Tutorial at the World Wide Web conference
Hyderabad, India, 2011
http://snap.stanford.edu/proj/socmedia-www

Analytics & Predictive Models for Social Media:

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

[LibenNowell-Kleinberg '03]

Link prediction via node distance

- Predict links evolving collaboration network

	training period			Core		
	authors	papers	collaborations[1]	authors	$\|E_{old}\|$	$\|E_{new}\|$
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 $\mathsf{paths}_{x,y}^{\langle\ell\rangle} := \{\text{paths of length exactly } \ell \text{ from } x \text{ to } y\}$

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

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

Γ(x) … degree of node *x*

[LibenNowell-Kleinberg '03]

Link prediction via proximity

- Rank potential links (x,y) based on:

hitting time	$-H_{x,y}$
stationary-normed	$-H_{x,y} \cdot \pi_y$
commute time	$-(H_{x,y} + H_{y,x})$
stationary-normed	$-(H_{x,y} \cdot \pi_y + H_{y,x} \cdot \pi_x)$
	where $H_{x,y}$:= expected time for random walk from x to reach y
	π_y := stationary distribution weight of y (proportion of time the random walk is at node y)
rooted PageRank$_\alpha$	stationary distribution weight of y under the following random walk: with probability α, jump to x. with probability $1-\alpha$, go to random neighbor of current node.
SimRank$_\gamma$	$\begin{cases} 1 & \text{if } x = y \\ \gamma \cdot \frac{\sum_{a \in \Gamma(x)} \sum_{b \in \Gamma(y)} \text{score}(a,b)}{\lvert\Gamma(x)\rvert \cdot \lvert\Gamma(y)\rvert} & \text{otherwise} \end{cases}$

$\Gamma(x)$ … degree of node x

[LibenNowell-Kleinberg ' 03]

Results: Improvement over random

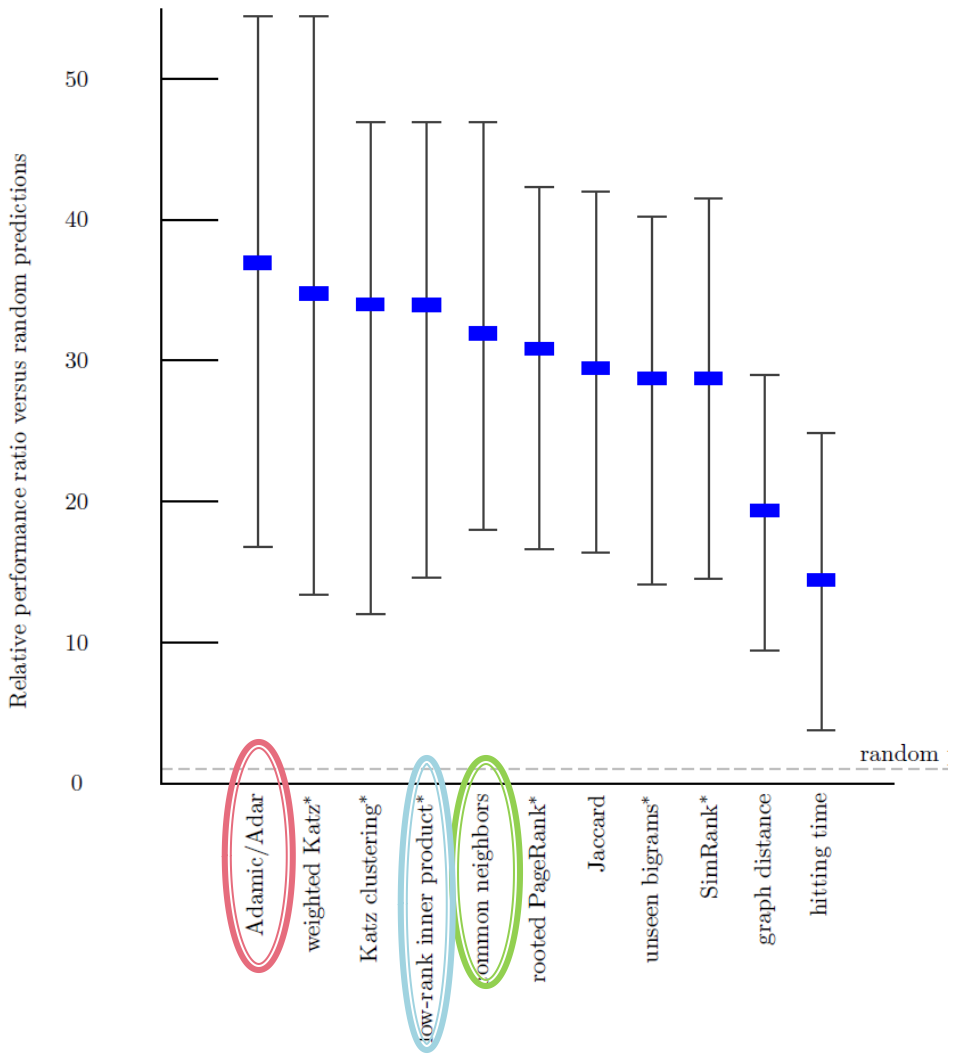

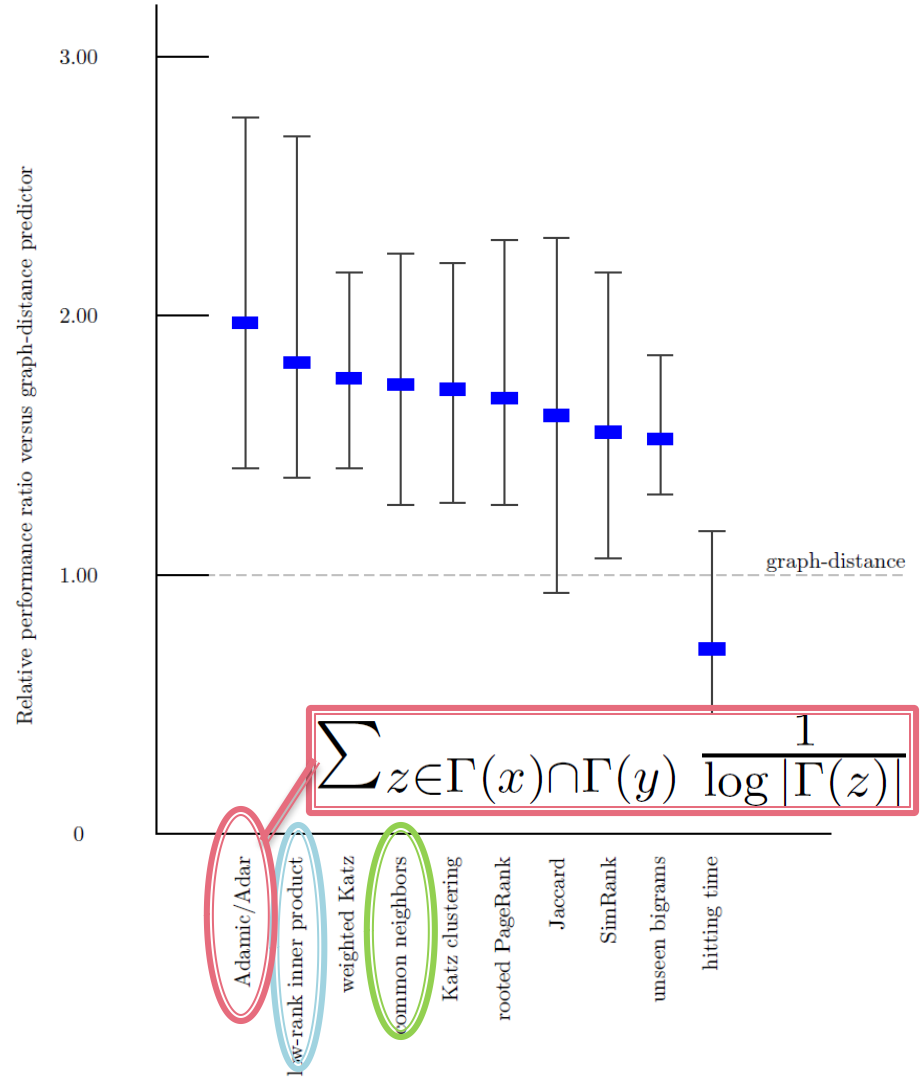

Results: Common neighbors

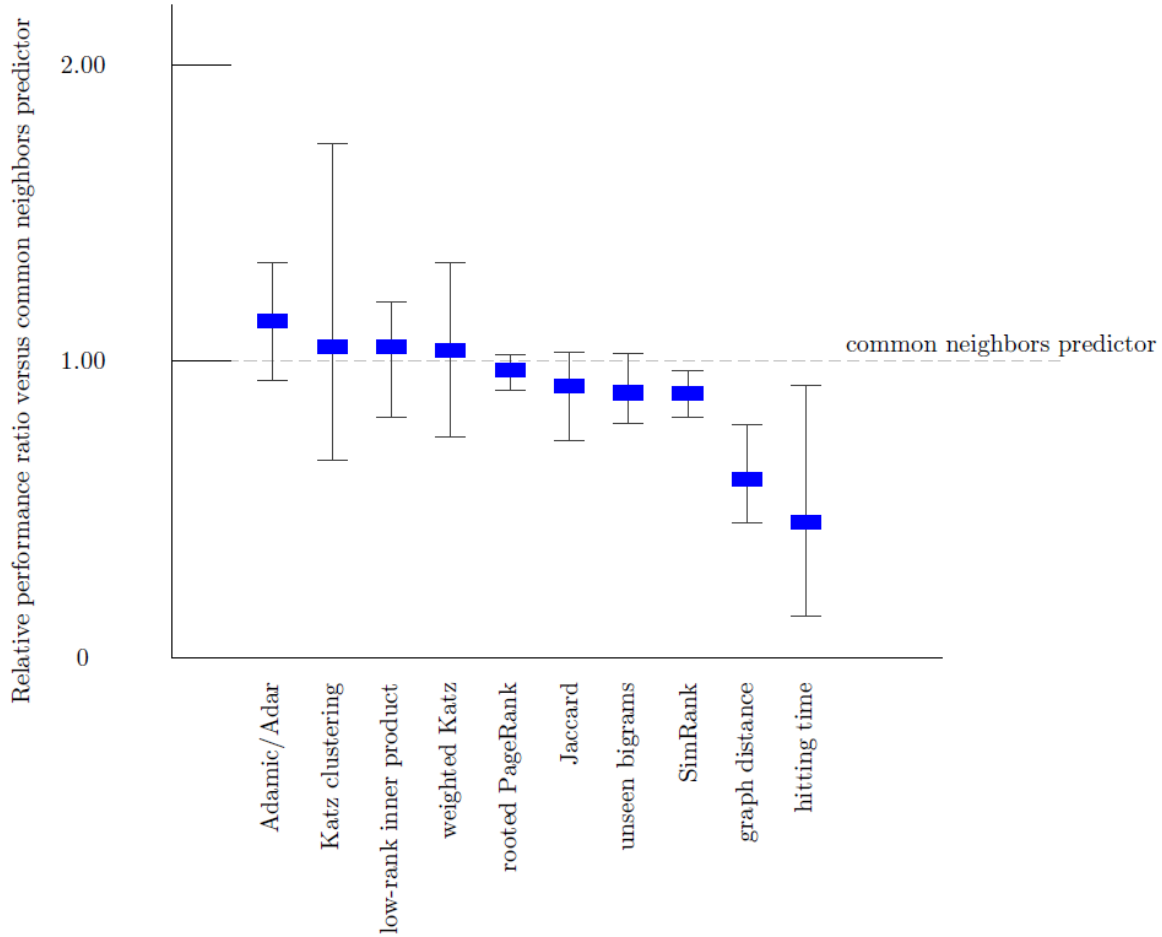

- Improvement over #common neighbors

[WSDM '11]

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

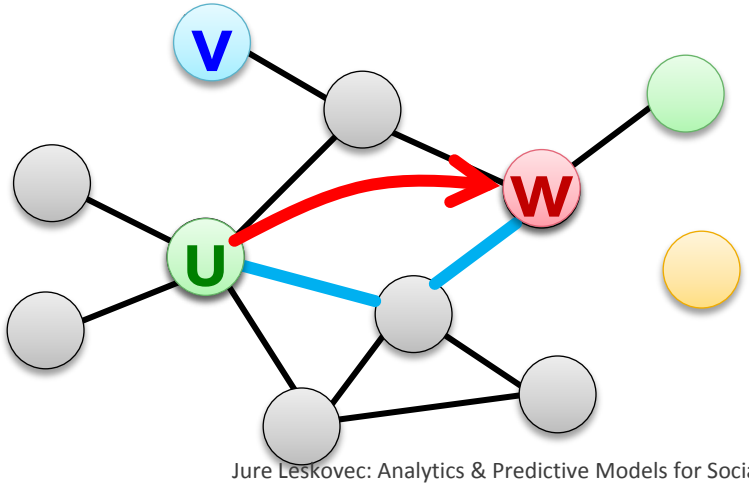

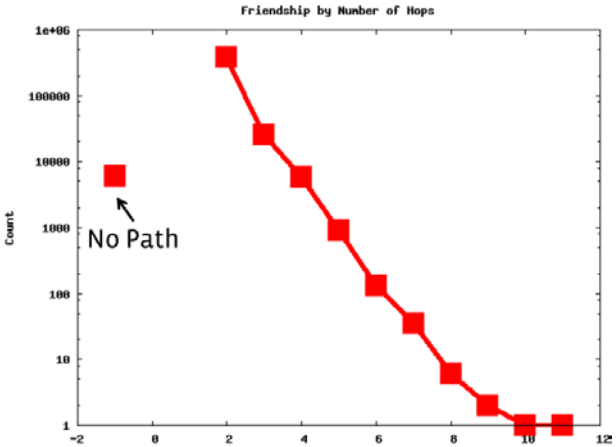

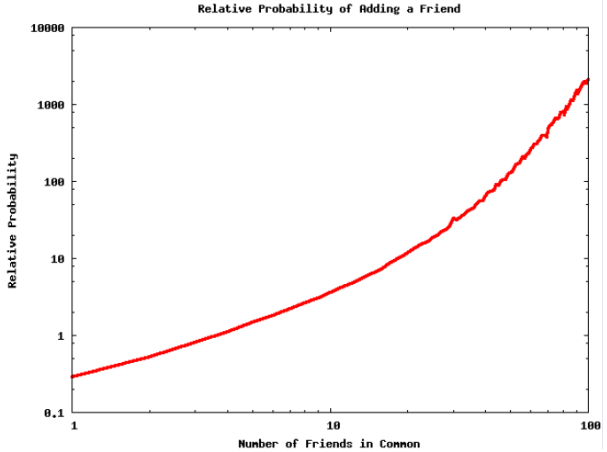

Supervised Link Prediction

- 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 $\{v_1, \ldots, v_k\}$
  - 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

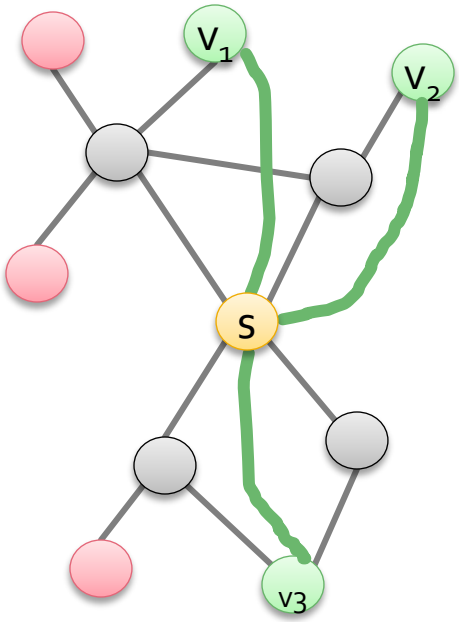

positive examples
negative examples

Supervised Link Prediction

- How to combine node/edge attributes and the network structure?
- Learn how to often visit green nodes:
  - Learn params. w of a func. $f_w(u,v)$ that assigns weights to edges (u,v) based on:
    - Profile of user u, profile of user v
    - Interaction history of u and v
  - Do a Personalized PageRank on this weighted graphs with a teleport set $\{s\}$
  - Rank nodes by their visiting prob. p_u

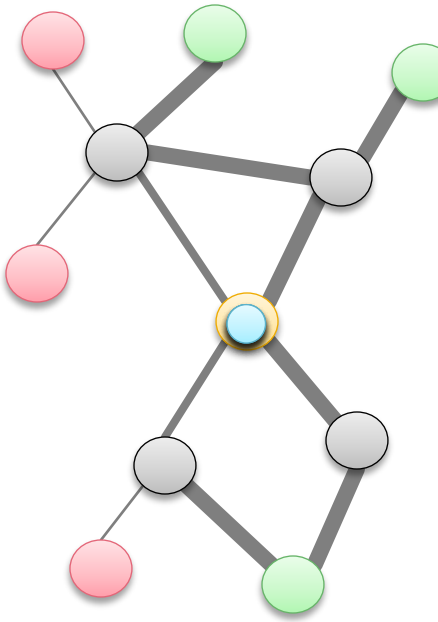

[WSDM '11]

Supervised Random Walks

- Let s be the center node
- Let $f_w(u,v)$ be a function that assigns a weight 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 a Personalized PageRank from s where transitions are according to edge weights
- How to learn $f_w(u,v)$?

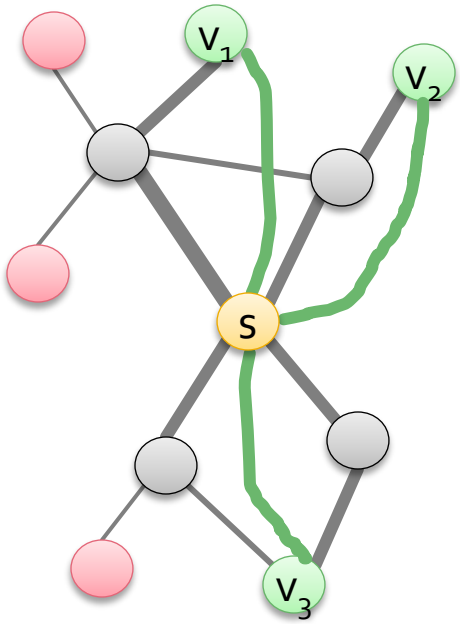

positive examples
negative examples

[WSDM '11]

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

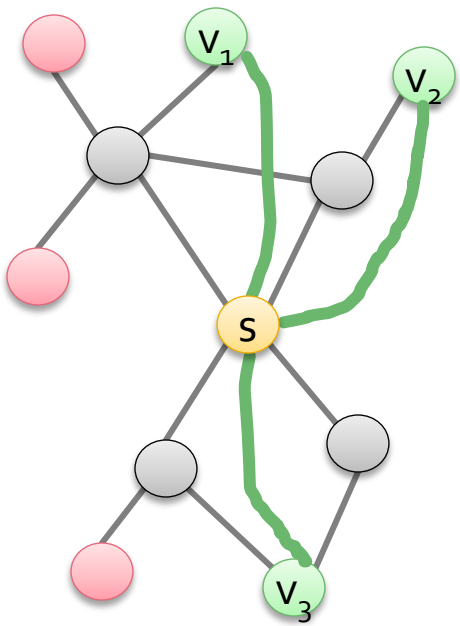

- Compute PageRank vector: $p = p^T Q$
- Rank nodes by p_u

[WSDM '11]

The Optimization Problem

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

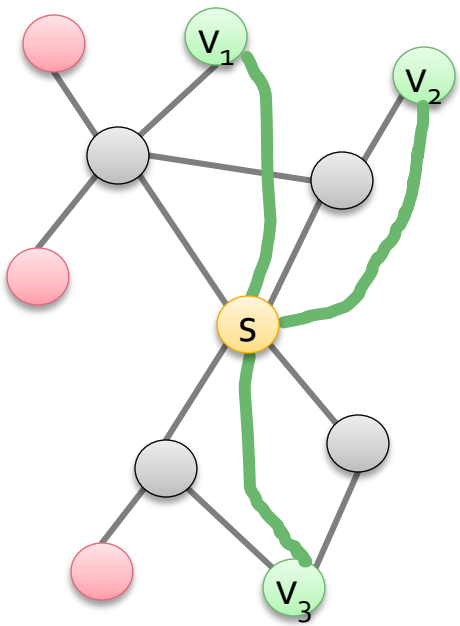

- What do we want?

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

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

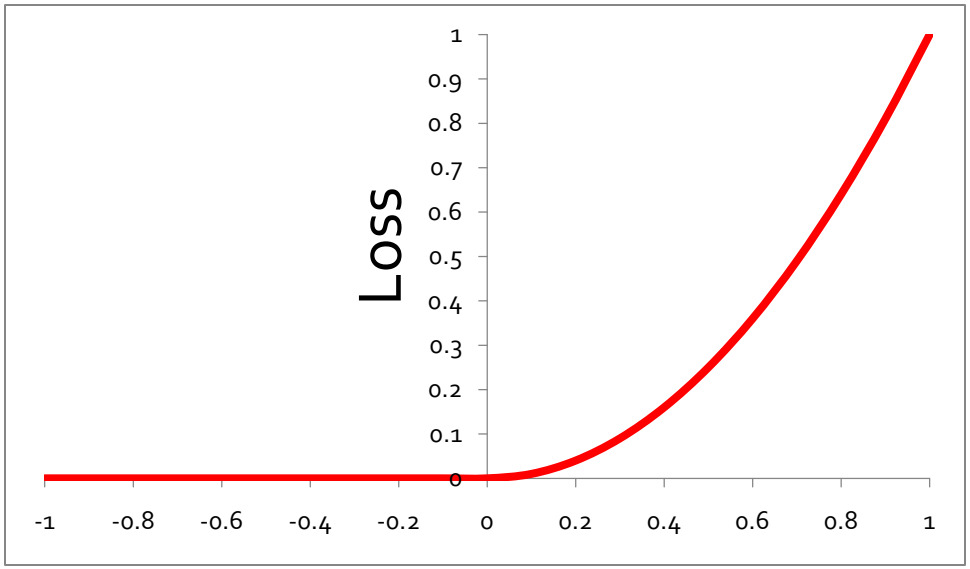

$p_l < p_d$ $\quad$ $p_l = p_d$ $\quad$ $p_l > p_d$

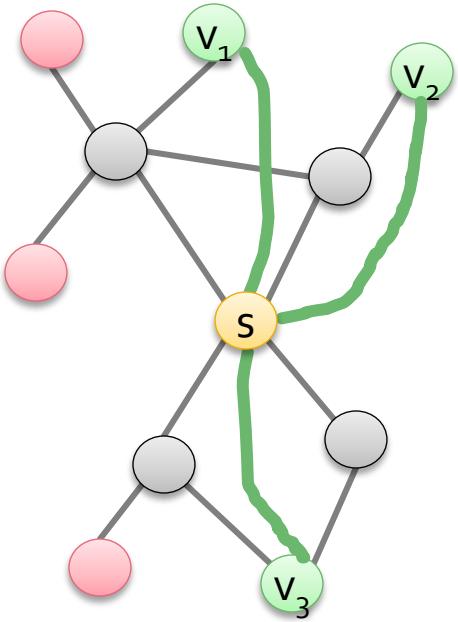

[WSDM '11]

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 set w such that $p_l < p_d$

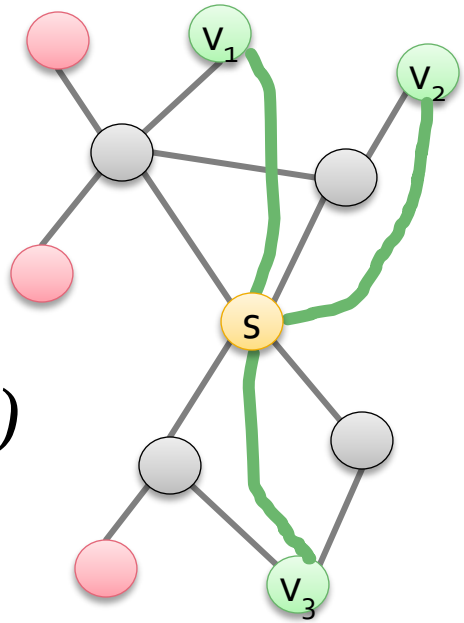

Gradient Descent

- How to minimize F?
  - Take the derivative!

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

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

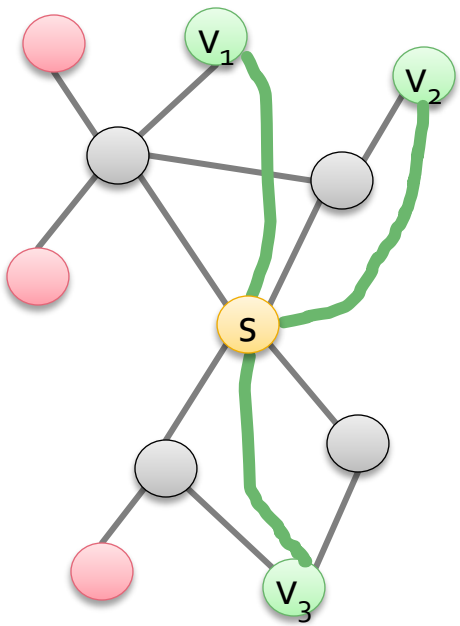

- We know:

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

- 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
  - Updated w
    - Optimize using quasi-Newton method

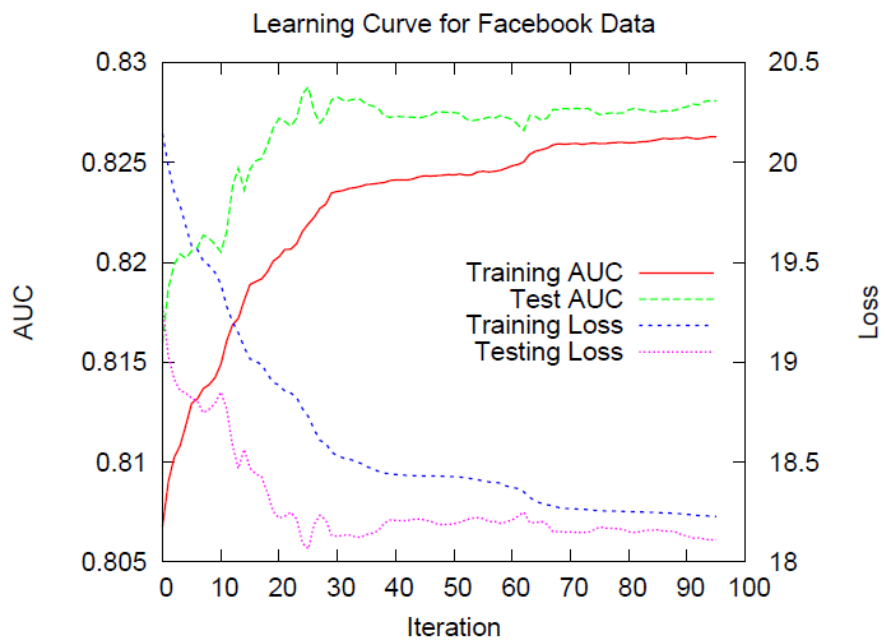

[WSDM '11]

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)!

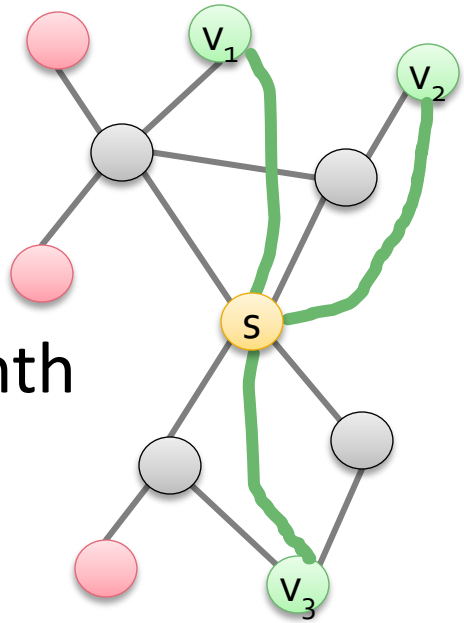

Experimental setting

- 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**

Results: Co-authorship

- Arxiv Hep-Ph collaboration network:
  - Poor performance of unsupervised methods
  - Logistic regression and decision trees don't work to well
  - SRW gives 10% boos in Prec@20

Learning Method	AUC	Prec@20
Random Walk with Restart	0.63831	3.41
Adamic-Adar	0.60570	3.13
Common Friends	0.59370	3.11
Degree	0.56522	3.05
DT: Node features	0.60961	3.54
DT: Network features	0.59302	3.69
DT: Node+Network	0.63711	3.95
DT: Path features	0.56213	1.72
DT: All features	0.61820	3.77
LR: Node features	0.64754	3.19
LR: Network features	0.58732	3.27
LR: Node+Network	0.64644	3.81
LR: Path features	0.67237	2.78
LR: All features	0.67426	3.82
SRW: one edge type	0.69996	4.24
SRW: multiple edge types	**0.71238**	**4.25**

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?

[Gilbert-Karahalios, '09]

Tie strength

- *"The strength of a tie is a (probably linear) combination of the amount of TIME, the emotional INTENSITY, the INTIMACY (mutual confiding), 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

[Gilbert-Karahalios, '09]

Five aspects of tie strength

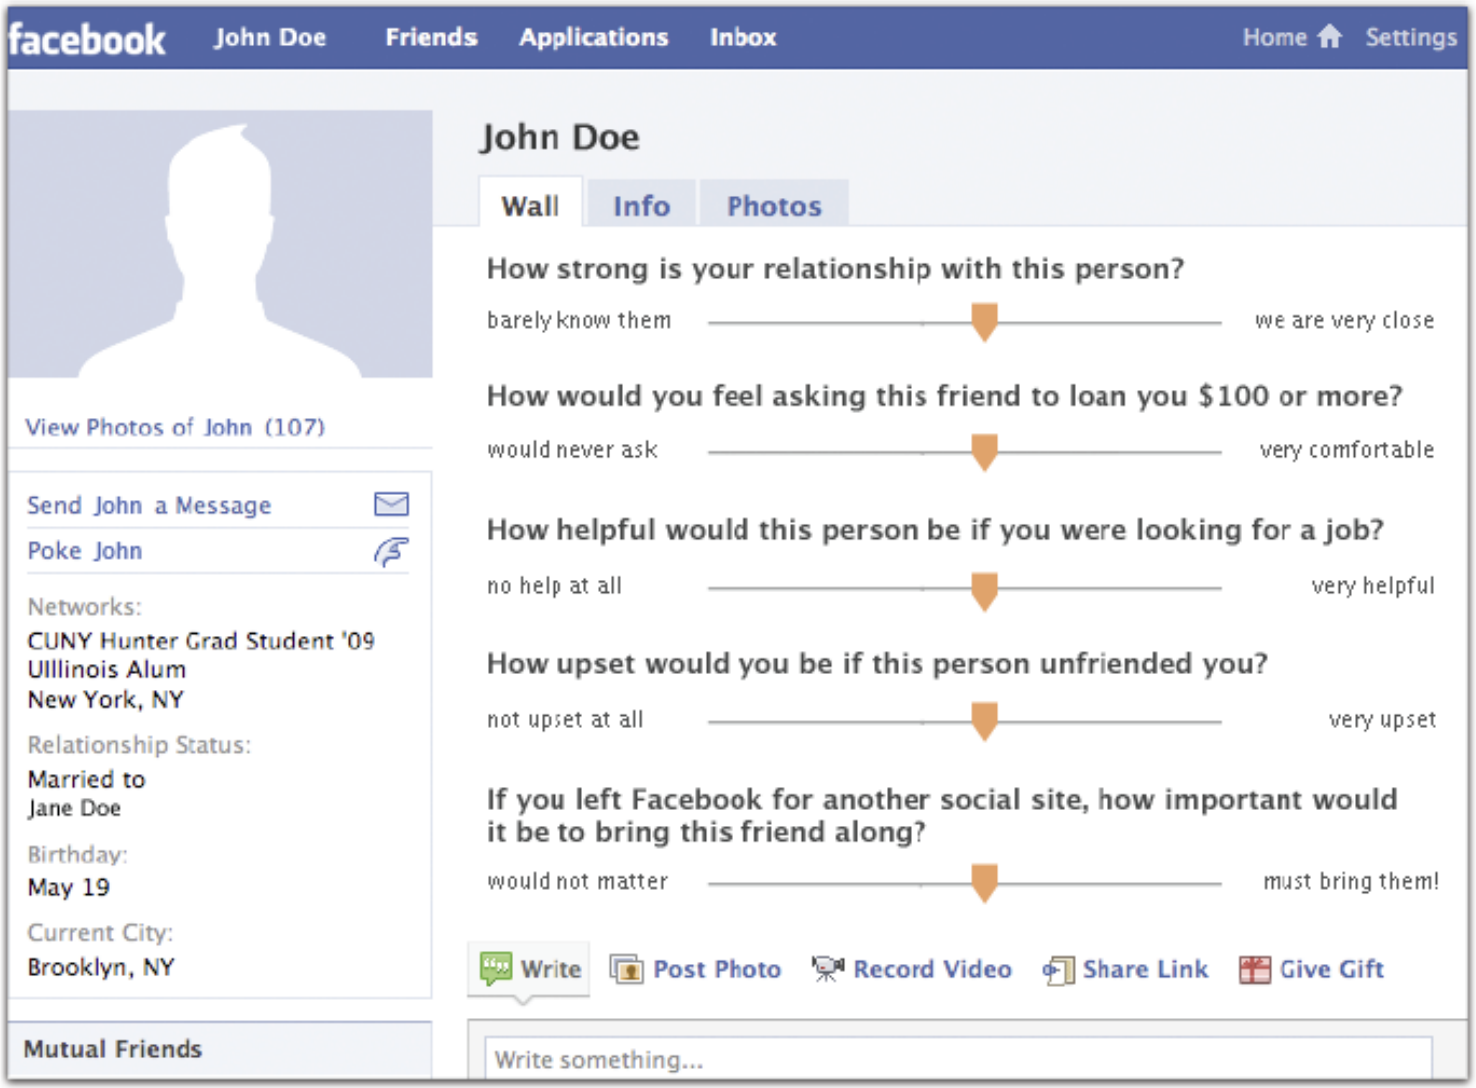

[Gilbert-Karahalios, '09]

Five aspects of tie strength

Dependent Variables	0–1 Scale	Median
How strong is your relationship?	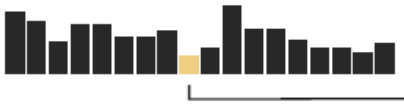	0.411
How comfortable asking for loan?		0.076
How helpful if looking for job?		0.362
How upset if unfriended?	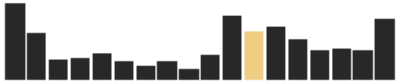	0.552
How important to bring friend?	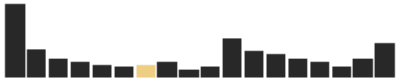	0.324

- Survey results for the 2,184 Facebook friendship links

[Gilbert-Karahalios, '09]

Attributes of the friendship

- 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

- 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

[Gilbert-Karahalios, '09]

Results: How strong is relationship

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

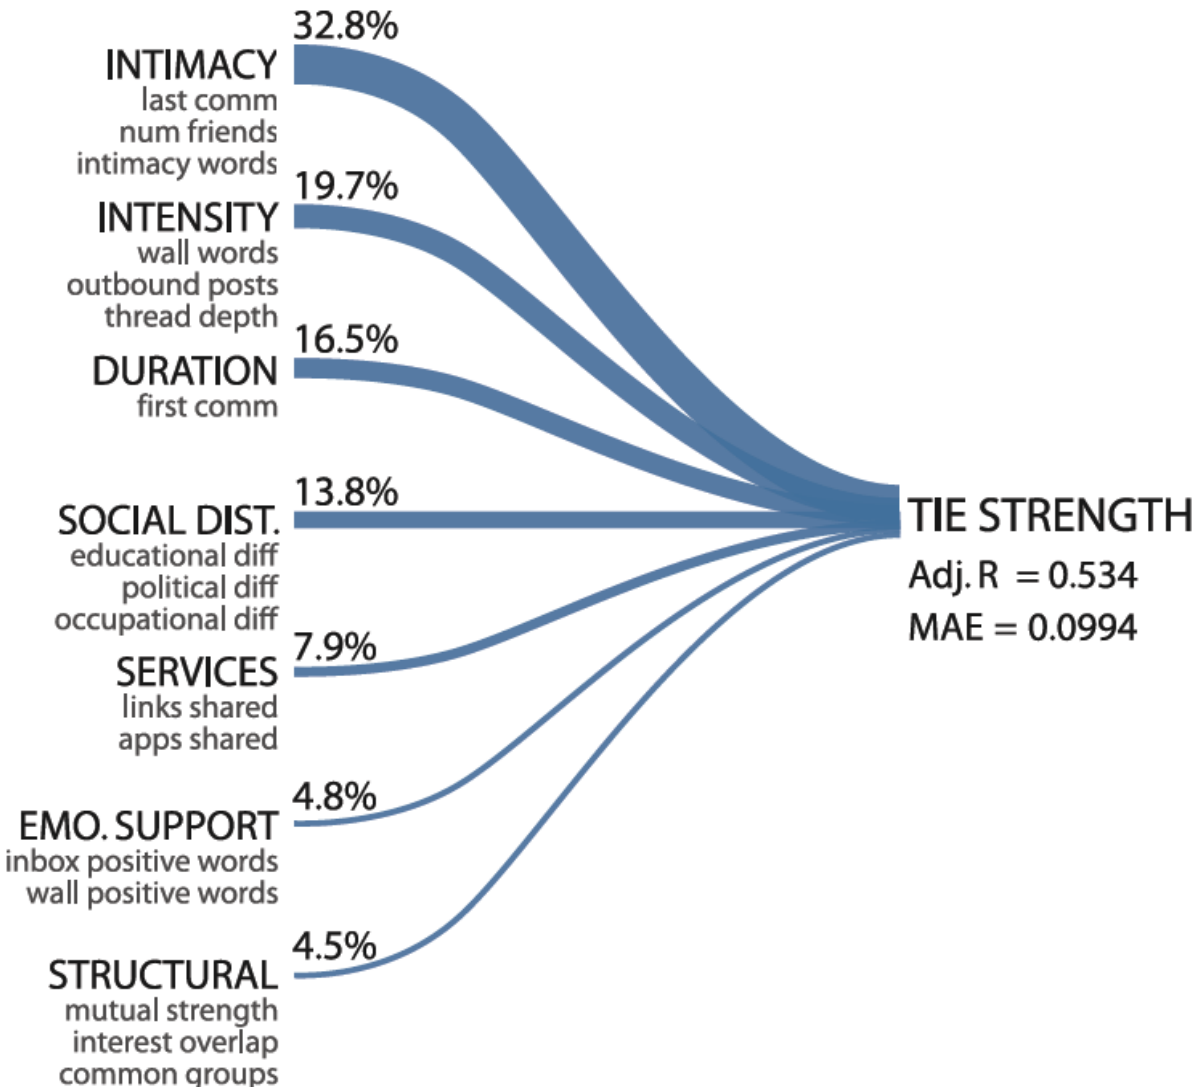

[Gilbert-Karahalios, '09]

Results: Most predictive features

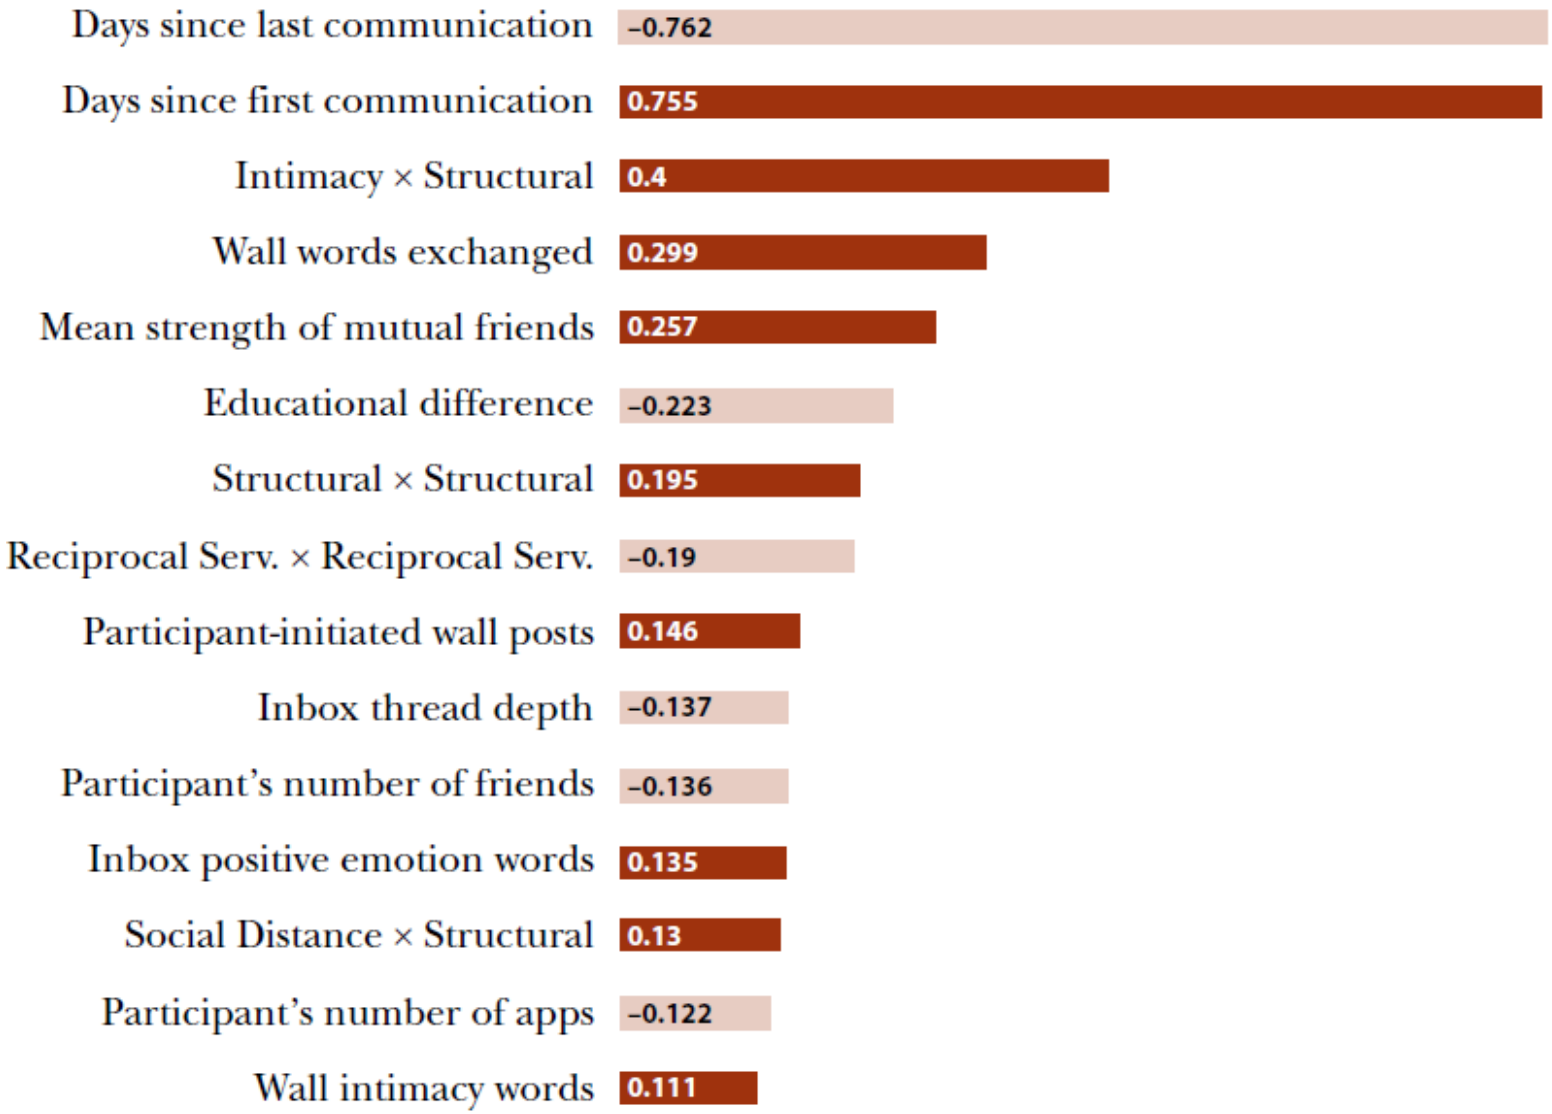

Tutorial Outline

- 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

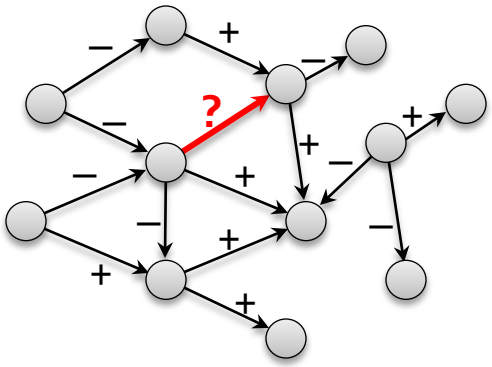

Networks with Explicit Signs

- 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?

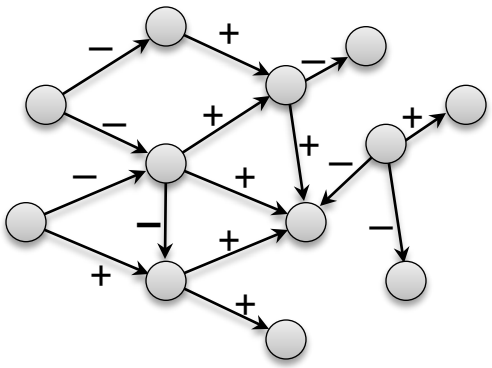

	Epinions	Slashdot	Wikipedia
Nodes	119,217	82,144	7,118
Edges	841,200	549,202	103,747
+ edges	85.0%	77.4%	78.7%
− edges	15.0%	22.6%	21.2%

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:

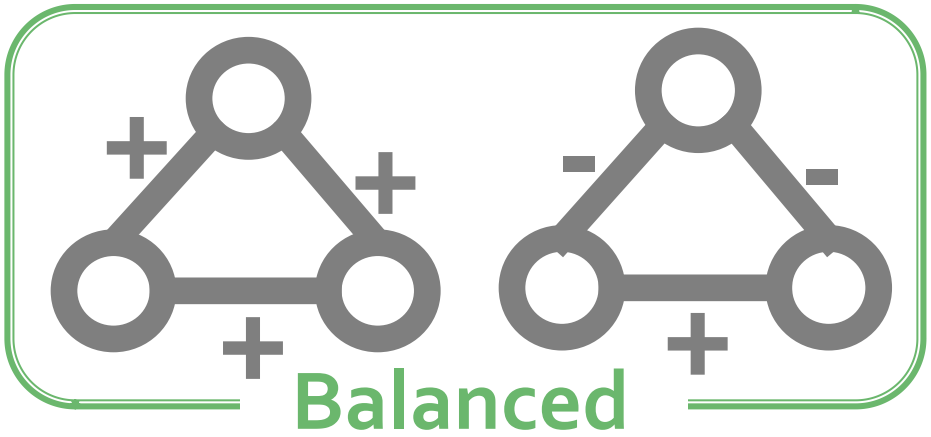

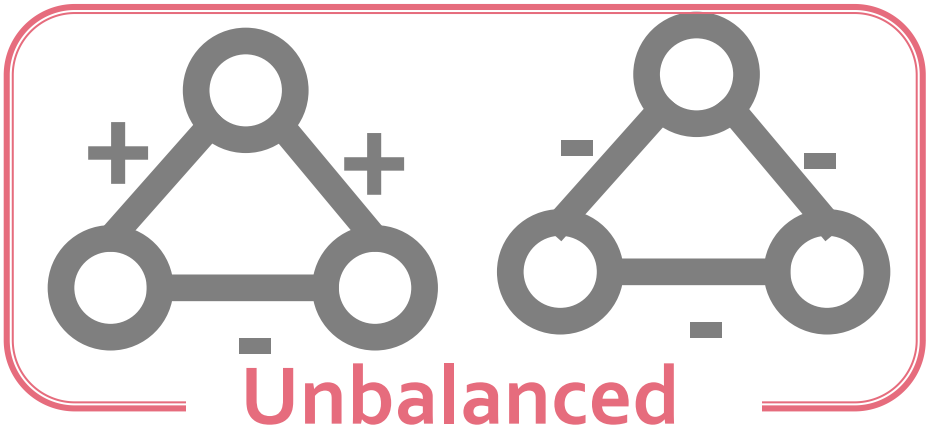

Consistent with "friend of a friend" or "enemy of the enemy" intuition

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

[CHI '10]

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:

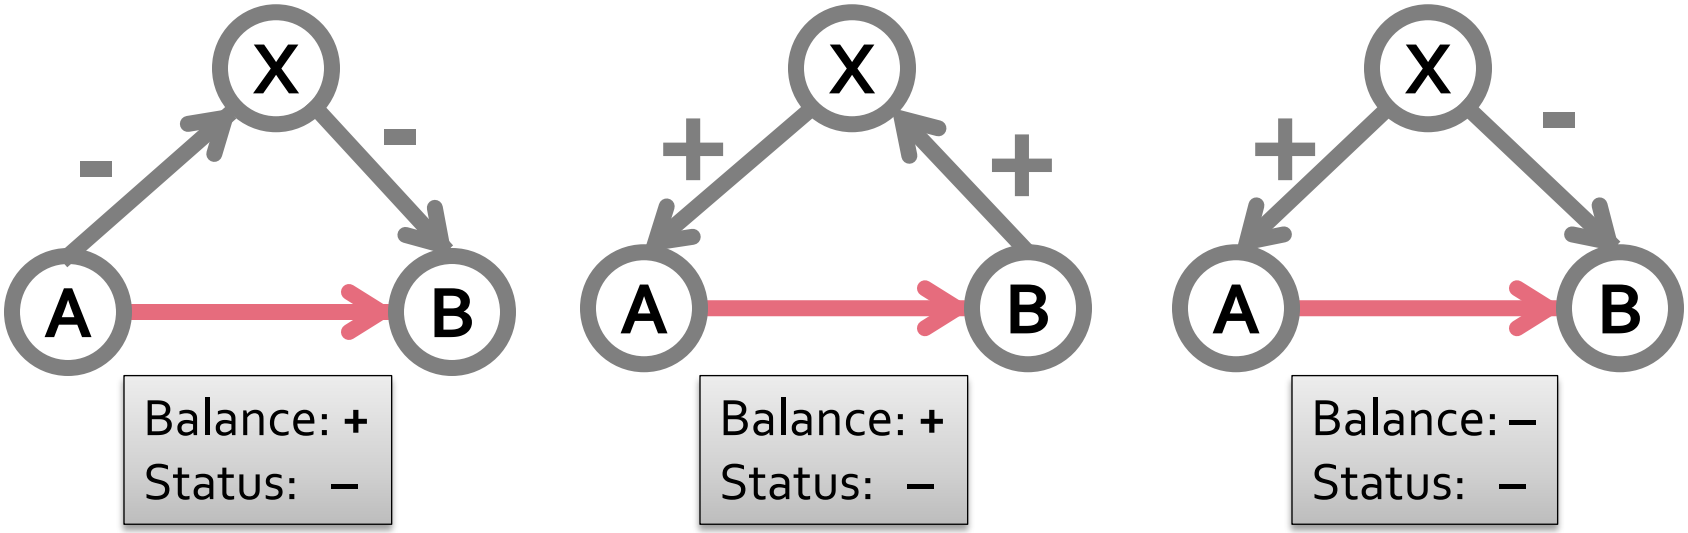

The Plan

- 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

Undirected Links: Balance

- Consider networks as undirected
- Compare frequencies of signed triads in real and shuffled data
  - 4 triad types *t*:

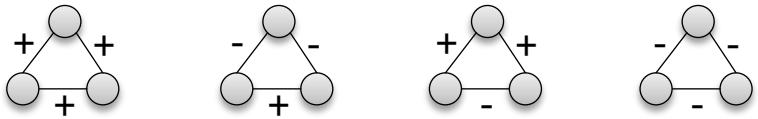

- **Surprise** value for triad type *t*:
  - Number of std. deviations by which number of occurrences of triad *t* differs from the expected number in shuffled data

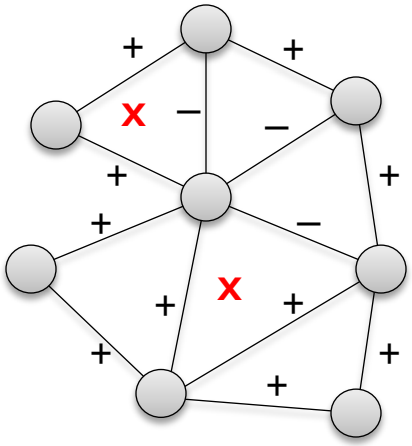

Real data

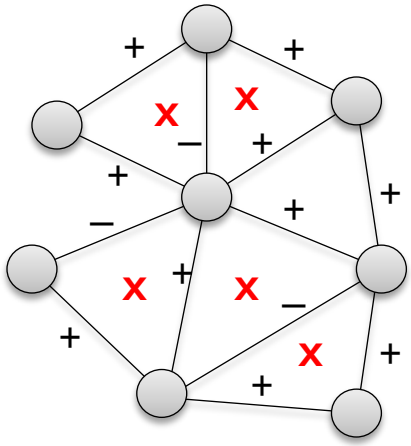

Shuffled data

[CHI '10]

Undirected Links: Balance

- Surprise values:

  i.e., z-score (deviation from random measured in the number of std. devs.)

	Triad	Epin	Wiki	Slashdot
Balanced	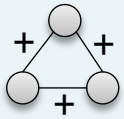	1,881	380	927
Balanced	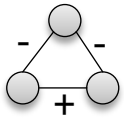	249	289	-175
Unbalanced	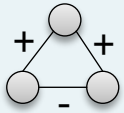	-2,105	-573	-824
Unbalanced	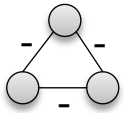	288	11	-9

- Observations:
  - Strong signal for balance
  - Epinions and Wikipedia agree on all types
  - Consistency with Davis's ['67] weak balance

[CHI '10]

Evolving directed networks

- Our networks are really directed
  - trust, opinion (, friendship)
- How many △ are now explained by balance?
  - Half (8 out of 16)

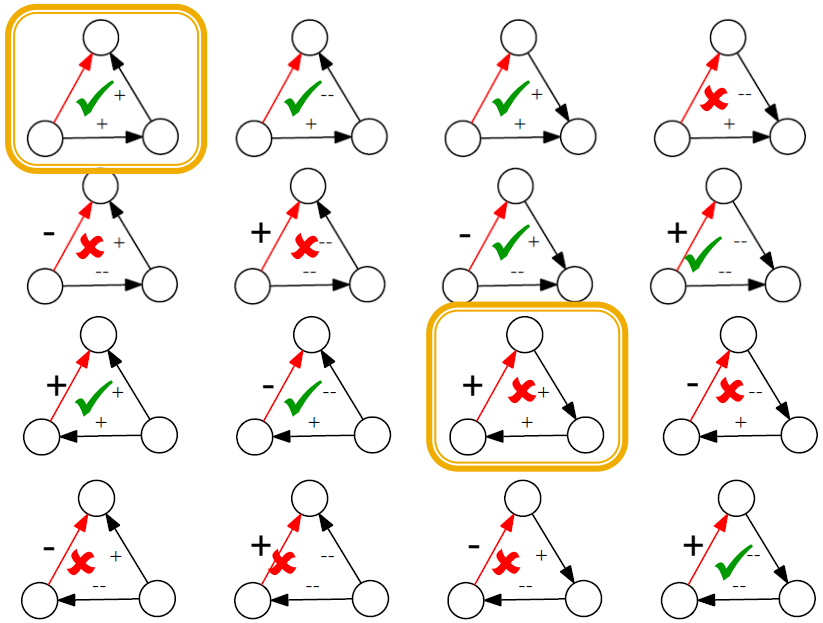

16 signed directed triads

- Is there a better explanantion?
  - Yes. **Theory of Status.**

[CHI '10]

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:

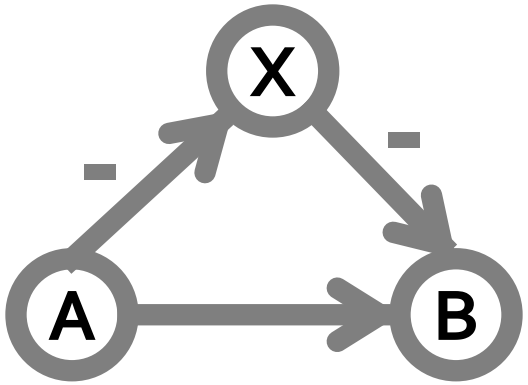

Balance: +
Status: –

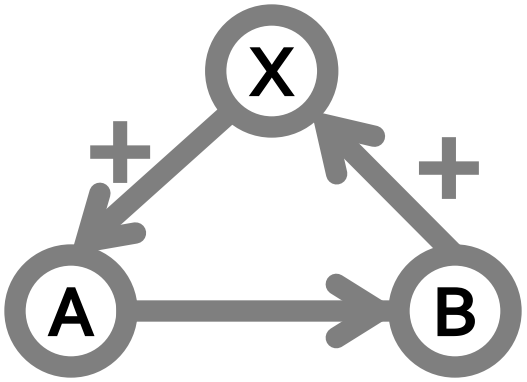

Balance: +
Status: –

[CHI '10]

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

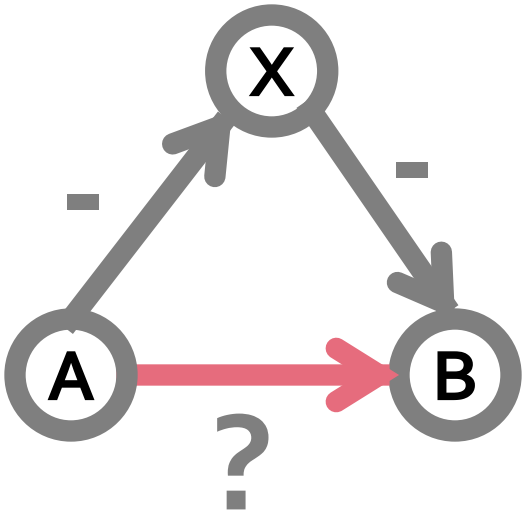

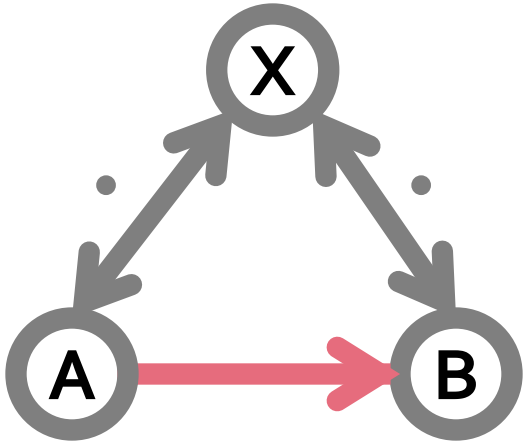

16 types of contextualized links

- **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

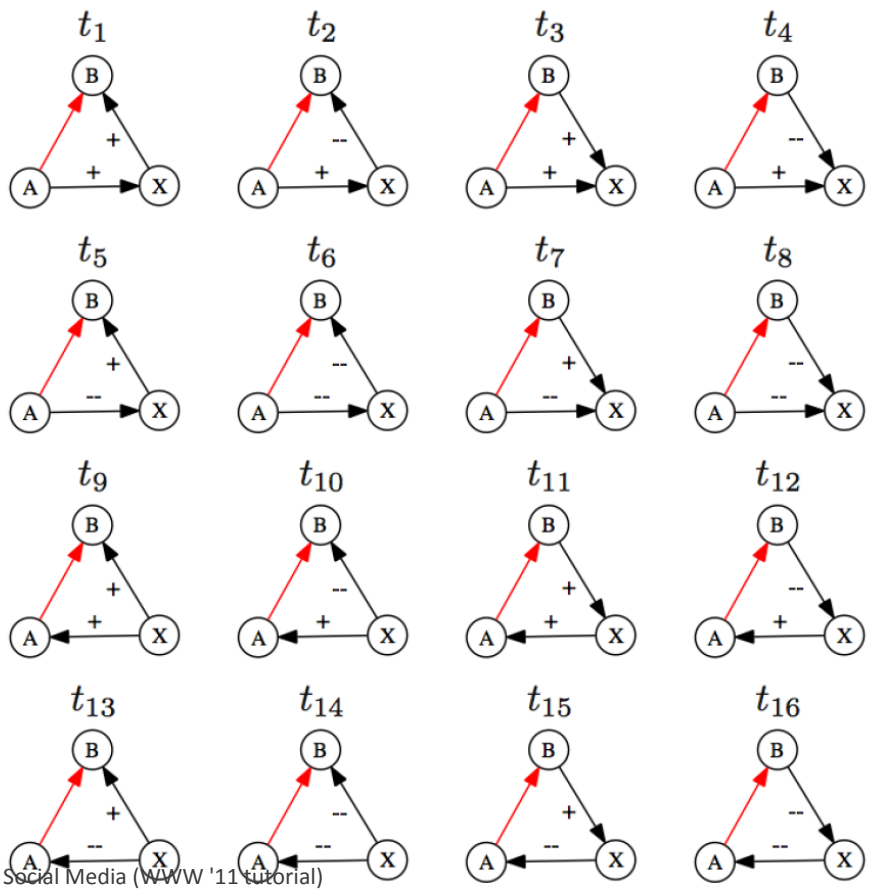

Heterogeneity in linking behavior

- 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**

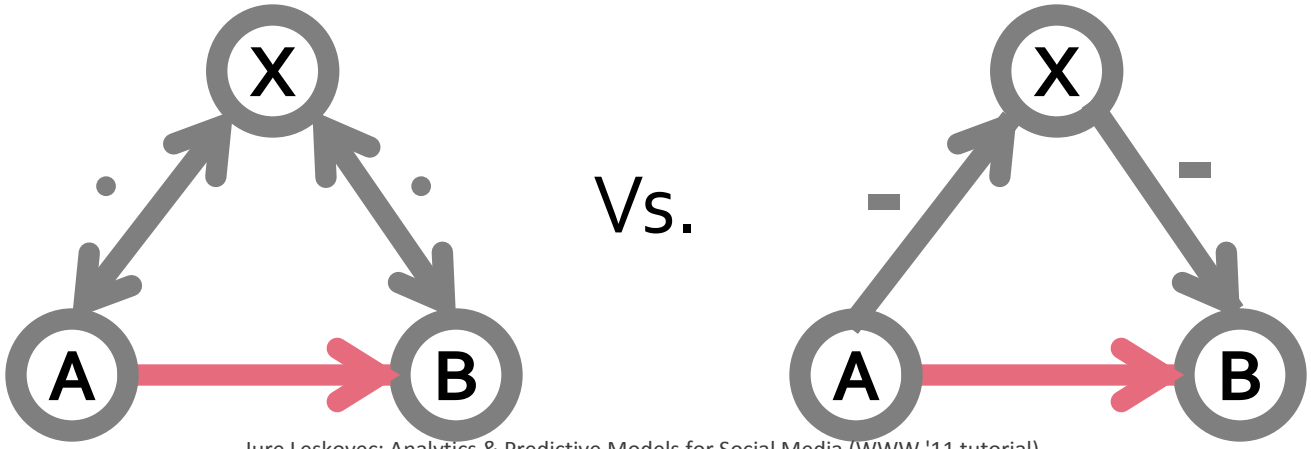

Status: Two Examples

- Two basic examples:

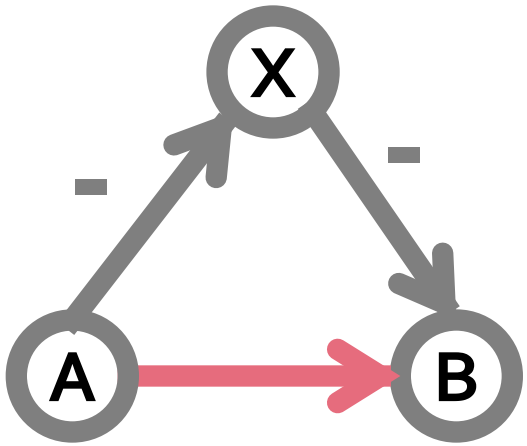

More **negative** than gen. baseline of A
More **negative** than rec. baseline of B

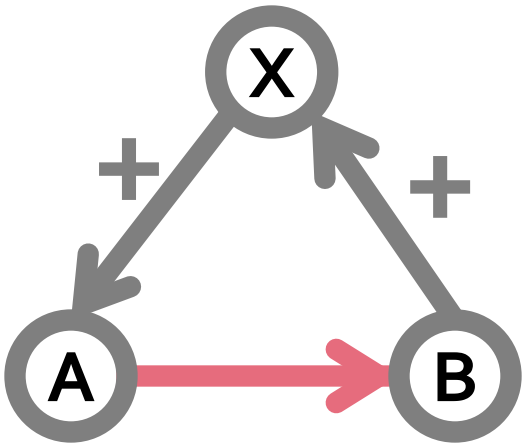

More **negative** than gen. baseline of A
More **negative** than rec. baseline of B

Consistency with Status

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

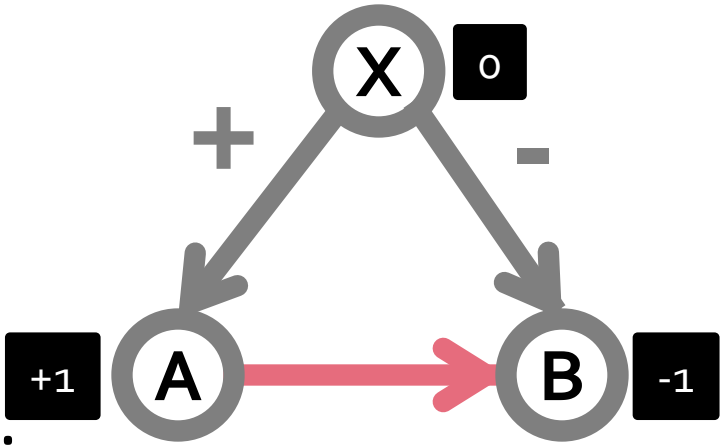

Status-consistent if:
Gen. surprise < 0
Rec. surprise < 0

- Surprise is **status**-consistent, if:
  - **G**en. surprise is status-consistent if it has **same** sign as status of B
  - **R**ec. 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

[CHI '10]

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

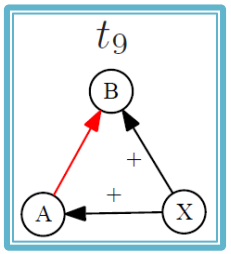

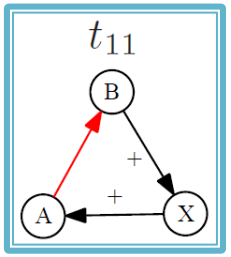

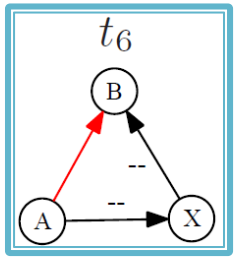

Cases where status fails

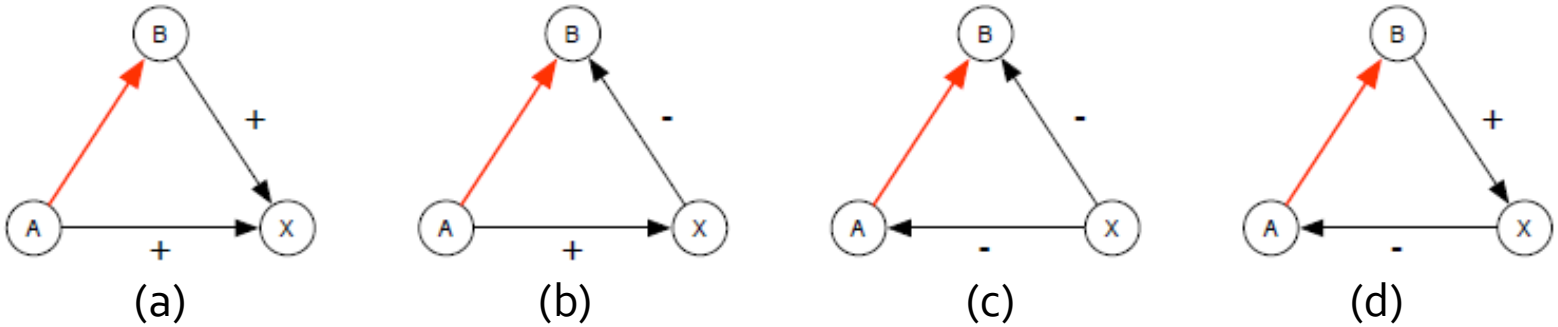

- Status seems to fail on triad types where A and B have both low status relative to X
  - Mistakes for gen. surprise (a) and (b) are “duals” of mistakes for receptive surprise (c) and (d)
  - Type (a) is one of most basic cases of balance theory

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

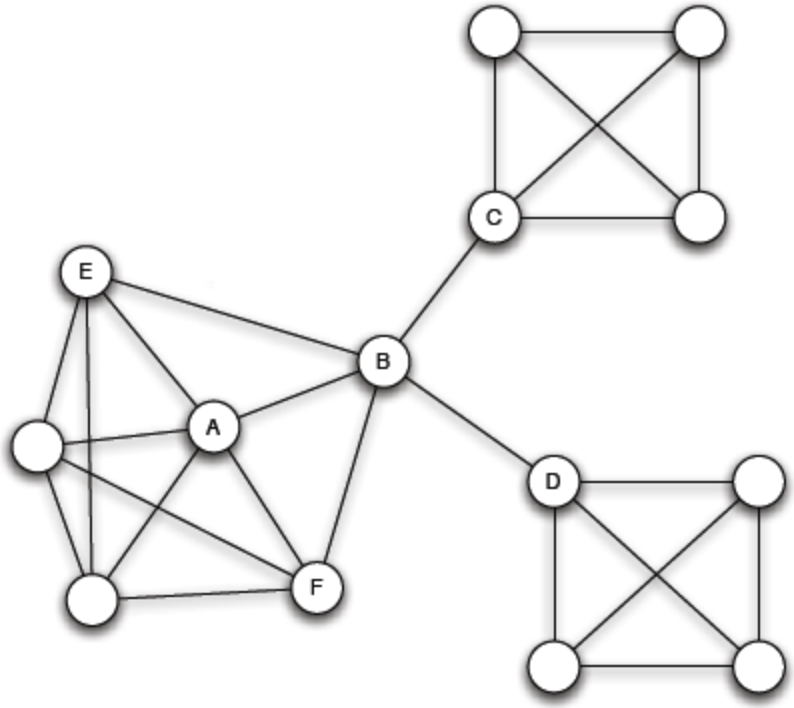

Global structure: Embeddedness

- 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

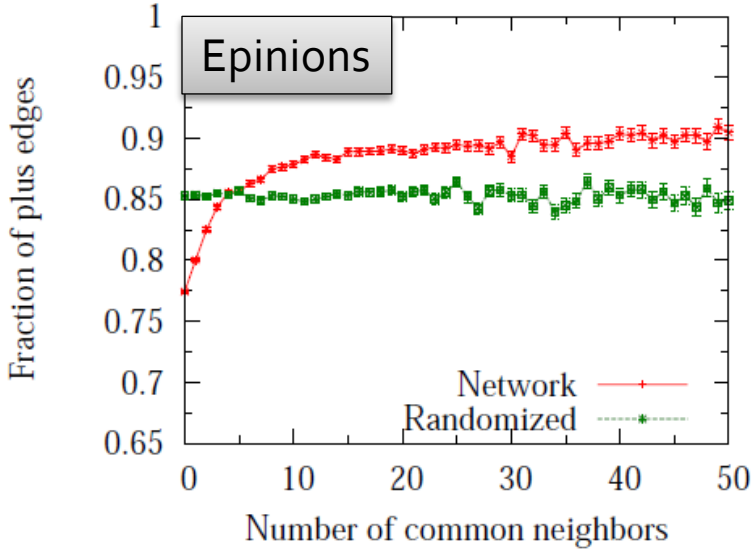

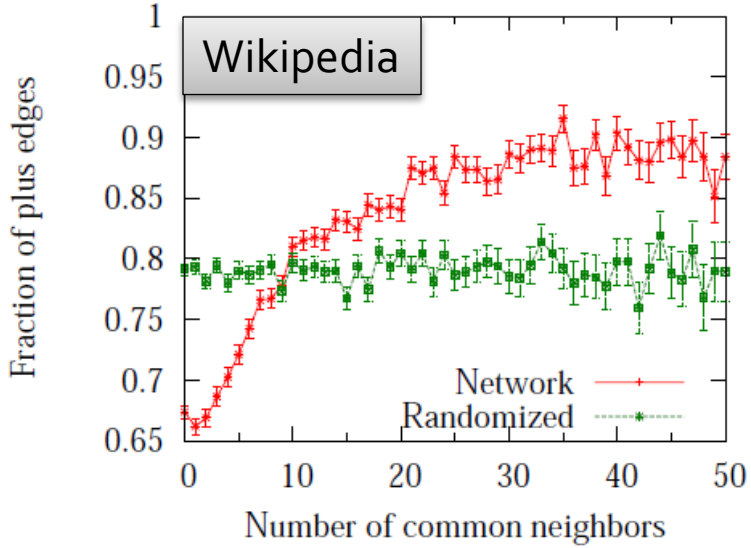

Predicting edge signs

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

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

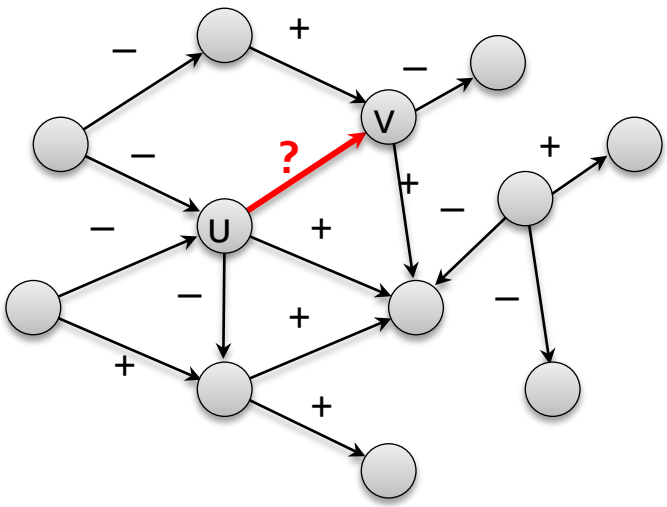

- Dataset:
  - Original: 80% +edges
  - Balanced: 50% +edges
- Evaluation:
  - Accuracy and ROC curves
- Features for learning:
  - Next slide

Features for learning

For each edge (u,v) create features:

- Triad counts (16):
  - Counts of signed triads edge u→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)

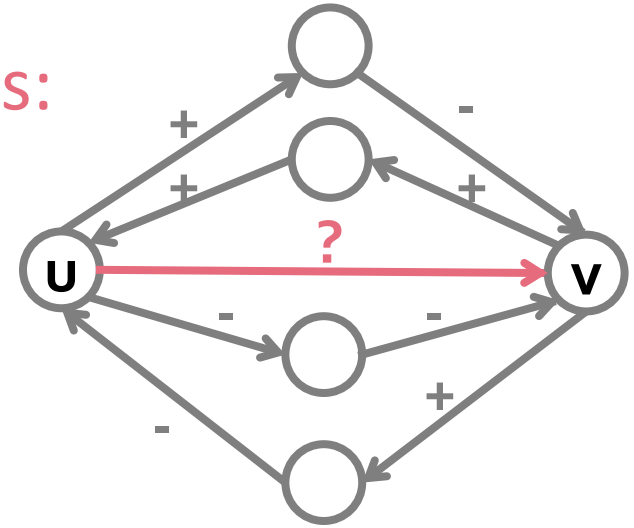

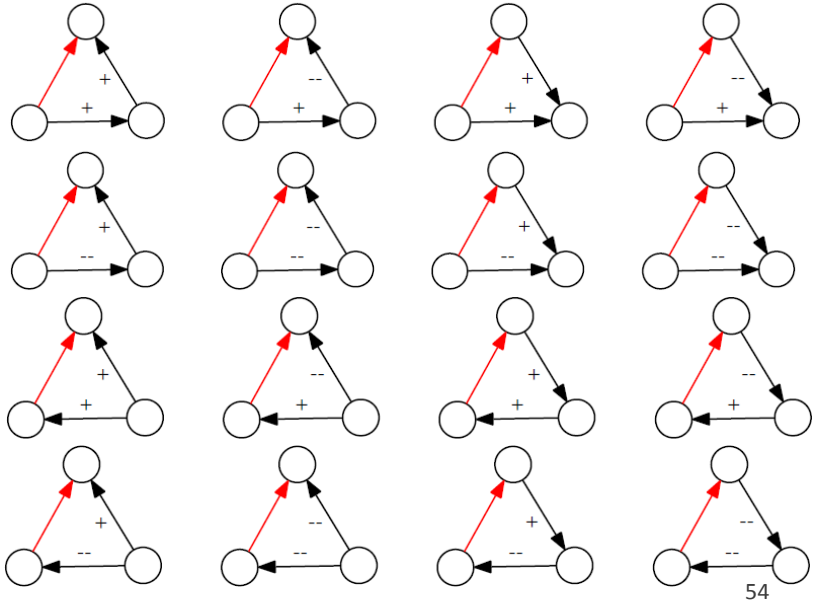

Edge sign prediction

- 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

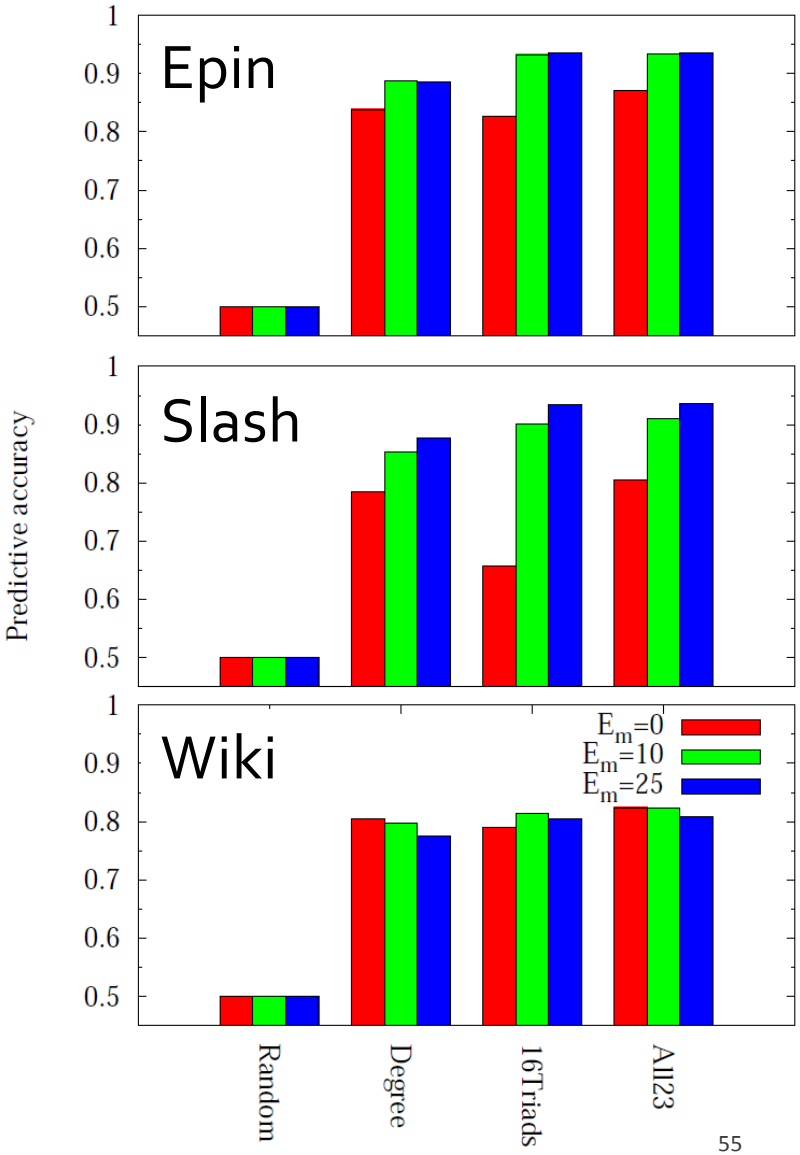

Generalization

- 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
$\bullet \xrightarrow{+} \bullet \xrightarrow{+} \bullet$	1	1	**0.5**	**0.9**	**0.3**
$\bullet \xrightarrow{+} \bullet \xrightarrow{-} \bullet$	-1	0	**-0.5**	**-0.9**	**-0.4**
$\bullet \xrightarrow{-} \bullet \xrightarrow{+} \bullet$	-1	0	**-0.4**	**-1.1**	**-0.3**
$\bullet \xrightarrow{-} \bullet \xrightarrow{-} \bullet$	1	-1	**-0.7**	**-0.6**	**-0.8**
$\circ \xrightarrow{+} \circ \xleftarrow{+} \circ$	1	0	0.3	0.4	0.05
$\circ \xrightarrow{+} \circ \xleftarrow{-} \circ$	-1	1	-0.01	-0.1	-0.01
$\bullet \xrightarrow{-} \bullet \xleftarrow{+} \bullet$	-1	-1	**-0.9**	**-1.2**	**-0.2**
$\circ \xrightarrow{-} \circ \xleftarrow{-} \circ$	1	0	0.04	-0.07	-0.03
$\circ \xleftarrow{+} \circ \xrightarrow{+} \circ$	1	0	0.08	0.4	0.1
$\bullet \xleftarrow{+} \bullet \xrightarrow{-} \bullet$	-1	-1	**-1.3**	**-1.1**	**-0.4**
$\circ \xleftarrow{-} \circ \xrightarrow{+} \circ$	-1	1	-0.1	-0.2	0.05
$\circ \xleftarrow{-} \circ \xrightarrow{-} \circ$	1	0	0.08	-0.02	-0.1
$\circ \xleftarrow{+} \circ \xleftarrow{+} \circ$	1	-1	-0.09	-0.09	-0.01
$\circ \xleftarrow{+} \circ \xleftarrow{-} \circ$	-1	0	-0.05	-0.3	-0.02
$\circ \xleftarrow{-} \circ \xleftarrow{+} \circ$	-1	0	-0.04	-0.3	0.05
$\circ \xleftarrow{-} \circ \xleftarrow{-} \circ$	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"

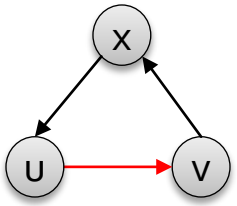

Tutorial Outline

- 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?**

[Danescu et al., 2009]

Recommendations & Opinions

- What do people think about our recommendations and opinions?

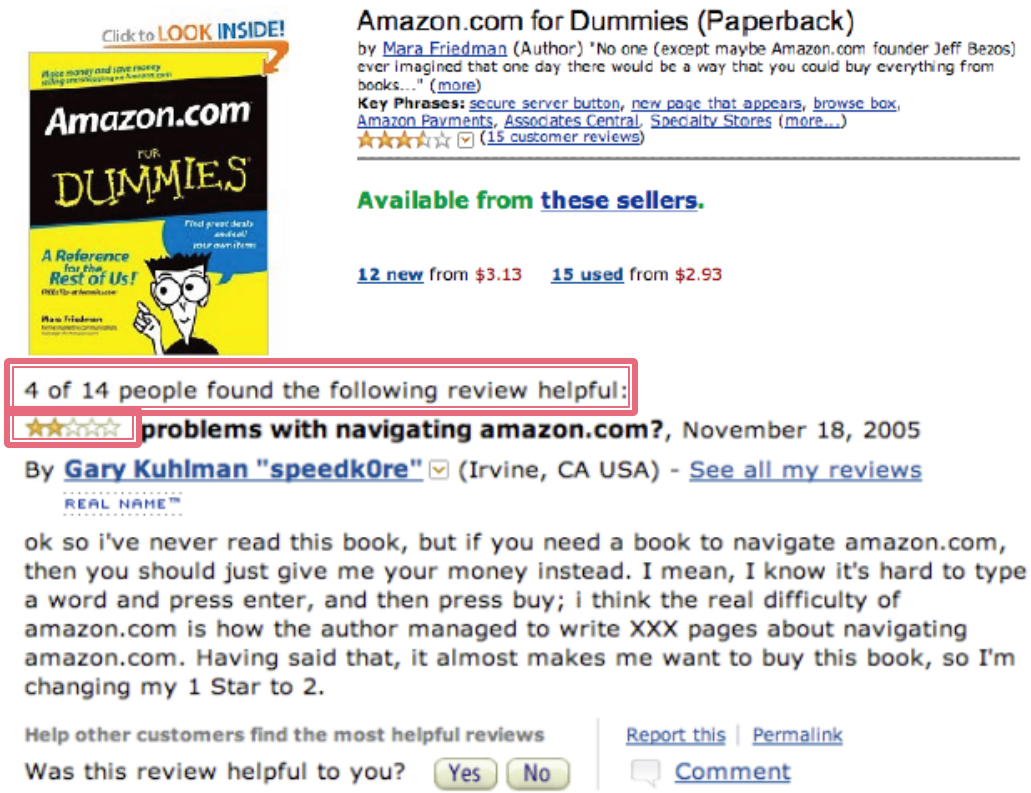

[Danescu et al., 2009]

Review helpfulness: Conformity

- People find conforming opinions more helpful

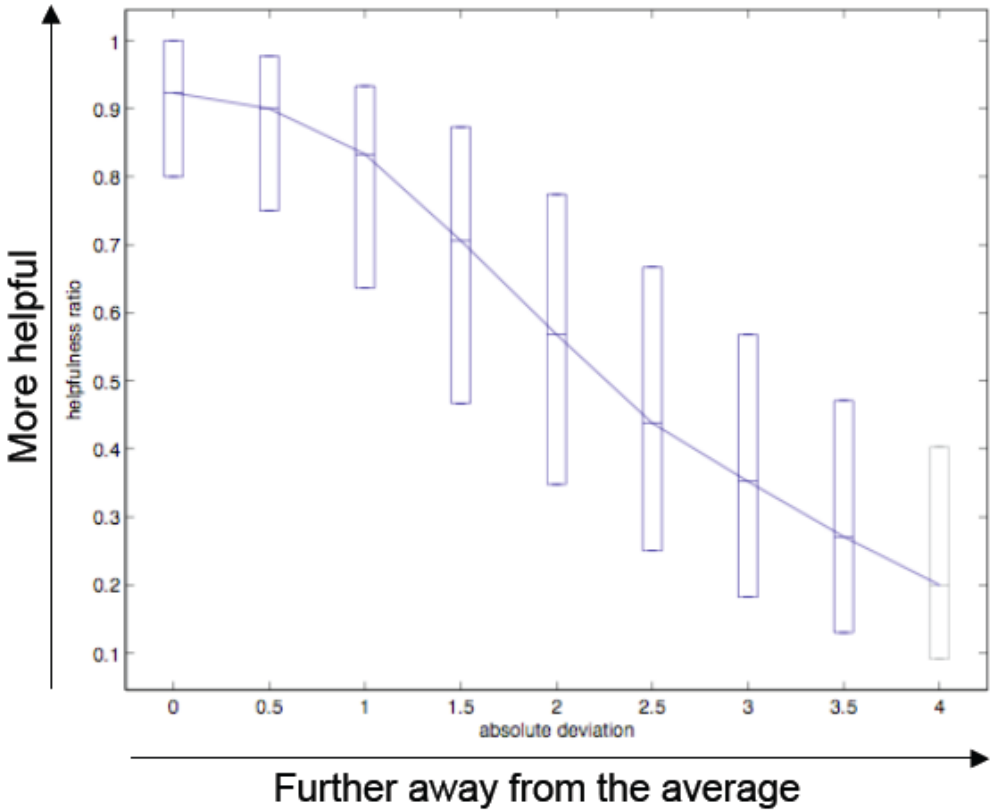

[Danescu et al., 2009]

Review helpfulness: Deviation

- Positive reviews are more helpful

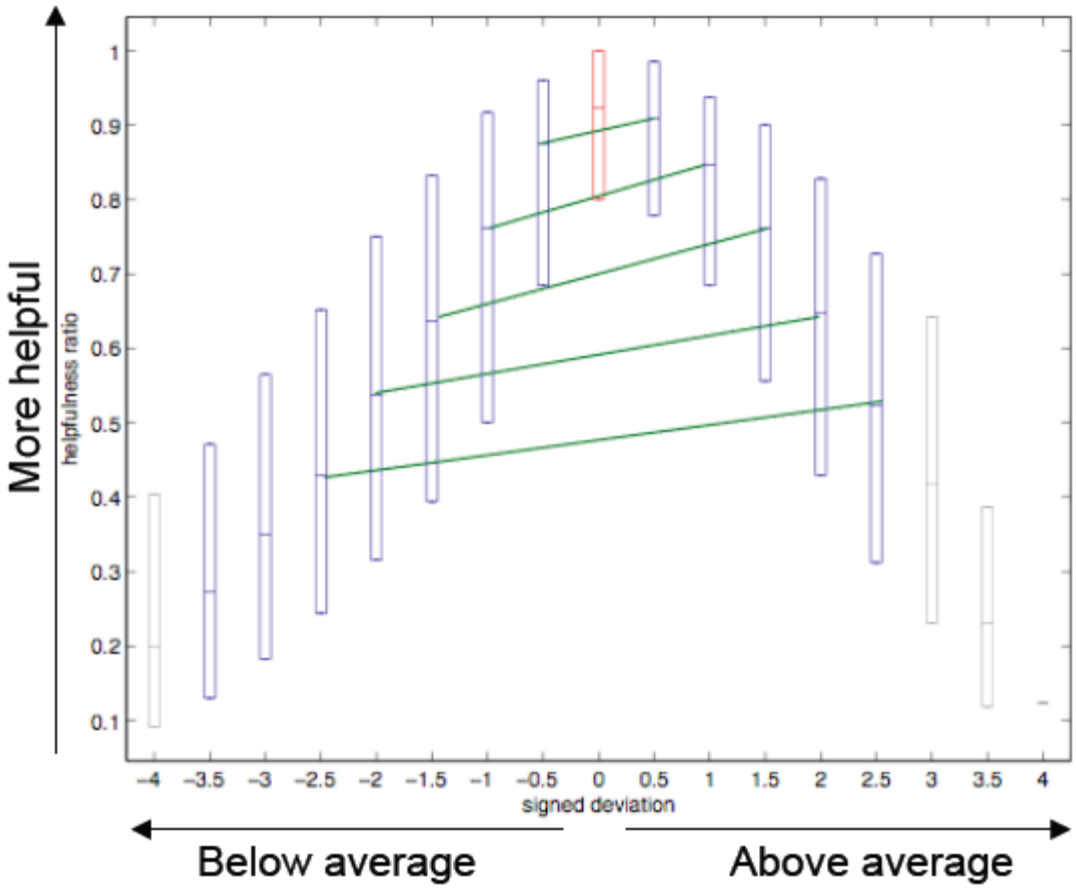

[Danescu et al., 2009]

Helpfulness: Varaince

- As variance of star ratings increases the two camps emerge

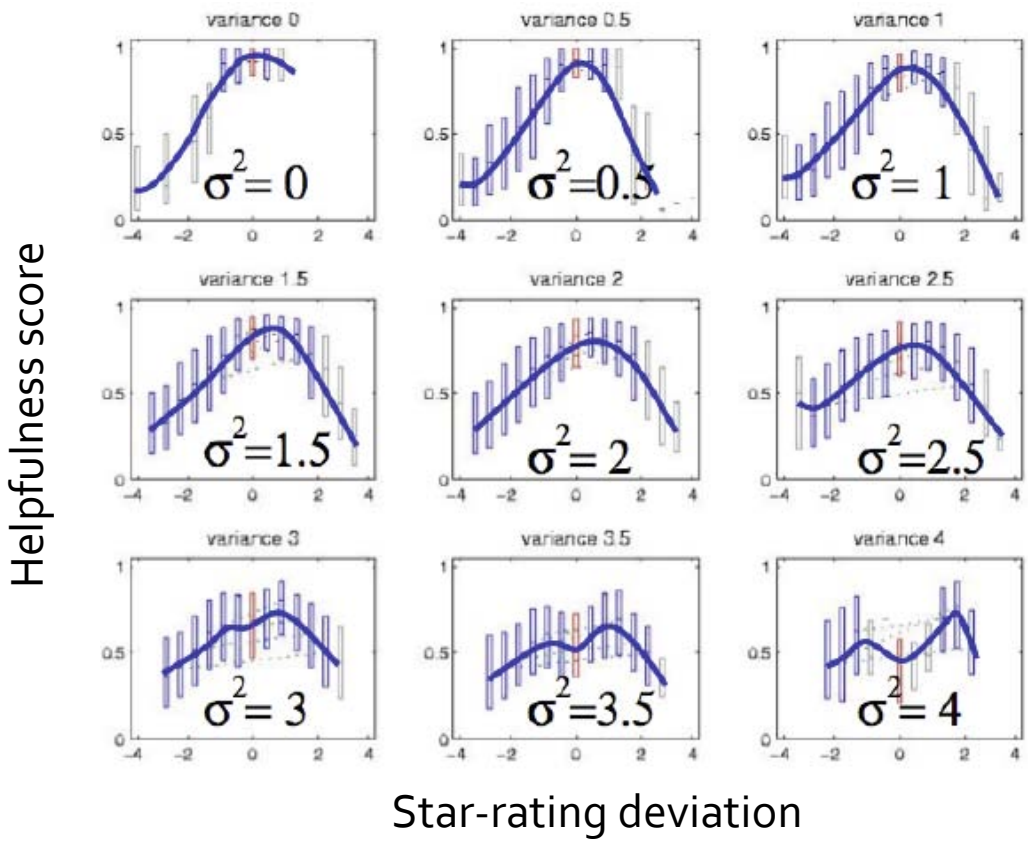

It is not about review quality but whether we agree with it or not

[ICWSM '10]

Aggregate vs. Individual

- So far we examined the aggregate behavior:
  - We observed what kinds of reviews the population as a whole find helpful
- What about behavior of individuals?
- Voting for admin promotion on Wikipedia
  - Vote is public and consequential
    - Do you support B to become Wikipedia administrator?
  - Data:
    - 2,800 elections between 2004-08
      - 44% resulted in promotion
    - 8,200 users
    - 114,000 votes (evaluations)
      - 80% of votes are positive

Effects of Status

- Voter (evaluator) V evaluates candidate C
- Prob. of positive evaluation of V as a function of status difference: $S_V - S_C$
  - Hypothesis: Monotonically decreases

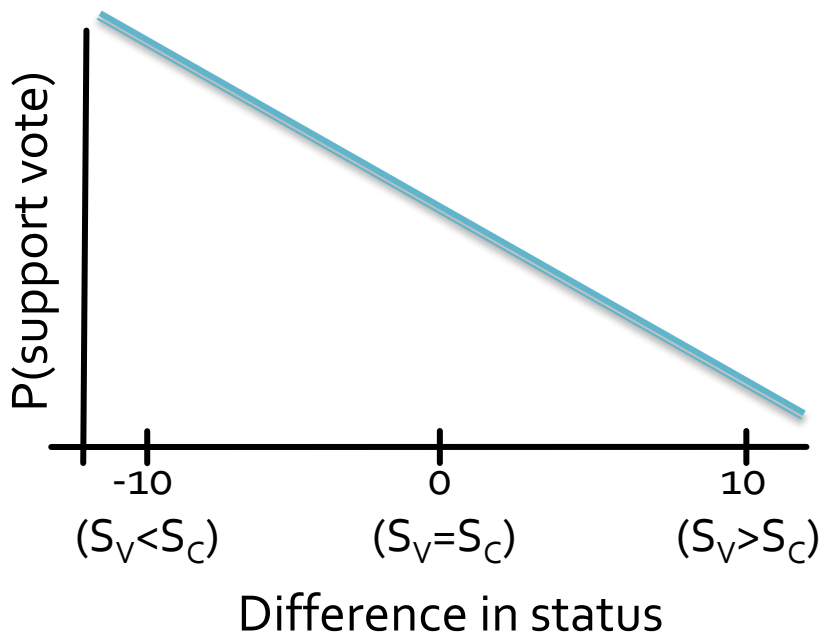

Effects of Status

- Prob. of positive evaluation of V as a function of status difference: $S_V - S_C$
  - Two notions of status:
    - Number of edits of a user
    - Number of barnstars
- Observations:
  - V is especially negative when status equals: $S_V = S_C$
  - "Mercy bounce" for $S_V > S_C$

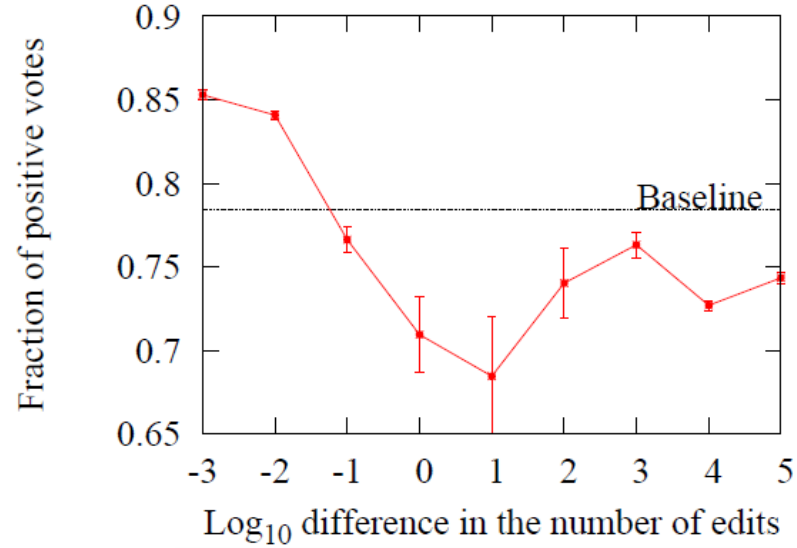

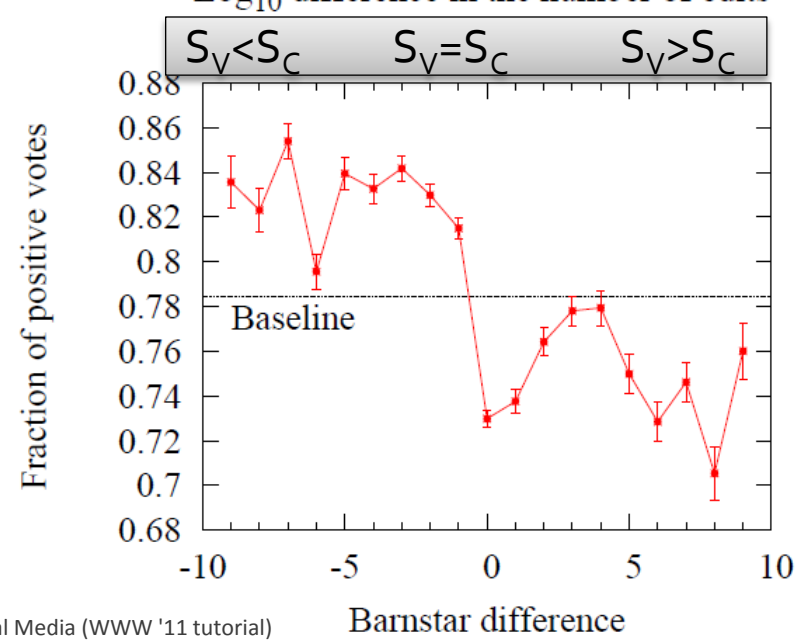

Direct V-C interaction

- Prob. of positive evaluation as a function of prior interaction of V and C

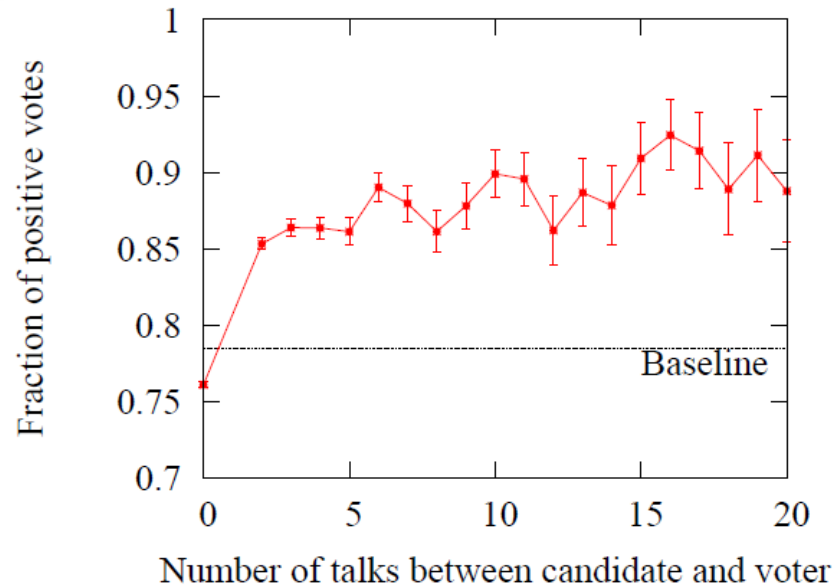

- Observation:
  Prior interaction increases prob. of a positive evaluation

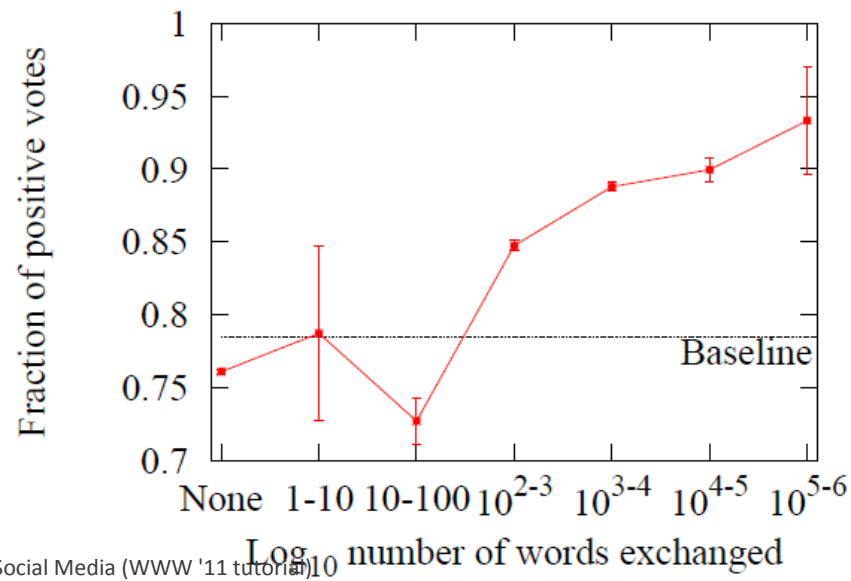

Thresholds and diversity of voters

- Aggregate response function:
  - How does prob. of positive evaluation depend on the frac. of positive evaluations so far?

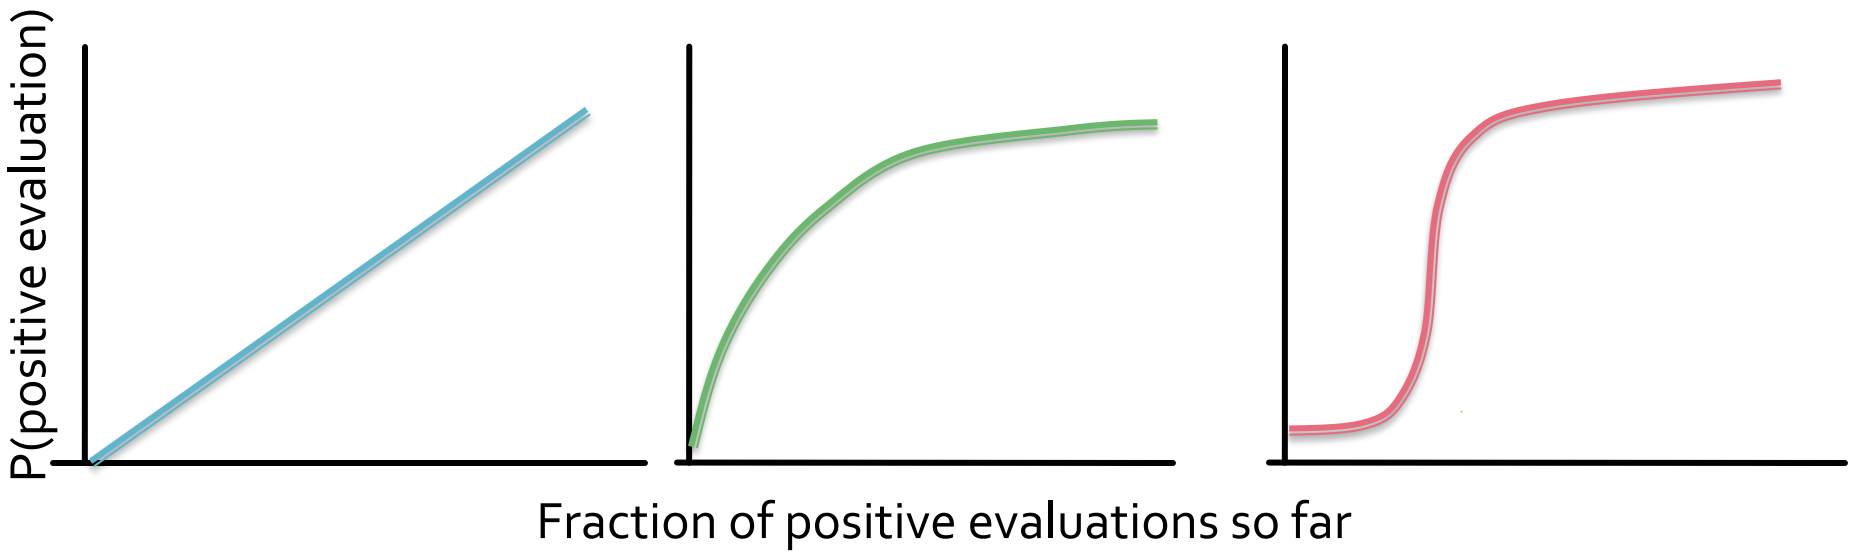

Linear? Diminishing returns? Threshold?

Thresholds and diversity of voters

- Aggregate response function:
  - Baseline: If voter were to flip a coin then $f(x)=x$
- Observation:
  - Voters more inclined to express opinion when it goes against the prevailing opinion
  - Consistent with [Wu-Huberman]

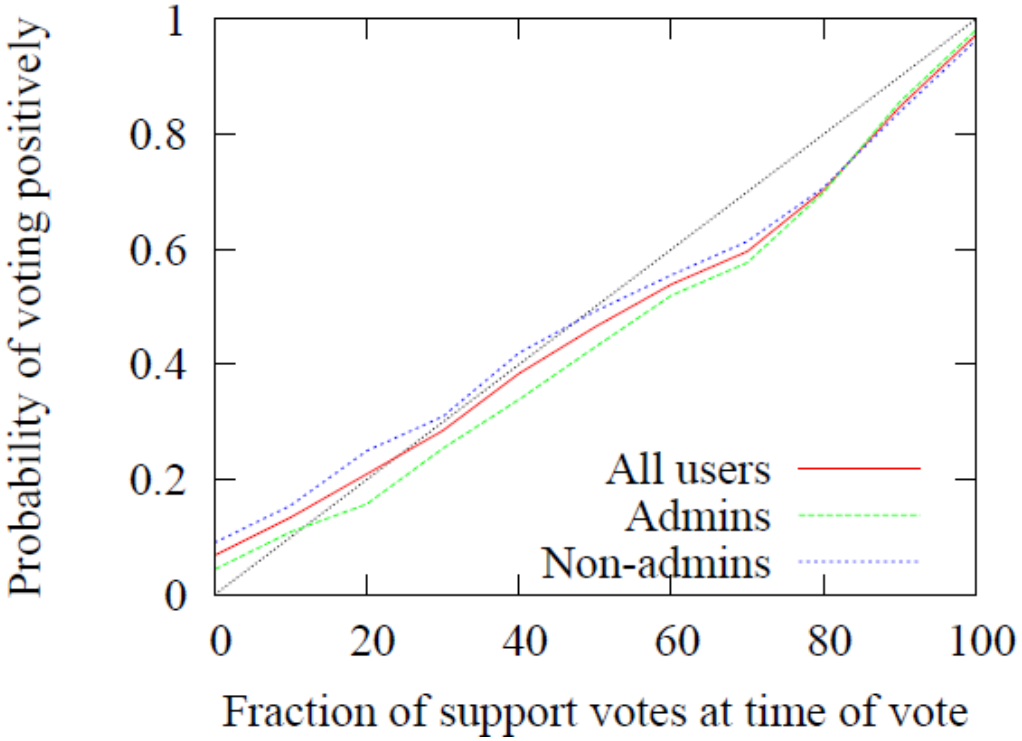

Thresholds and diversity of voters

- Personal response functions:
  - How does prob. of voter V positively evaluating C depend on frac. of positive evaluations so far?
  - Enough data that we can build models of individuals

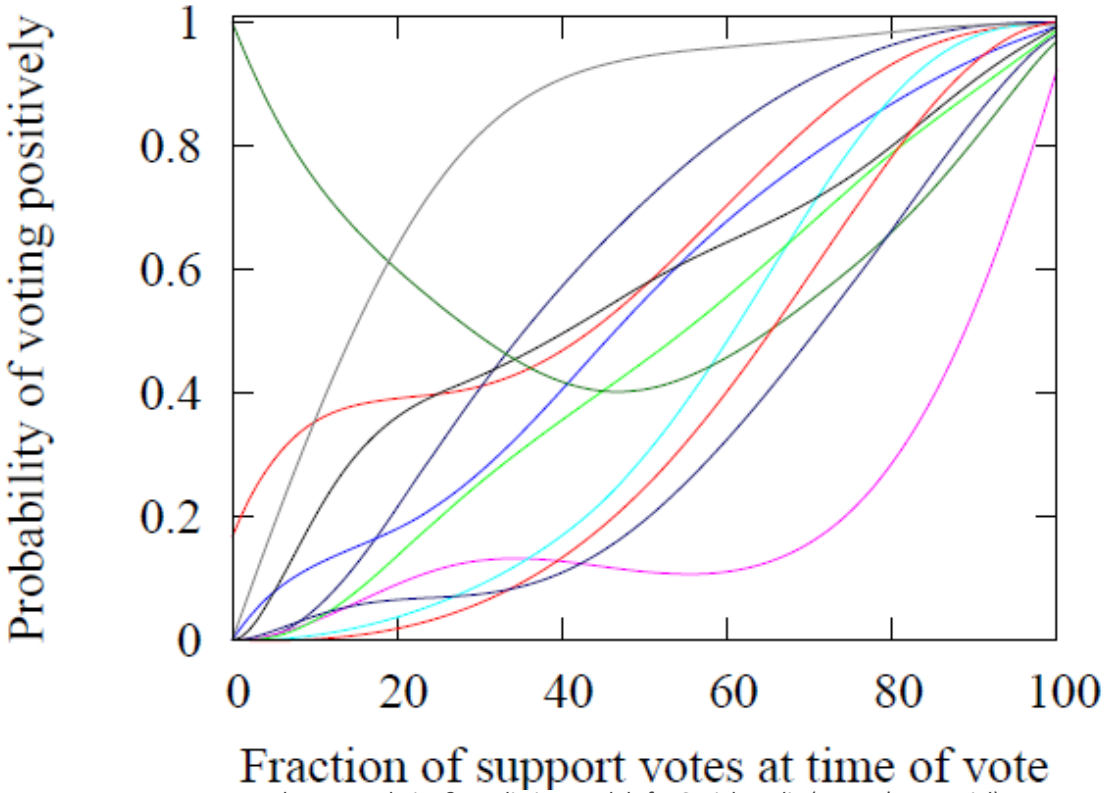

11 users that took part in >400 elections

Thresholds and diversity of voters

- Personal response functions:
  - Average is close to baseline but individual variation in shape of response function is large

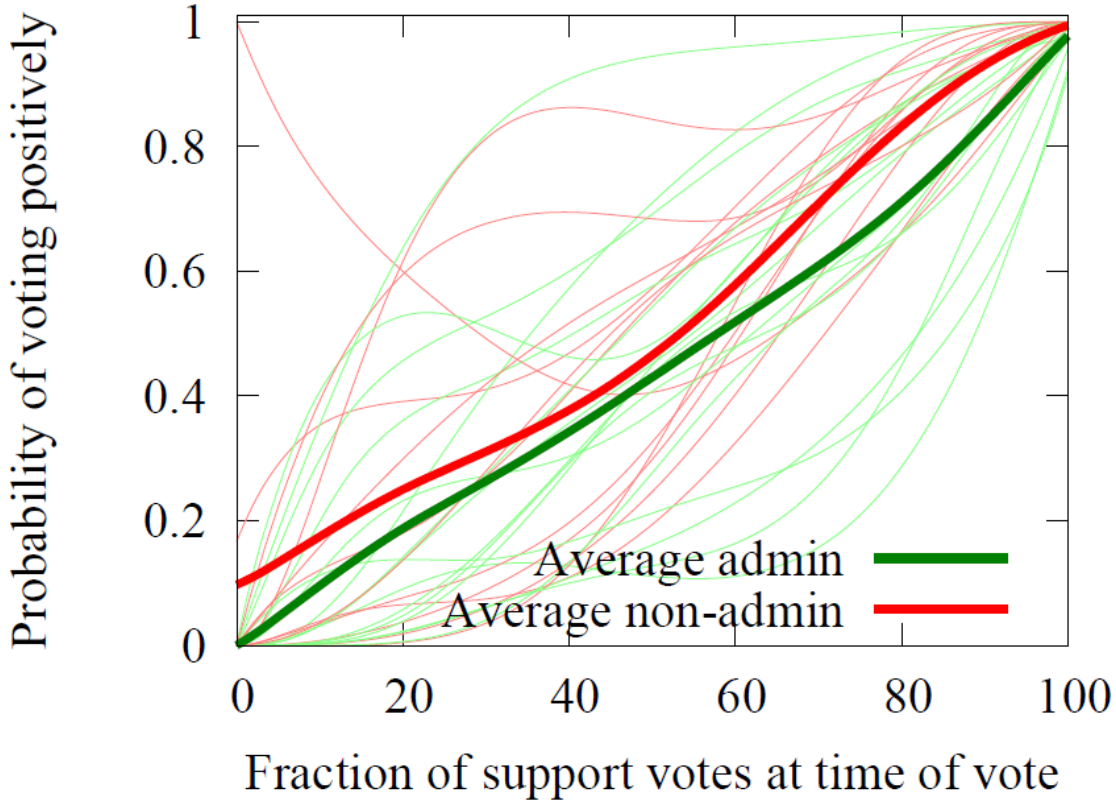

28 users that took part in >300 elections

Thresholds and diversity of voters

- Personal response functions:
  - Over time voters become more conservative
    - Response functions shift down and left

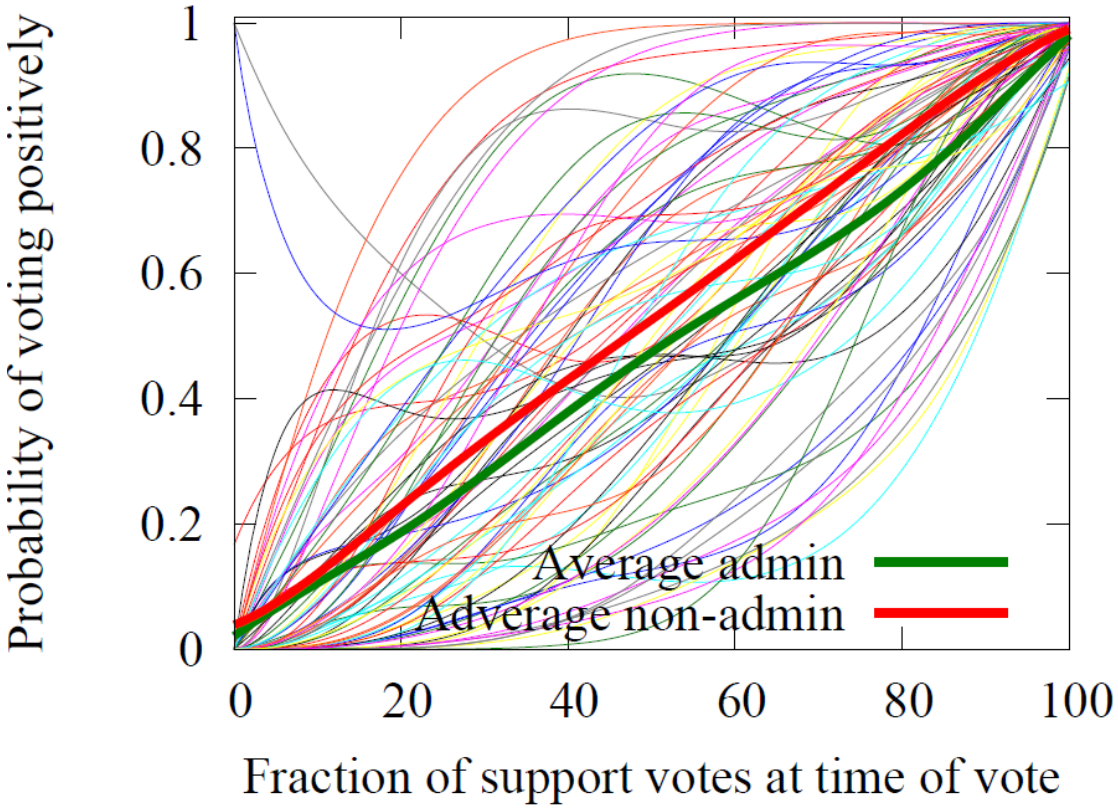

78 users that took part in >200 elections

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

New Setting: Stack Overflow

- Stack Overflow:
  - Q&A website for questions about programming
  - Users can up/down vote others questions/answers
  - Questions are annotated with tags describing relevant topics
  - Complete data:
    - 1.1M questions, 3.2M answers, and 7.5M votes (1.8M on questions, 5.7M on answers)
    - 93.4% of the votes are positive

Effects of Status & Similarity

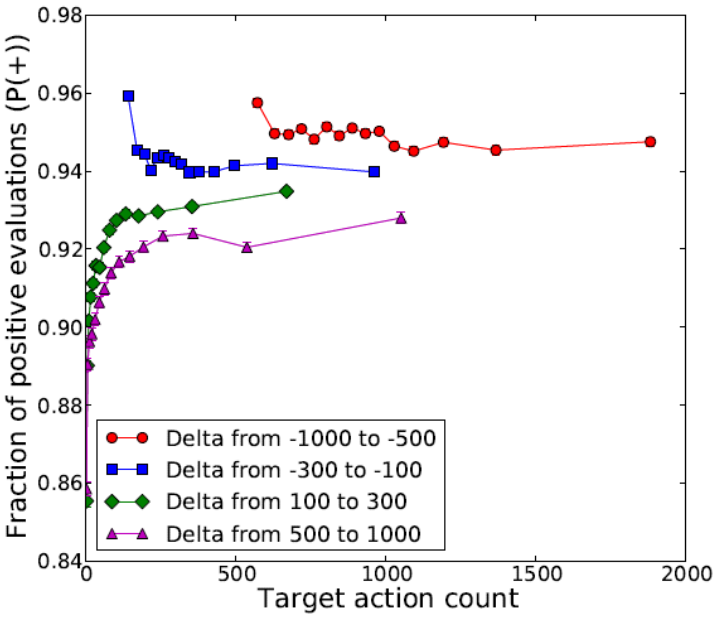

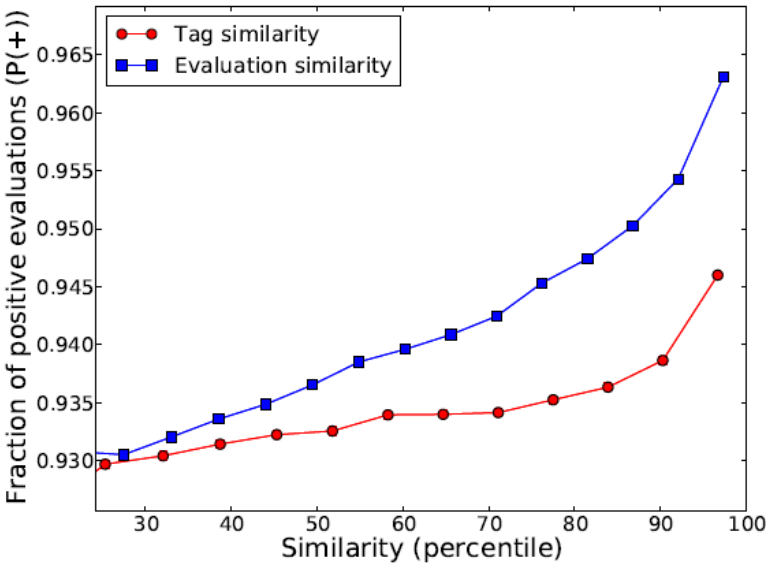

- For a fixed status difference curves are **flat**
  - Prob. of positive evaluation does not depend on target status
- Prob. of positive eval. increases in similarity both in topic (tag) and social circle (eval. similarity)

Status vs. Similarity

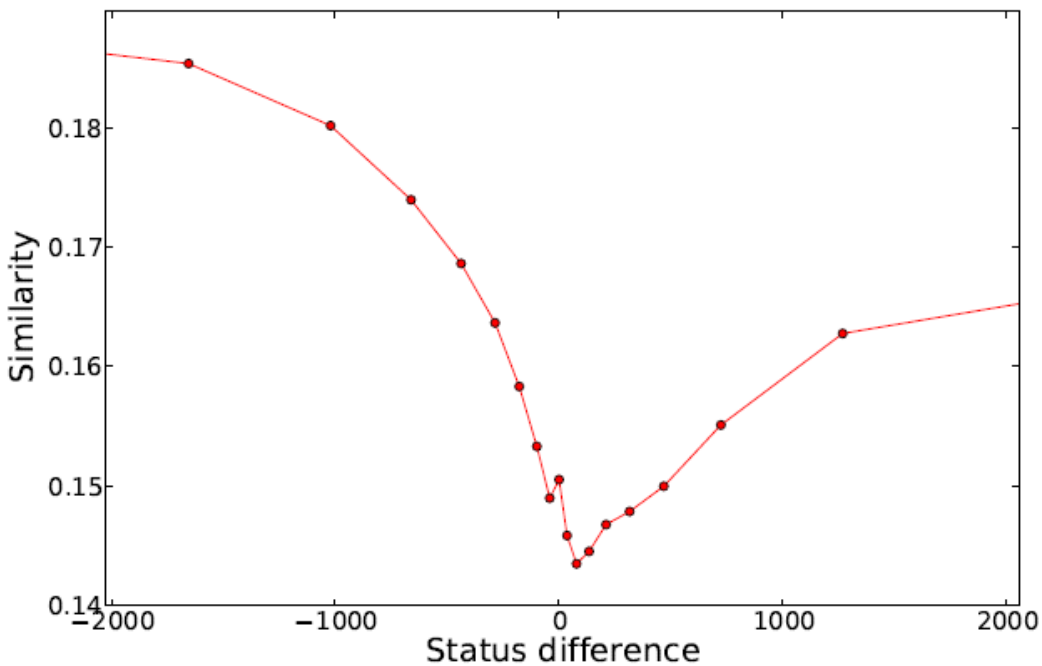

- Evaluators with lower status than the target tend to be more similar than when they have higher status.

Why most harsh at zero difference?

- How to explain low aggregate evaluations given by users to others of similar status?
  - This cannot be due to colleagues being overly tough on each other since the similarity tends to increase the positivity of evaluations, not decrease it.
- Two possible explanations:
  - (A) Low-status targets tend to garner evaluations that depend mainly on the target's absolute status
  - (B) Evaluations of higher-status targets are functions of the differential status between the pair

Explanation (A): Absolute Status

- Remember:
  - Evaluation behavior is different for low status targets
  - Most evaluators have low status
- Model (explains StackOverflow):
  - Evaluators do not discriminate among high-status targets
    - For high-status ($>T_{min}$) target vote + with const prob
  - But for low status targets they are increasingly negative as the target's status decrease
    - Vote + with prob. $\propto T/T_{min}$

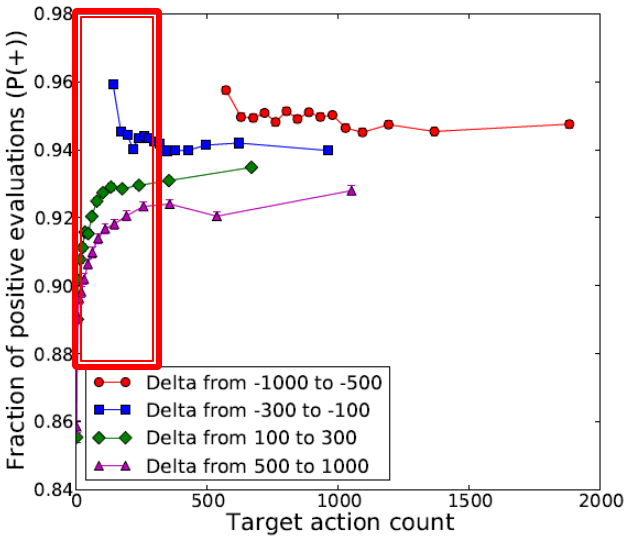

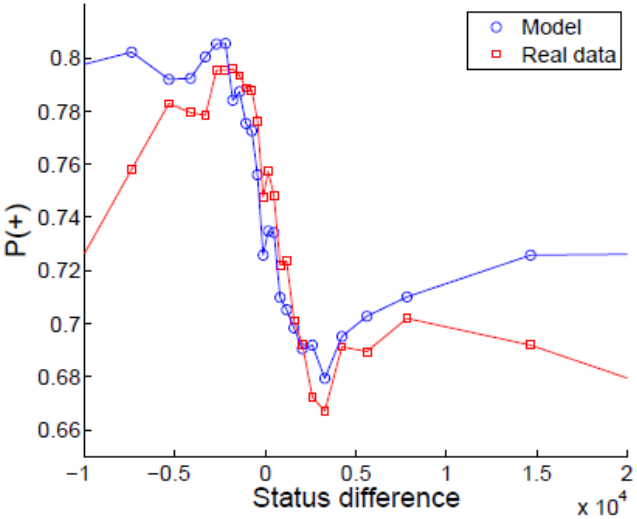

Explanation (B): Differential Status

- In Wikipedia the dip at zero persists even when removing low status targets
- Model ingredients:
  - High similarity between evaluator and target yields more positive votes
  - Selection bias effect: Higher-status evaluators are more similar to the target than lower-status evaluators are

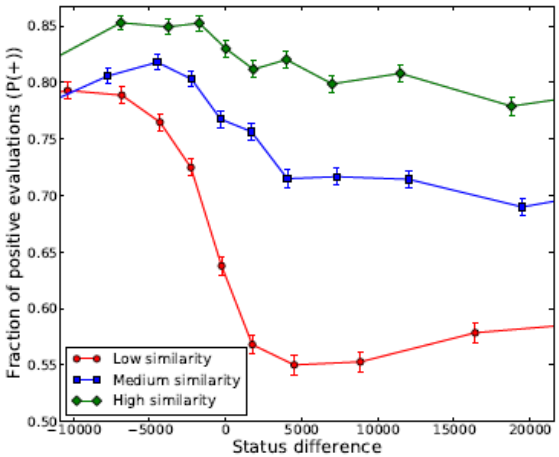

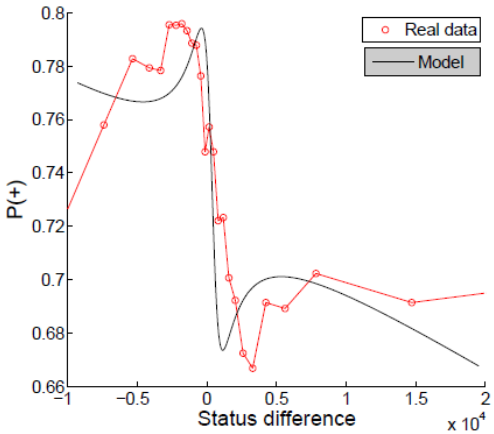

Application: Ballot-blind prediction

- Based on only who showed to up to evaluate predict the outcome of the Wiki election

Number of votes	E	Relative gain over LogReg
5	71.4%	12.8%
10	75.0%	23.8%
all	75.6%	25.7%

- Method:
  - Divide the Status-Similarity space, each cell prob. + vote
- Baseline:
  - Guessing gives 50% accuracy
  - Logistic Regression based on the target status (67% acc)

Conclusions: Status and Similarity

- Going beyond unlabeled edges
- Networks are globally organized based on status
- People use signed edges consistently regardless of a particular application
  - Near perfect generalization across datasets
- Assessing tie strength
- Effects of Status and Similarity
  - Implicit user evaluations

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