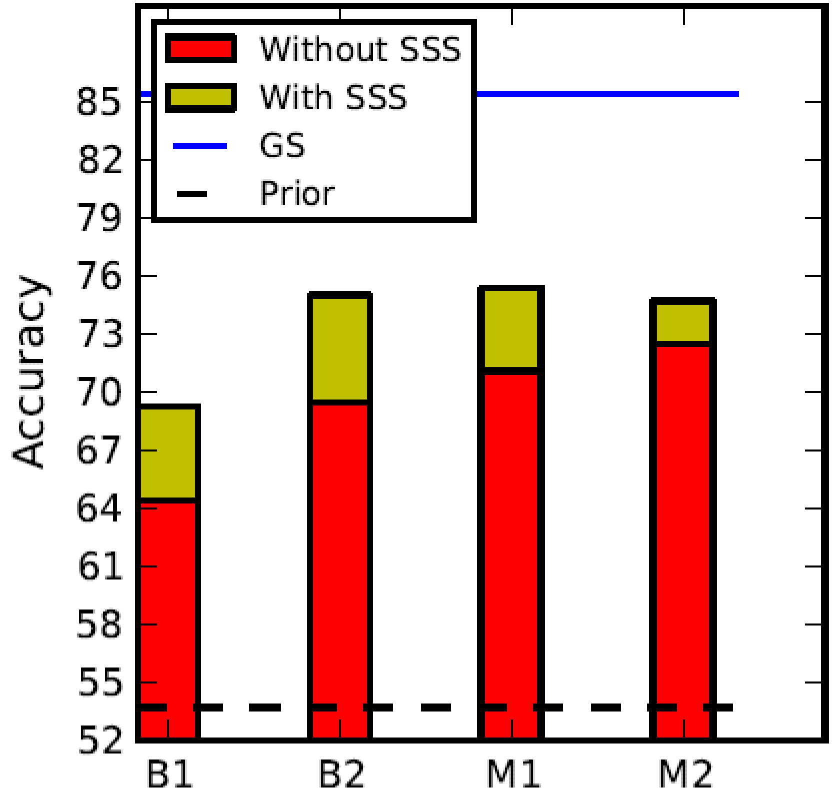

Audience composition predict audience's reaction

Conclusion and reflections

- Online social systems are globally organized based on status
- Users use evaluations consistently regardless of a particular application
  - Near perfect generalization across datasets
- Audience composition helps predict audience's reaction
- What kinds of opinions do people find helpful?

What do people find helpful?

- What do people think about our recommendations and opinions?

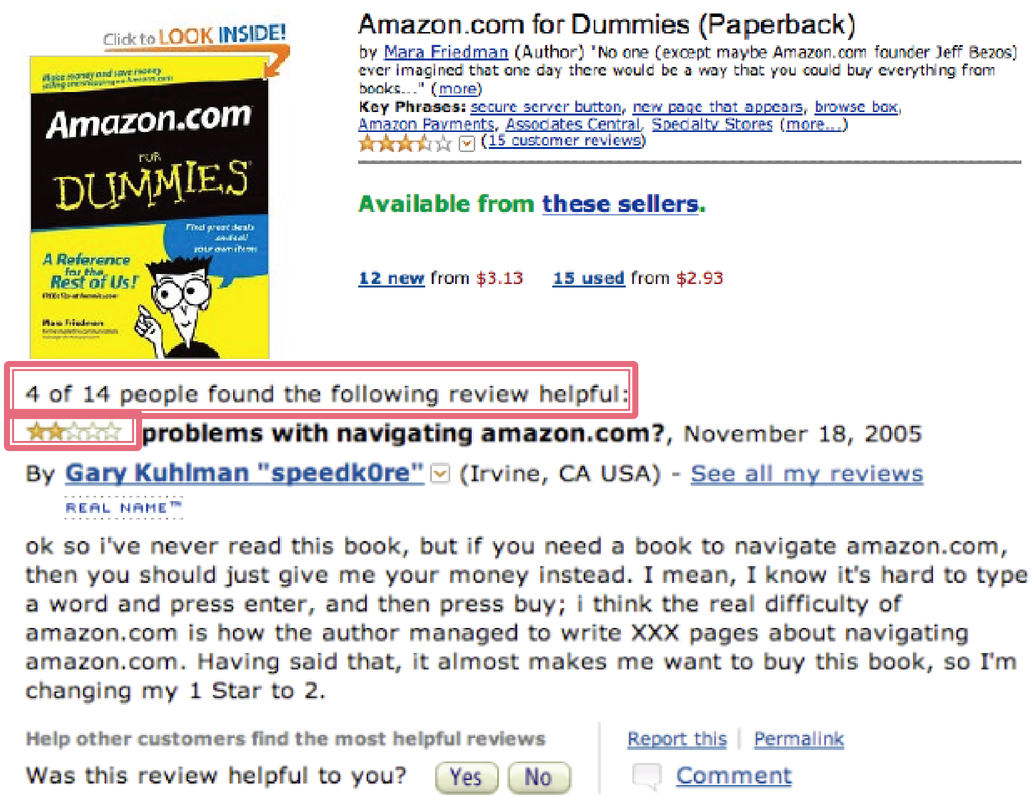

Review helpfulness: Conformity

- People find conforming opinions more helpful

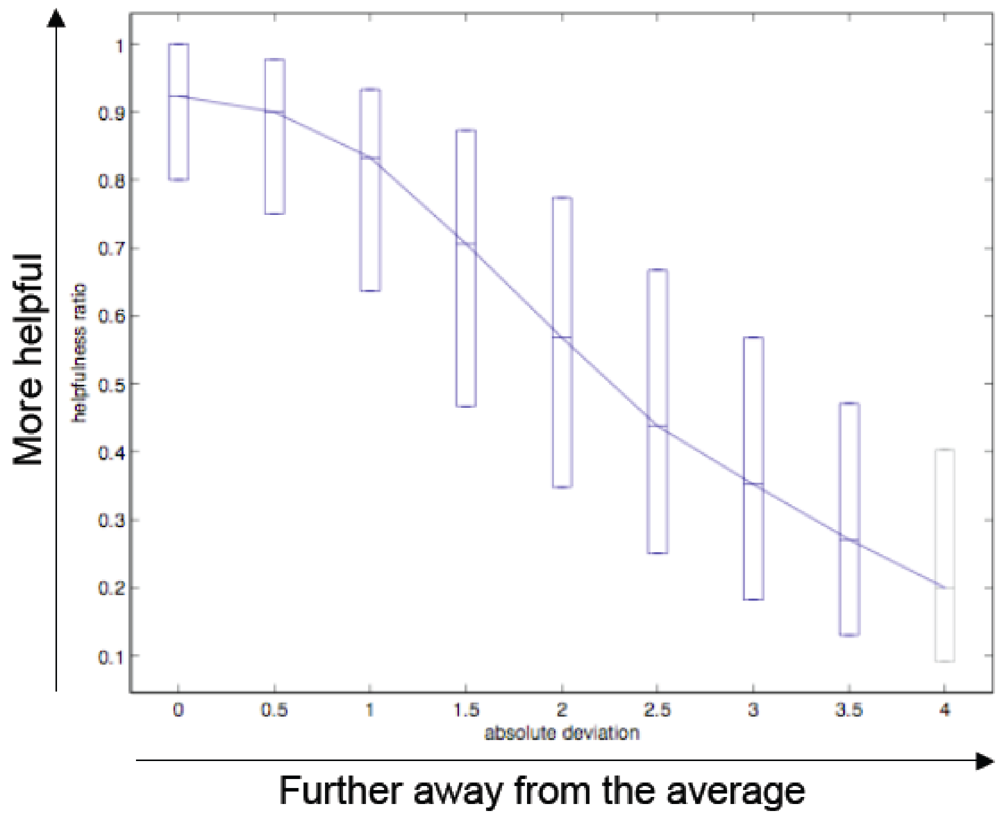

Review helpfulness: Deviation

- Positive reviews are more helpful

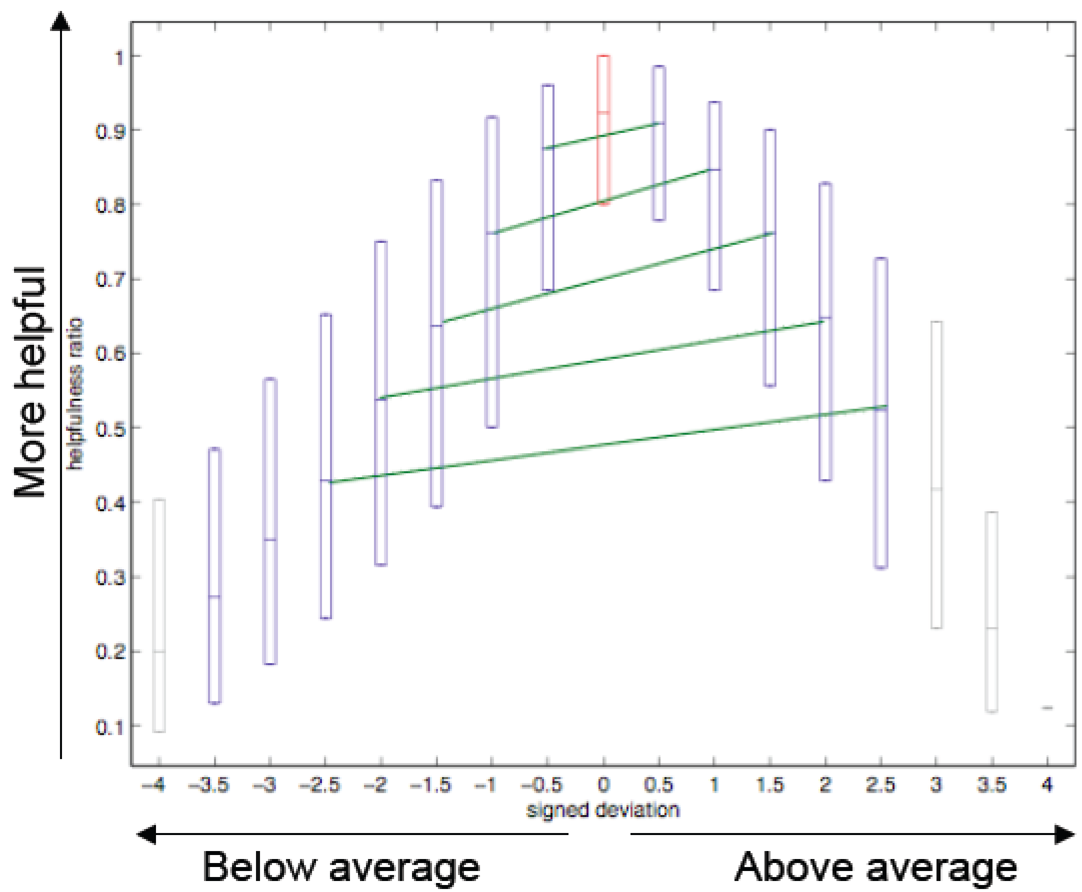

Tutorial Outline

- Part 1: Information flow in networks
- Part 2: Rich interactions
- **Conclusion and reflections**

Part 1: Information flow

- Messages arriving through networks from real-time sources requires new ways of thinking about information dynamics and consumption
- "Tell me about X" vs.
  "Tell me what I need to know now"

Part 1: Conclusion and Connections

- Diffusion of Topics and Sentiment
  - How news cascade through on-line networks
  - Do we need new notions of rank/importance?
- Incentives and Diffusion
  - Using diffusion in the design of on-line systems
  - Connections to game theory
- When will one cascade overtake the other?

Part 1: Opportunities

- A number of novel opportunities:
  - Predictive modeling of the spread of new ideas and behaviors
  - Opportunity to design systems that make use of diffusion process
- Applications:
  - Search
  - Real-time search
  - Social search

Part 2: Rich Interactions

- Links are more than just links
  - Strengths
  - Sentiment
  - They reveal what we think of others

- Main characteristics:
  - Importance of relative assessment: **Status**
  - Importance of prior interaction: **Similarity**

Part 2: Connections

- Don’t predict just who we link to but also what we think of them
- Evaluations range from evaluating a person to the content they produced
- Different dimensions of the evaluation:
  - Is the content technically correct?
  - Do I agree/disagree with the answer?

Part 2: Opportunities

- Composition of an audience can tell us something about the audience's reaction
  - Predict outcomes simply from the statuses and similarities of the users who show up to provide evaluations, without ever seeing the values of the evaluations themselves
  - Connections to collaborative filtering
- Design reputation systems that account for status and similarity and encourage interaction

The Road Ahead

Strengths	Weaknesses
▪ free form facilitates capturing the true voice of customer, wisdom of crowd ▪ can be expressed through voice, text messaging on mobile phones, etc.	▪ language analysis and mining are challenging ▪ susceptible to spam, self-serving use by companies ▪ Behavior, predictive models need more research
Threats	**Opportunities**
▪ privacy and security issues: possible to assimilate detailed knowledge about person's activities, whereabouts ▪ can lead to anti-social behavior!	▪ promise of collective problem solving: coordination, cooperation ▪ mobile use supports dealing with societal problems, disaster situations: social network is geospatial proximity