Recommendation in social networks

Shuang Hong Yang

CSE, Georgia Tech

Apr. 19th, 2011

from Social network to interest network

Levchin and Gurley Say That Next Big Company Will Capture The Interest Graph

Rip Empson

Like 398 Buzz 232 Tweet 737 4 Digg

As such, what everyone in Silicon Valley and "Venture Land" conceive of as the real game-changing model involves capturing and capitalizing on the "interest graph," he says. The company that succeeds in doing so would be "close to the Google search paradigm because it would be right in line with demand generation and with discovery that relates to product purposes." Thus, it is the interest graph that defines the middle ground between Google and Facebook — between search, advertising, and the social graph.

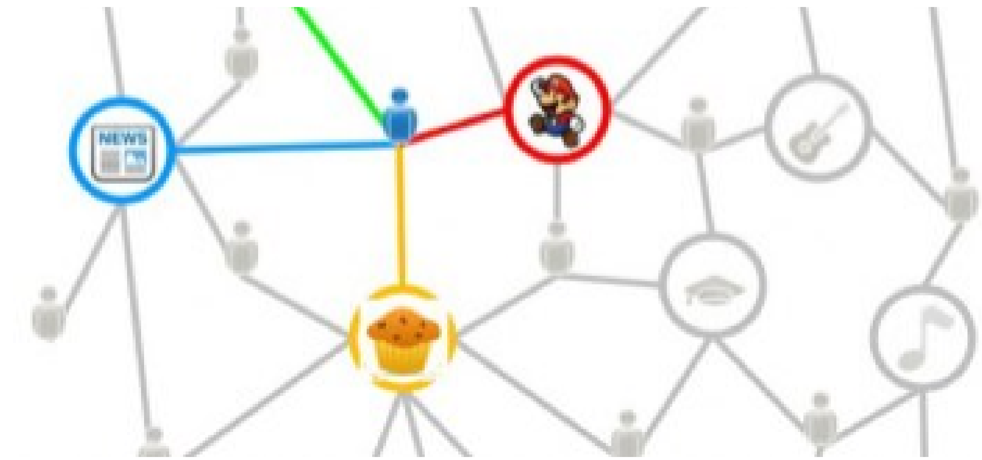

changing technology" and the future of the Web. Emblematic of today's mindset, they attacked this rather large topic by comparing the strengths and objectives of Google and Facebook, using the latter's jaw-dropping stats (500+ million users, 1 in every 13 people *on Earth* logs into Facebook each day) and its promotion of the social graph as a measure of what's to come.

from Social network to interest network

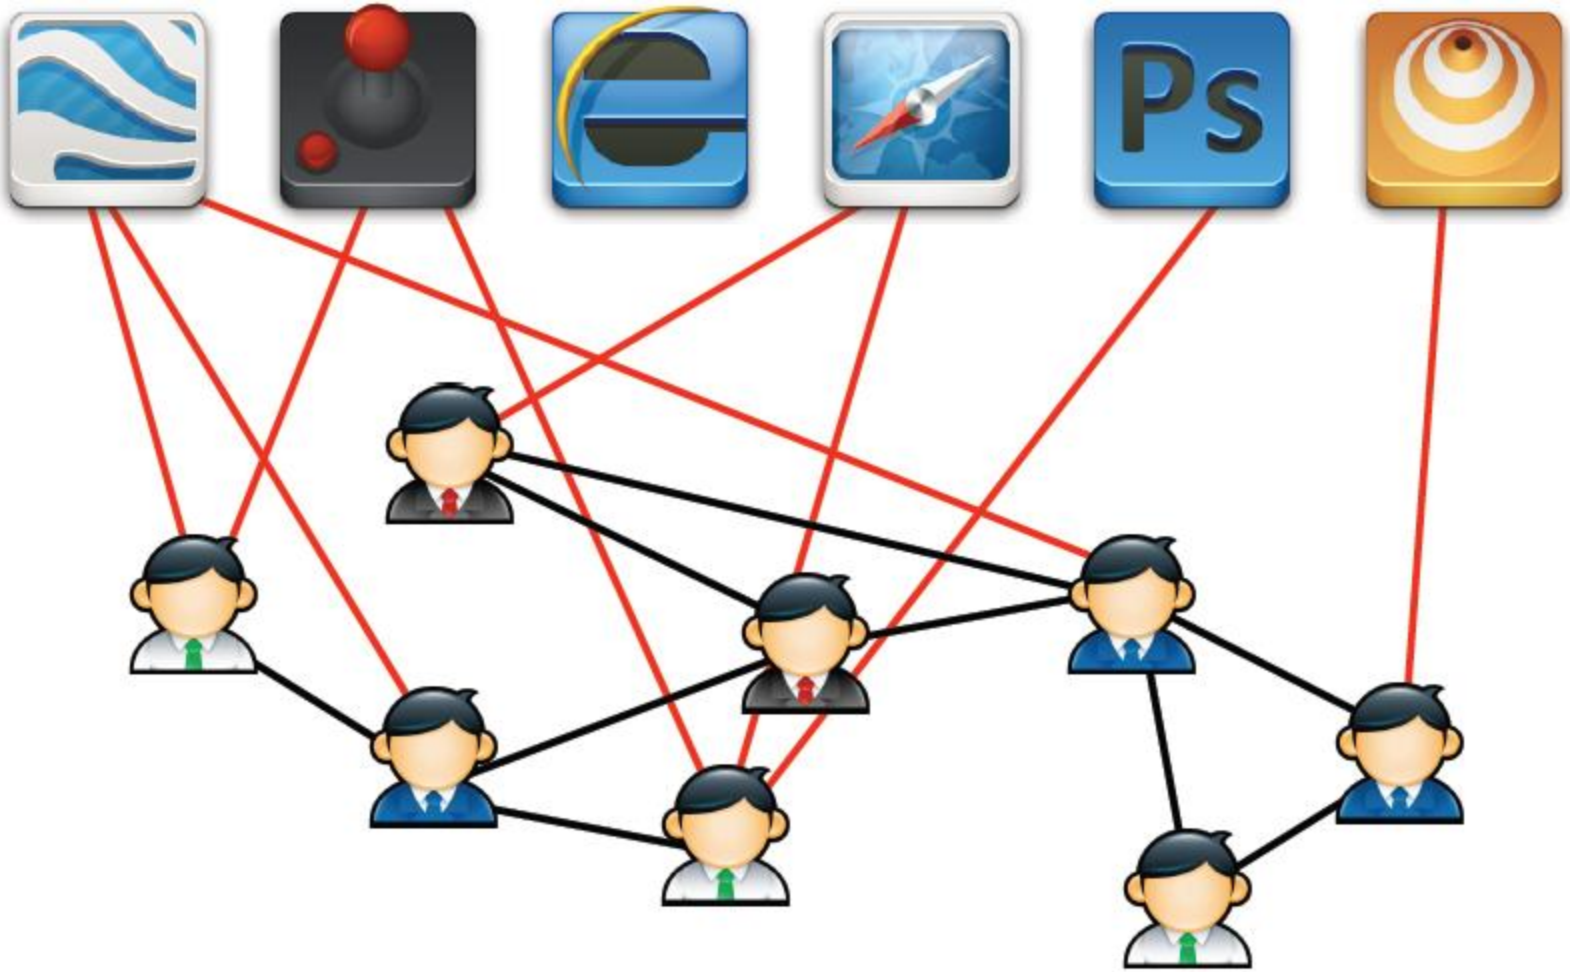

users connect to their friends

users interact with service items (applications, ads, games, movies,...)

from Social network to interest network

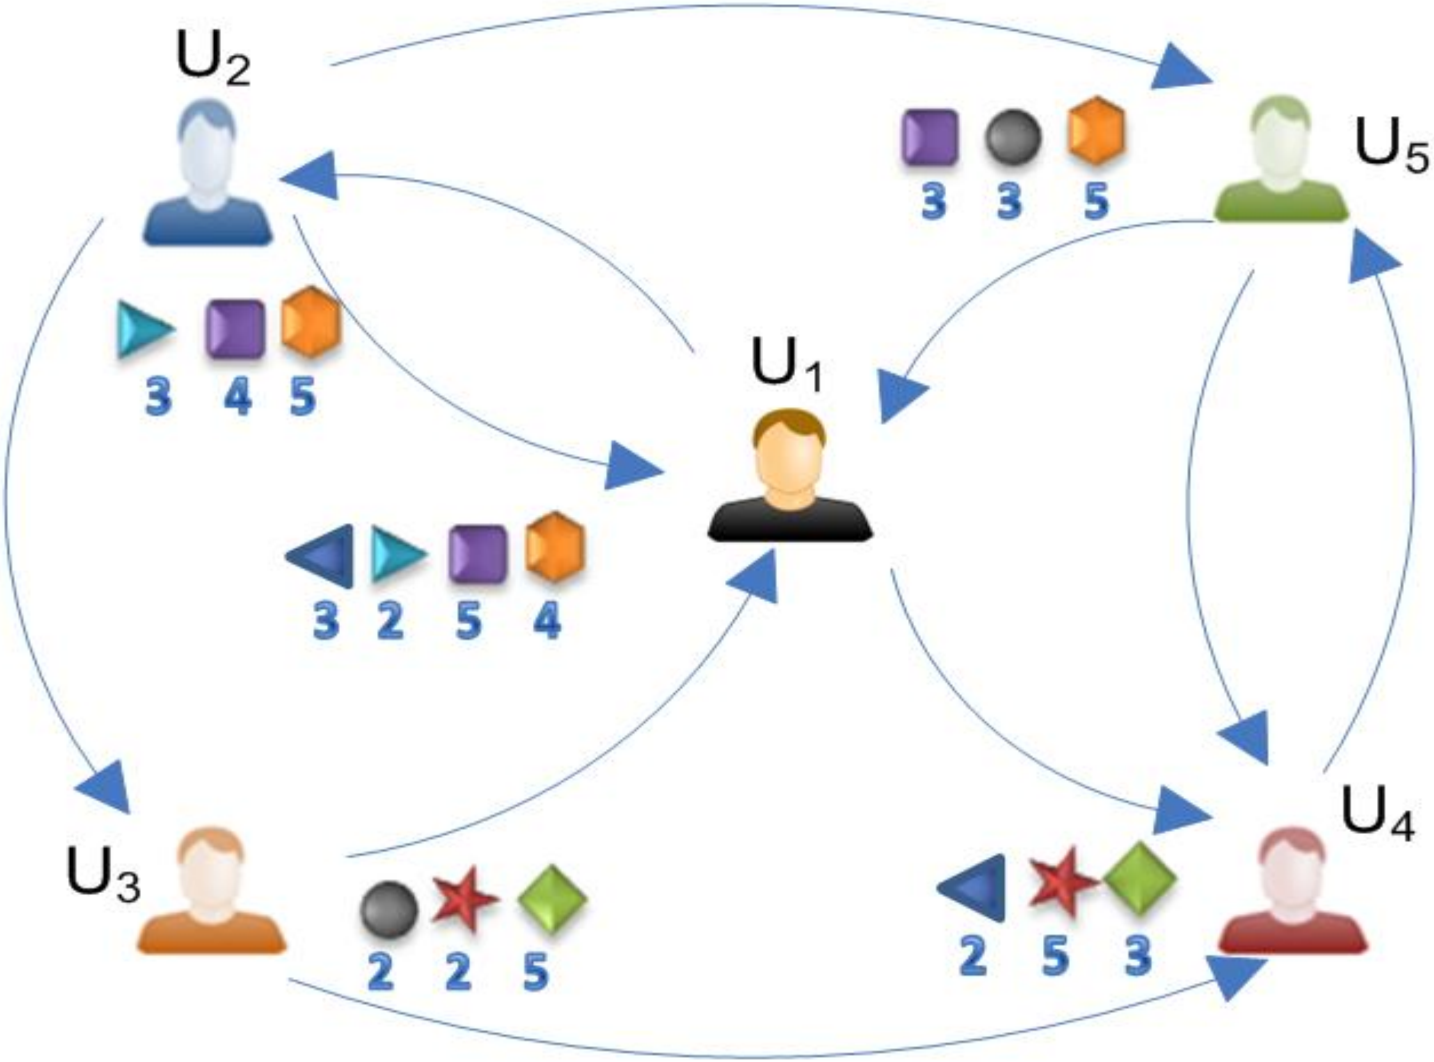

social rating network

What can we gain from social graph

- Homophily
  - People connected to each other tend to have similar interest;
- Influence
  - trust, agreement, approval
  - distrust, disagreement, opposement

Representative work

- Homophily
  - [Ma et al, SIGIR' 2009]

    Learning to Recommend with Social Trust Ensemble
  - [Jamali & Ester et al, RecSys' 2010]

    A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks
  - [Yang et al, WWW' 2011]

    Like like alike -- Joint friendship and interest propagation in social networks
- Influence
  - [Ma et al, RecSys' 2009]

    Learning to Recommend with Trust and Distrust Relationships

Representative work

- Homophily
  - [Ma et al, SIGIR' 2009]

    Learning to Recommend with Social Trust Ensemble
  - [Jamali & Ester et al, RecSys' 2010]

    A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks
  - [Yang et al, WWW' 2011]

    Like like alike -- Joint friendship and interest propagation in social networks
- Influence
  - [Ma et al, RecSys' 2009]

    Learning to Recommend with Trust and Distrust Relationships

Social trust ensemble … [Ma et al *SIGIR' 2009*]

- Trust graph + rating matrix

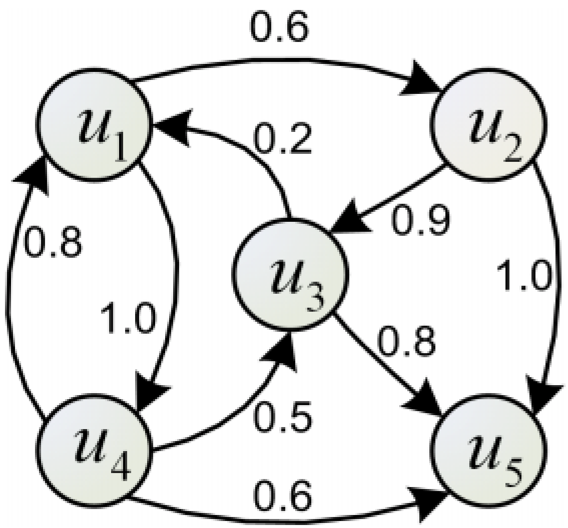

(a) Social Trust Graph

	v_1	v_2	v_3	v_4	v_5	v_6
u_1		5	2		3	
u_2	4			3		4
u_3			2			2
u_4	5			3		
u_5		5	5			3

(b) User-Item Rating Matrix

Figure 1: Example for Trust based Recommendation

Social trust ensemble ... [Ma et al *SIGIR' 2009*]

- Users makes decisions by:
  - *either* following her own taste

$$p(R|U,V,\sigma_R^2) = \prod_{i=1}^{m}\prod_{j=1}^{n}\left[\mathcal{N}\left(R_{ij}|g(U_i^T V_j),\sigma_R^2\right)\right]^{I_{ij}^R}$$

$$p(U|\sigma_U^2)=\prod_{i=1}^{m}\mathcal{N}(U_i|0,\sigma_U^2\mathbf{I}),\ \ p(V|\sigma_V^2)=\prod_{j=1}^{n}\mathcal{N}(V_j|0,\sigma_V^2\mathbf{I})$$

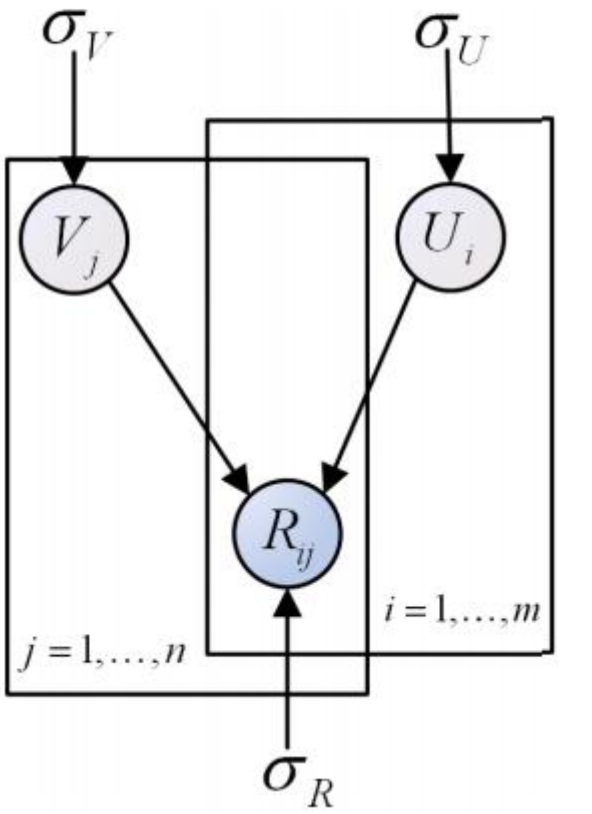

Social trust ensemble ... [Ma et al *SIGIR' 2009*]

- Users makes decisions by:
  - *either* following her own taste
  - *or* learning from people she trusts

$$\widehat{R}_{ik} = \sum_{j \in \mathcal{T}(i)} R_{jk} S_{ij} = \sum_{k \in \mathcal{T}(i)} S_{ik} U_k^T V_j$$

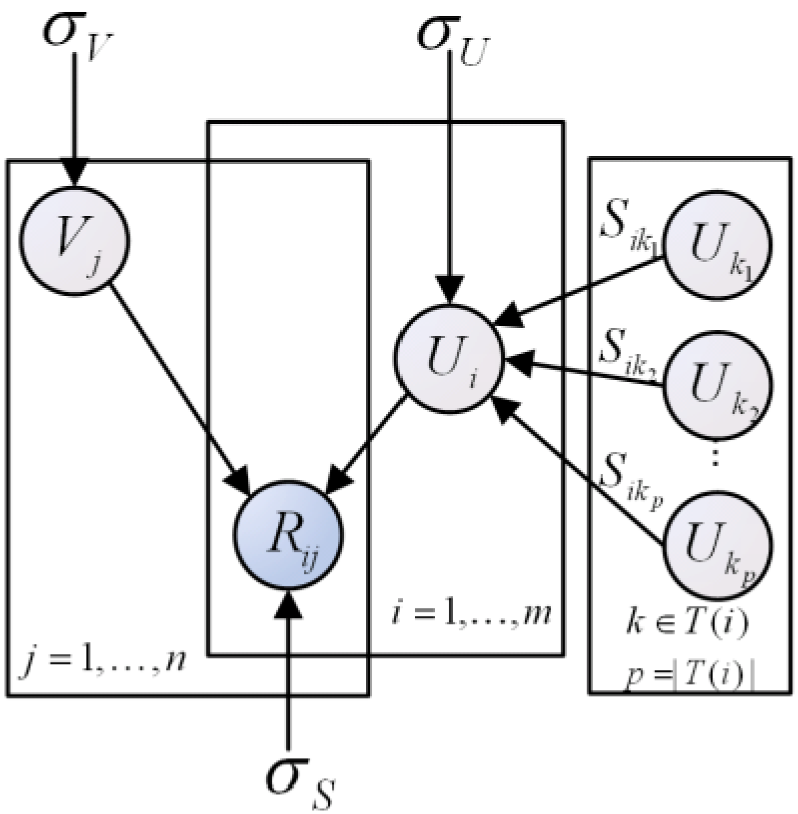

Social trust ensemble ... [Ma et al *SIGIR' 2009*]

- Users makes decisions by:
  - *either* following her own taste
  - *or* learning from people she trusts

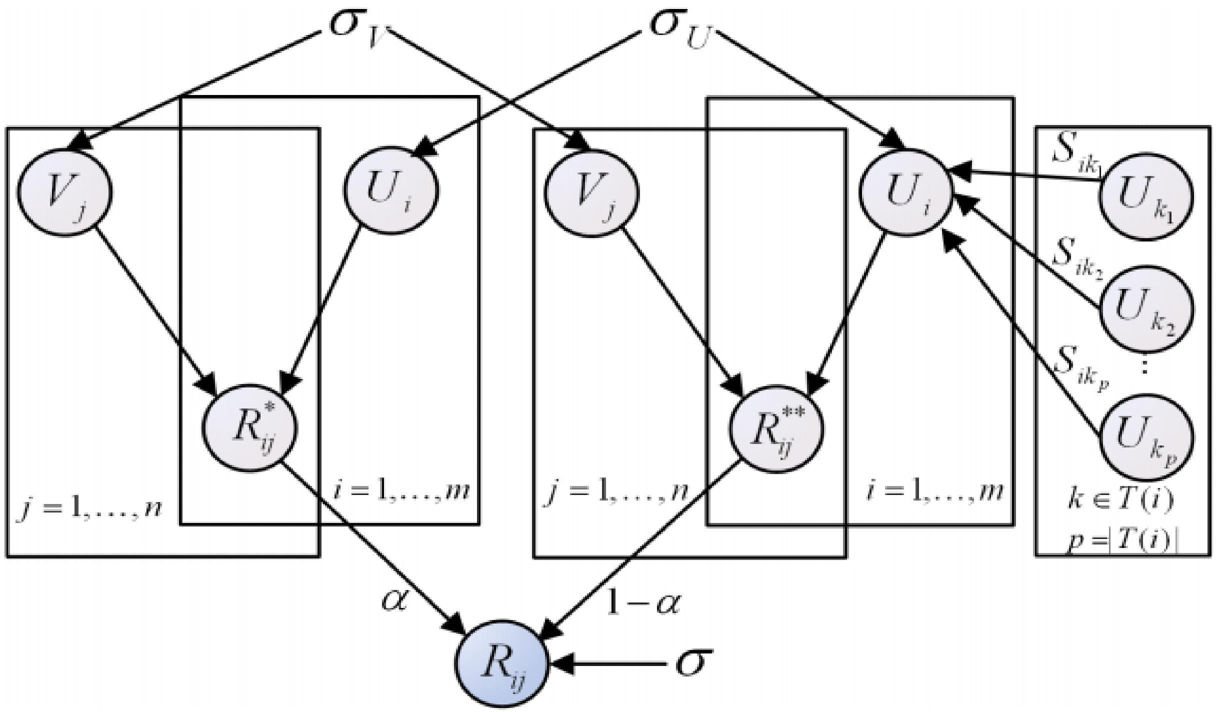

$$R_{ij} = \alpha U_i^T V_j + (1-\alpha) \sum_{k \in \mathcal{T}(i)} S_{ik} U_k^T V_j)$$

Social trust ensemble ... [Ma et al *SIGIR' 2009*]

- Experiments
  - Epinion data: a trust network + a rating matrix

Table 1: Statistics of User-Item Rating Matrix of Epinions

Statistics	User	Item
Max. Num. of Ratings	1960	7082
Avg. Num. of Ratings	12.21	7.56

Table 2: Statistics of Social Trust Network of Epinions

Statistics	Trust per User	Be Trusted per User
Max. Num.	1763	2443
Avg. Num.	9.91	9.91

Social trust ensemble ... [Ma et al *SIGIR' 2009*]

- Experiments
  - Epinion data: a trust network + a rating matrix
  - Metrics: <u>*RMSE, MAE*</u> for rating prediction
  - Comparison: PMF, trust-only, trust ensemble

Social trust ensemble … [Ma et al *SIGIR' 2009*]

- Experiments
  - Results

Training Data	Metrics	Dimensionality = 5				Dimensionality = 10			
		Trust	PMF	SoRec	RSTE	Trust	PMF	SoRec	RSTE
90%	MAE	0.9054	0.8676	0.8442	**0.8377**	0.9039	0.8651	0.8404	**0.8367**
	RMSE	1.1959	1.1575	1.1333	**1.1109**	1.1917	1.1544	1.1293	**1.1094**
80%	MAE	0.9221	0.8951	0.8638	**0.8594**	0.9215	0.8886	0.8580	**0.8537**
	RMSE	1.2140	1.1826	1.1530	**1.1346**	1.2132	1.1760	1.1492	**1.1256**

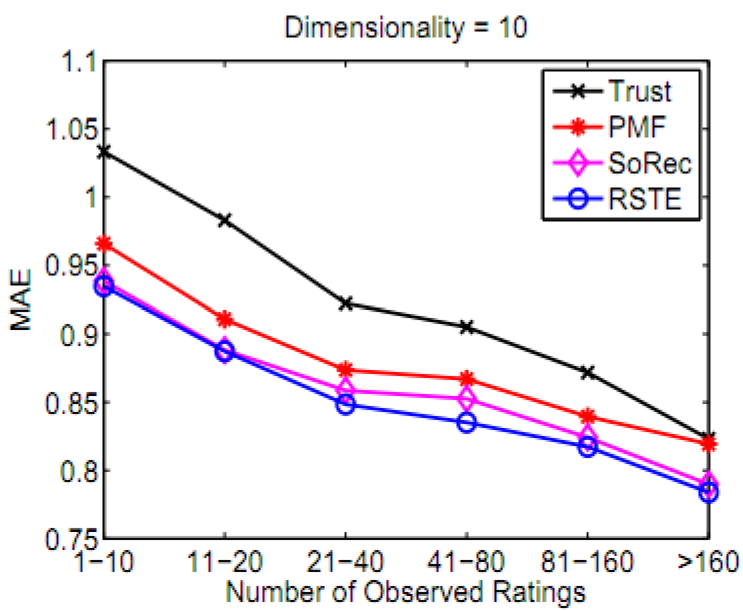

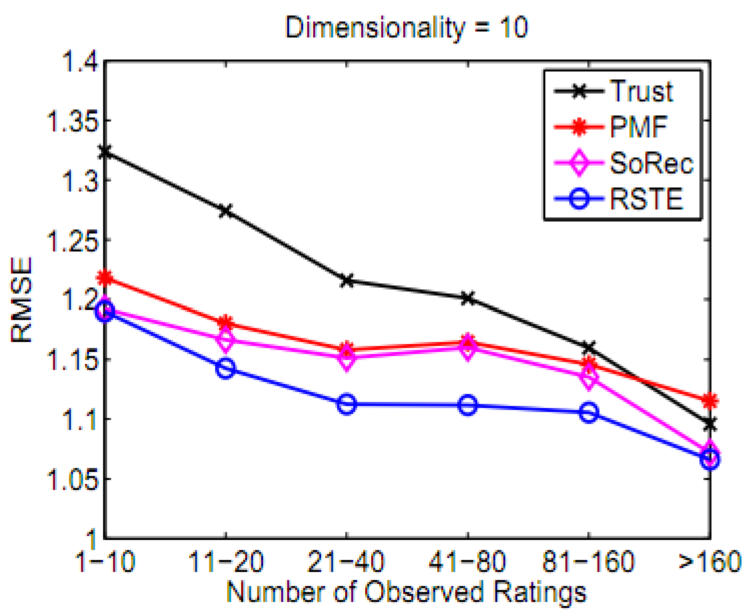

Representative work

- Homophily
  - [Ma et al, SIGIR' 2009]
    Learning to Recommend with Social Trust Ensemble
  - [Jamali & Ester et al, RecSys' 2010]
    A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks
  - [Yang et al, WWW' 2011]
    Like like alike -- Joint friendship and interest propagation in social networks
- Influence
  - [Ma et al, RecSys' 2009]
    Learning to Recommend with Trust and Distrust Relationships

MF with trust propagation ... [Jamali & Ester *RecSys' 2010*]

- Feature vectors of neighbors should influence the feature vect

Bingo, best paper award!

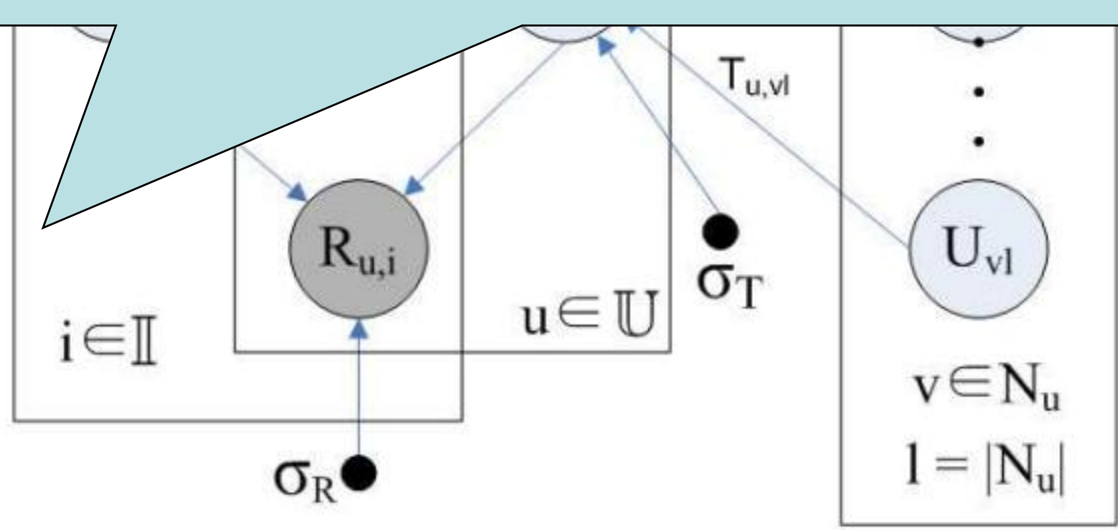

$$\boxed{\widehat{U}_u = \sum_{v \in N_u} T_{u,v} U_v}$$

MF with trust propagation … [Jamali & Ester *RecSys' 2010*]

- Experiments
  - Results on Epinion and Flixster

Method	K=5	K=10
CF	1.180	1.180
BaseMF	1.175	1.195
STE	1.145	1.150
SocialMF	1.075	1.085

Method	K=5	K=10
CF	0.911	0.911
BaseMF	0.878	0.863
STE	0.864	0.852
SocialMF	0.821	0.815

Method	Epinions	Flixster
CF	1.361	1.228
BaseMF	1.352	1.213
STE	1.295	1.152
SocialMF	1.159	1.057

RMSE values on cold start users (K=5).

Representative work

- Homophily
  - [Ma et al, SIGIR' 2009]

    Learning to Recommend with Social Trust Ensemble
  - [Jamali & Ester et al, RecSys' 2010]

    A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks
  - [Yang et al, WWW' 2011]

    Like like alike -- Joint friendship and interest propagation in social networks
- Influence
  - [Ma et al, RecSys' 2009]

    Learning to Recommend with Trust and Distrust Relationships

Like like alike ... [Yang et al *WWW' 2011*]

- Tie strength

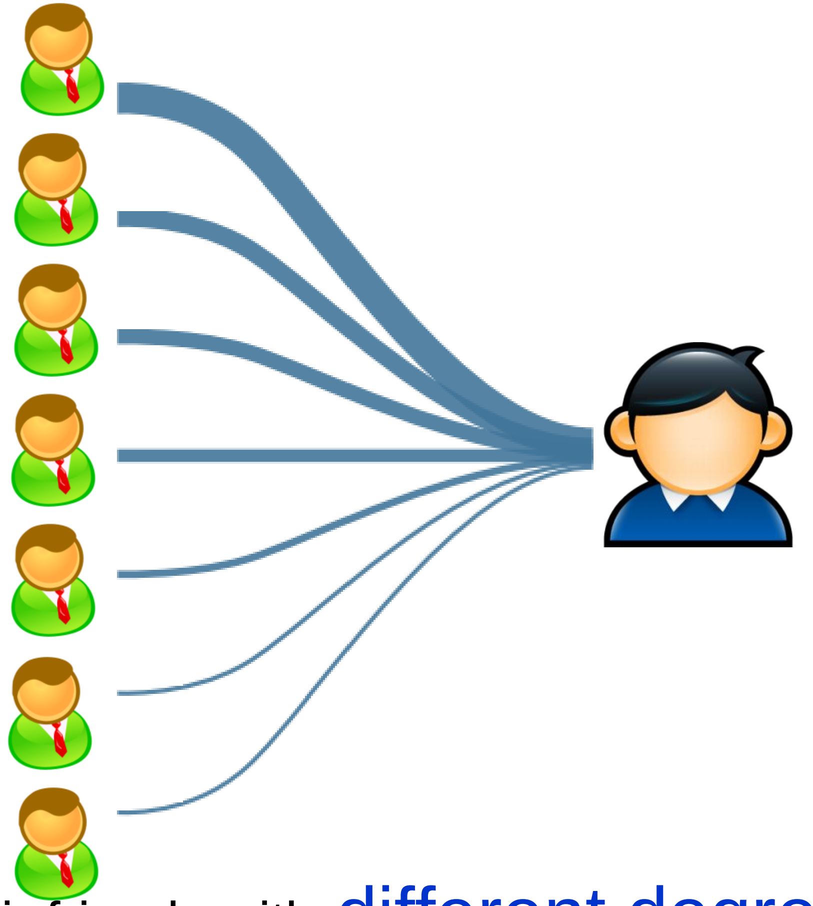

users connect to their friends with different degree of kinship

Like like alike ... [Yang et al *WWW' 2011*]

- Social network = friend network + interest network

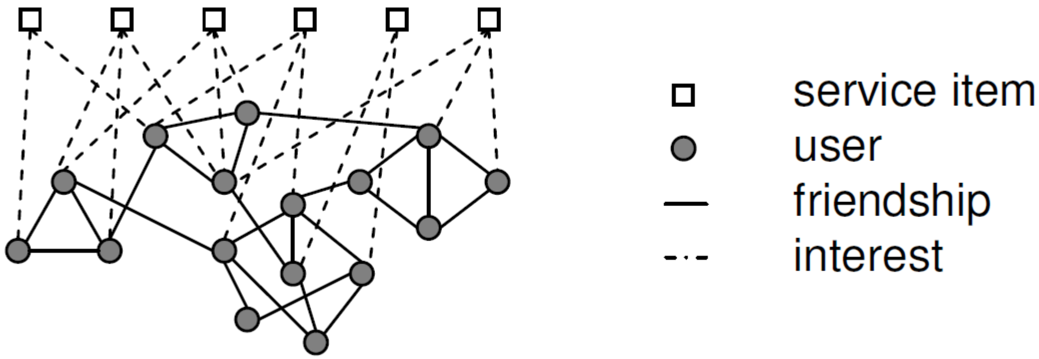

users connect to their friends

users interact with service items (applications, ads, games, movies,...)

Two key tasks

- Friendship Propagation

  connecting people to *real friends*

  – boost traffic & user population, make the social graph denser ...

- Interest Propagation

  targeting services to people *interested*

  – boost revenue, increase user participation, make the interest graph denser ...

These two tasks are usually addressed *separately with different methodologies*.

Homophily:

- The social effect:
  - People connected to each other tend to have similar interest;
  - People with similar interest are more likely to be friends.

- Hints:

  Freindship and *interest* evidences are
  - highly correlated (Y! pulse: higher interest-correlation between connected users)
  - mutually reinforcing if modeled jointly

Friendship and interest should be propagated **jointly**!

Friendship-Interest Propagation (FIP)

Exploit *Homophily* to established an *integrated network* for ***joint*** propagation of friendship and interest.

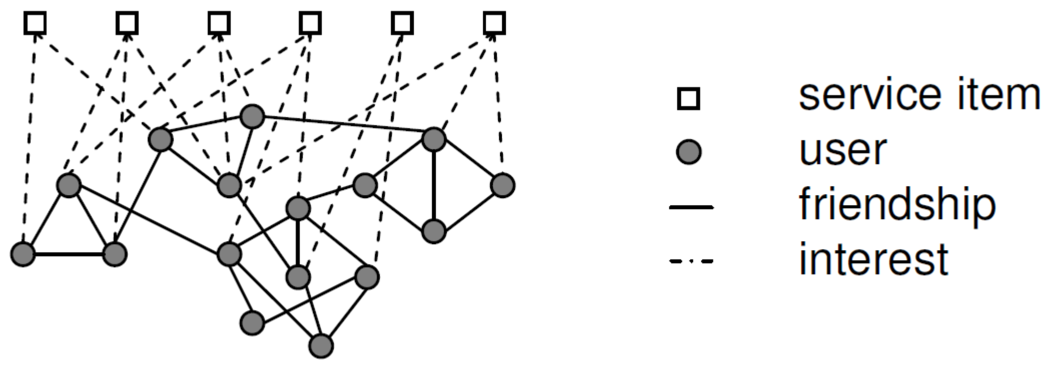

The FIP Model

- Modeling interests: *collaborative filtering*

$$\phi_i \sim p(\phi_i | x_i)$$

$$\phi_j \sim p(\phi_j | x_j)$$

$$y_{ij} \sim p(y_{ij} | \phi_i, \phi_j, x_i, x_j, \Theta)$$

i: user
j: item
y: interest indication
φ: latent profiles
x_i: user features (age, gender, income)
x_j: item features (words, visual features)

The FIP Model

- Modeling friendship: *latent-factor-based random walk*

$$\phi_i \sim p(\phi_i | x_i)$$
$$s_{ii'} \sim p(s_{ii'} | \phi_i, \phi_j, x_i, x_j, \Theta)$$

i, i': user
s: friendship connection
φ: latent profiles
x_i: user features (age, gender, income)

The FIP Model

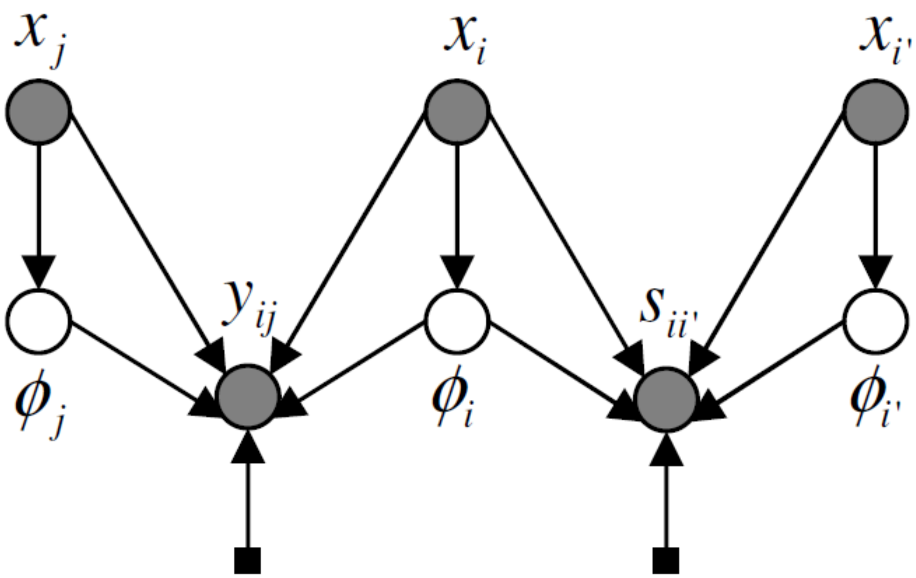

i: user
j: item
y: interest indication
s: friendship connection
φ: latent profiles
x_i: user features (age, gender, income)
x_j: item features (words, visual features)

The Friendship-Interest Propagation (FIP) model.	
$\forall\ i \in \mathcal{I}$	$\phi_i \sim p(\phi_i \mid x_i, \Theta)$
$\forall\ j \in \mathcal{J}$	$\phi_j \sim p(\phi_j \mid x_j, \Theta)$
$\forall\ i \in \mathcal{I}, j \in \mathcal{J}$	$y_{ij} \sim p(y_{ij} \mid \phi_i, \phi_j, x_i, x_j, \Theta)$
$\forall\ i, i' \in \mathcal{I}$	$s_{ii'} \sim p(s_{ii'} \mid \phi_i, \phi_{i'}, x_i, x_{i'}, \Theta)$

The FIP model

- Model specification:

$$\phi_i = Ax_i + \epsilon_i \qquad \phi_j = Bx_j + \epsilon_j$$

$$y_{ij} \sim p(y_{ij}|f_{ij}) \text{ where } f_{ij} = \phi_i^\top \phi_j + x_i^\top W x_j$$

$$s_{ii'} \sim p(s_{ii'}|h_{ii'}) \text{ where } h_{ii'} = \phi_i^\top \phi_{i'} + x_i^\top M x_{i'}$$

Optimization

- Overall objective:

$$\min \lambda_y \boxed{\sum_{i,j \in \mathcal{O}_{ij}} \ell(y_{ij}, f_{ij}) + \lambda_s \sum_{(i,i') \in O_{i,i'}} \ell(s_{ii'}, h_{ii'})}$$

Dyadic factorization

$$+ \boxed{\lambda_{\mathcal{I}} \sum_i \gamma(\phi_i | x_i) + \lambda_{\mathcal{J}} \sum_j \gamma(\phi_j | x_j)}$$

Content factorization

$$+ \boxed{\lambda_W \Omega[W] + \lambda_M \Omega[M] + \lambda_A \Omega[A] + \lambda_B \Omega[B]},$$

Regularization

Optimization

- Loss functions

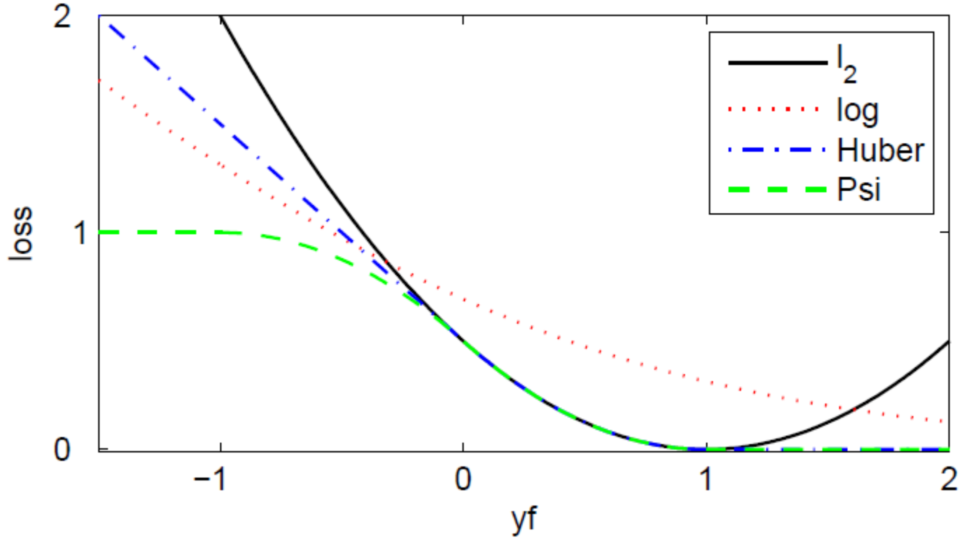

Figure 3: Least mean squares (ℓ_2), logistic (log), Langford-Huber (Huber) and Ψ-loss (Psi). We use these four and the lazy ℓ_2 (omitted since its shape in parameter space is essentially identical to ℓ_2) loss for binary classification.

- Regularizer
  – L2, L1, Ky-Fan, etc

- Bias Correction
  - Observations (for both interest and friendship) are sparse with exclusively positive interactions
  - Absence of negatives leads to inevitable overfitting, e.g., all the incoming dyadic interactions are predicted *positive*
  - *Selection bias correction:* treat missing observations as very-weak negative observations:

For every positive observation, e.g. $y_{ij} = 1$, we randomly sample a handful set of missing (unobserved) entries $\{y_{ij'}\}_{j'=1:m}$ and treat them as negative examples (e.g. $y_{ij'} = -1$,) with credibility $1/m$ each. Since the sampling procedure is random during the SGD process, the set of pseudo-negatives changes at each iteration and consequently each missing entry is treated as a potentially *very weak* negative instance.

Experiments

- Data
  - A subset of Yahoo! Pulse data.
  - 1.2M users, 386 items
  - 6.1M friend connections
  - 29M interest indications

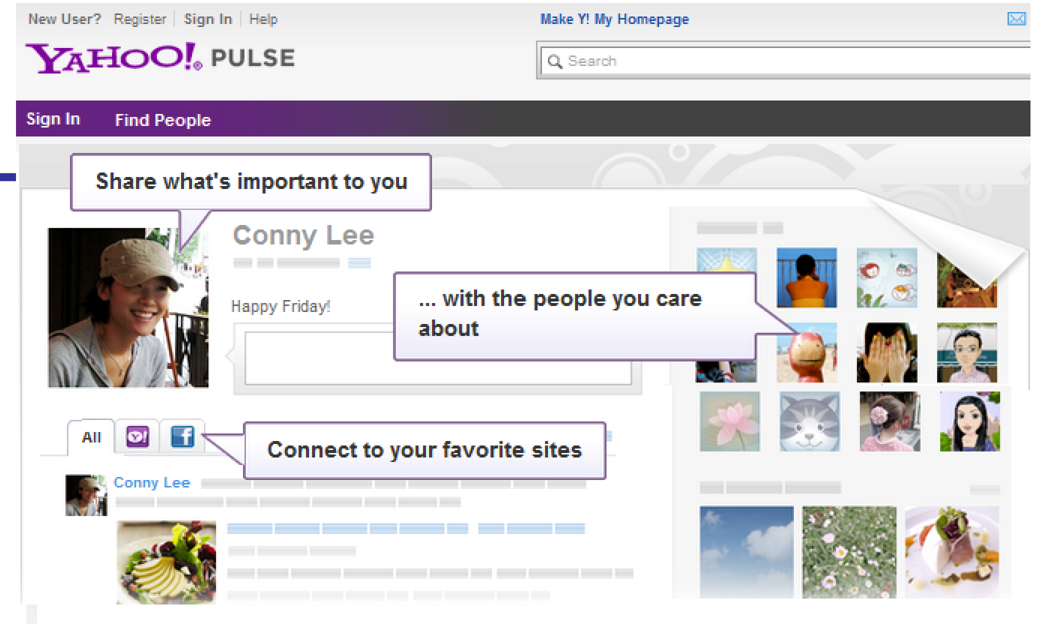

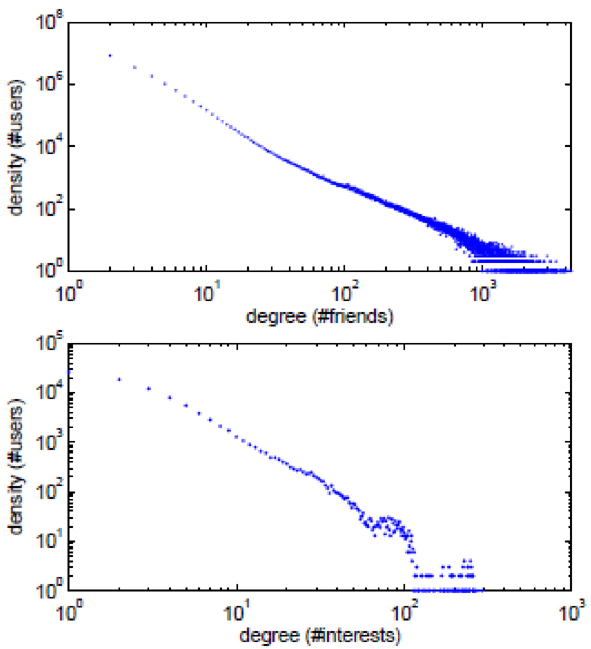

Figure 4: Degree distributions of Yahoo! Pulse friendship (top) and interest (bottom) networks.

Experiments

- Interest propagation [in terms of service recommendation]

Table 1: Service recommendation performance.

Models	loss	$\Omega[\cdot]$	AP@5	AR@5	nDCG@5
SIM			0.630	0.186	0.698
RLFM			0.729	0.211	0.737
NLFM			0.748	0.222	0.761
FIP	ℓ_2	ℓ_2	0.768	0.228	0.774
FIP	lazy ℓ_2	ℓ_2	0.781	0.232	0.790
FIP	logistic	ℓ_2	0.781	0.232	0.793
FIP	Huber	ℓ_2	0.781	0.232	0.794
FIP	Ψ	ℓ_2	0.777	0.231	0.771
FIP	ℓ_2	ℓ_1	0.778	0.231	0.787
FIP	lazy ℓ_2	ℓ_1	0.780	0.231	0.791
FIP	logistic	ℓ_1	0.779	0.231	0.792
FIP	Huber	ℓ_1	**0.786**	**0.233**	**0.797**
FIP	Ψ	ℓ_1	0.765	0.215	0.772

Experiments

- Friendship Propagation [in terms of friend suggestion]

Table 2: Friendship prediction performance.

Models	loss	$\Omega[\cdot]$	AP@5	AR@5	nDCG@5
RLFM			0.164	0.202	0.174
FIP	ℓ_2	ℓ_2	**0.359**	**0.284**	**0.244**
FIP	lazy ℓ_2	ℓ_2	0.193	0.269	0.200
FIP	logistic	ℓ_2	0.174	0.220	0.189
FIP	Huber	ℓ_2	0.210	0.234	0.215
FIP	Ψ	ℓ_2	0.187	0.255	0.185
FIP	ℓ_2	ℓ_1	0.186	0.230	0.214
FIP	lazy ℓ_2	ℓ_1	0.180	0.223	0.194
FIP	logistic	ℓ_1	0.183	0.217	0.189
FIP	Huber	ℓ_1	0.188	0.222	0.200
FIP	Ψ	ℓ_1	0.178	0.208	0.179

Experiments

- Bias correction

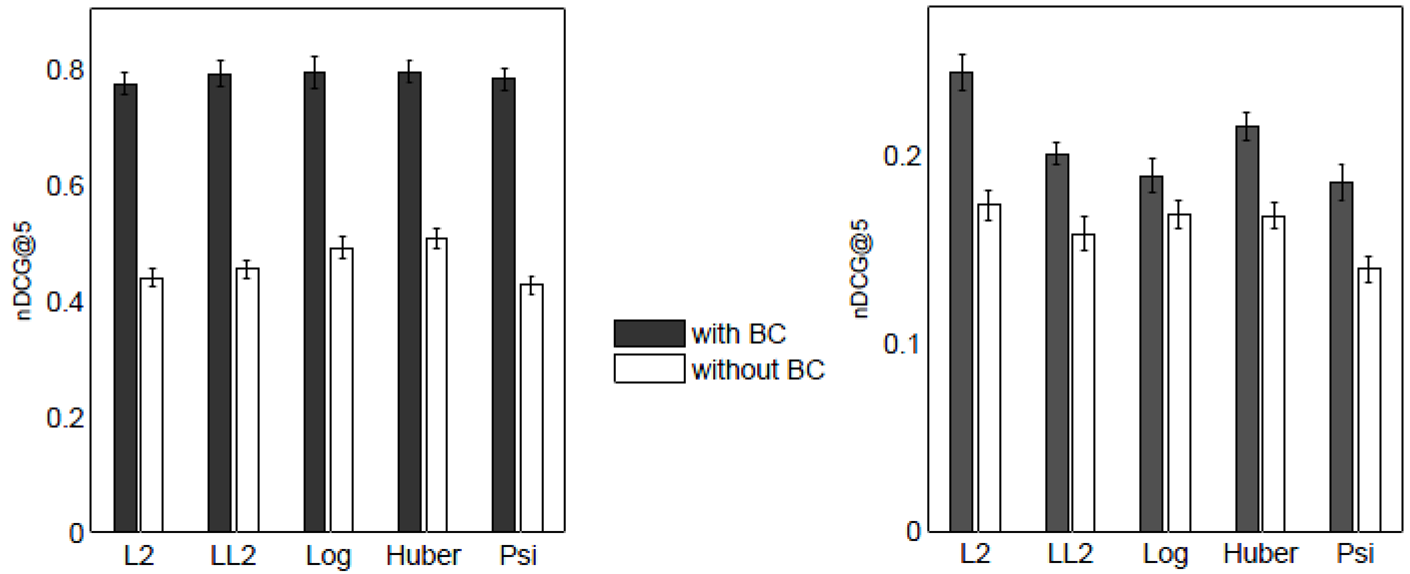

Figure 7: Recommendation performance in terms of nDCG@5 with and without bias-correction (BC) when applied to service recommendation (left) and friendship prediction (right).

Representative work

- Homophily
  - [Ma et al, SIGIR' 2009]
    Learning to Recommend with Social Trust Ensemble
  - [Jamali & Ester et al, RecSys' 2010]
    A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks
  - [Yang et al, WWW' 2011]
    Like like alike -- Joint friendship and interest propagation in social networks
- Influence
  - [Ma et al, RecSys' 2009]
    Learning to Recommend with Trust and Distrust Relationships

Recommendation w trust/distrust [Ma et al *2009*]

- Trust graph + distrust graph + rating matrix

	v_1	v_2	v_3	v_4	v_5
u_1	5		3		5
u_2				1	
u_3		4			
u_4	3		4	2	
u_5		5			4

(b) User-Item Rating Matrix

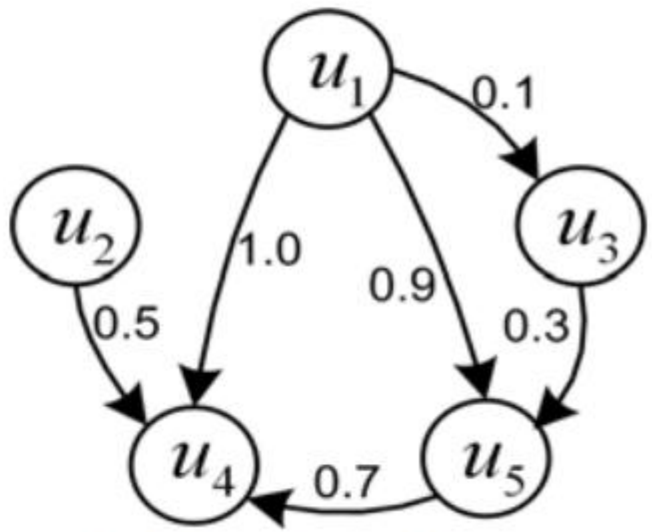

(c) User Trust Graph

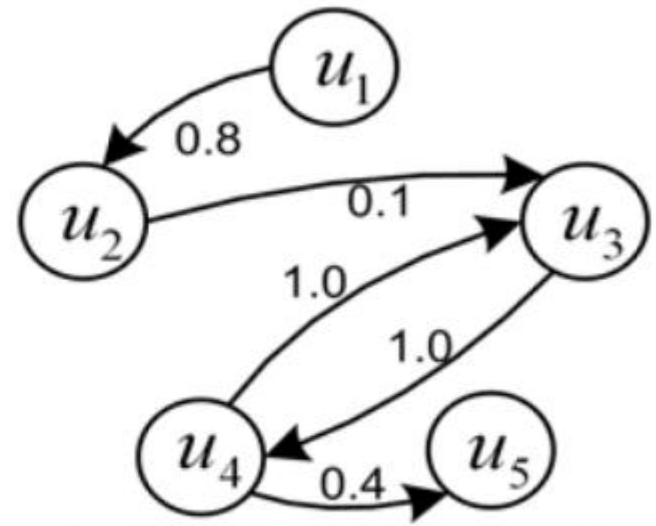

(d) User Distrust Graph

Recommendation w trust/distrust [Ma et al *2009*]

- matrix factorization

$$
\begin{aligned}
\min_{U,V} \mathcal{L}(R,U,V) &= \frac{1}{2}\sum_{i=1}^{m}\sum_{j=1}^{n} I_{ij}^{R}(R_{ij} - g(U_i^T V_j))^2 \\
&+ \frac{\lambda_U}{2}\|U\|_F^2 + \frac{\lambda_V}{2}\|V\|_F^2,
\end{aligned}
$$

Recommendation w trust/distrust [Ma et al *2009*]

- matrix factorization + trust regularization

$$
\begin{aligned}
\min_{U,V} \mathcal{L}(R,U,V) &= \frac{1}{2}\sum_{i=1}^{m}\sum_{j=1}^{n} I_{ij}^{R}(R_{ij} - g(U_i^T V_j))^2 \\
&+ \frac{\lambda_U}{2}\|U\|_F^2 + \frac{\lambda_V}{2}\|V\|_F^2, \\
&\boxed{+ \quad \frac{\alpha}{2}\sum_{i=1}^{m}\sum_{t\in\mathcal{T}^+(i)} (S_{it}^{\mathcal{T}}\|U_i - U_t\|_F^2)}
\end{aligned}
$$

if user u_i trusts user u_t, U_i and U_d <u>should be close to each other</u>

Recommendation w trust/distrust [Ma et al *2009*]

- matrix factorization + distrust regularization

$$
\begin{aligned}
\min_{U,V} \mathcal{L}(R,U,V) &= \frac{1}{2}\sum_{i=1}^{m}\sum_{j=1}^{n} I_{ij}^{R}(R_{ij} - g(U_i^T V_j))^2 \\
&+ \frac{\lambda_U}{2}\|U\|_F^2 + \frac{\lambda_V}{2}\|V\|_F^2, \\
&\boxed{+ \frac{\beta}{2}\sum_{i=1}^{m}\sum_{d\in\mathcal{D}^+(i)} (-S_{id}^{\mathcal{D}}\|U_i - U_d\|_F^2)}
\end{aligned}
$$

if user u_i distrusts user u_d, then U_i and U_d will have a large distance

Recommendation w trust/distrust [Ma et al *2009*]

- Experiments
  - Epinion: a trust network + a distrust network + a rating matrix

Table 1: Statistics of User-Item Rating Matrix of Epinions

Statistics	User	Item
Min. Num. of Ratings	1	1
Max. Num. of Ratings	162169	1179
Avg. Num. of Ratings	102.07	17.79

Table 2: Statistics of Trust Network of Epinions

Statistics	Trust per User	Be Trusted per User
Max. Num.	2070	3338
Avg. Num.	5.45	5.45

Table 3: Statistics of Distrust Network of Epinions

Statistics	Distrust per User	Be Distrusted per User
Max. Num.	1562	540
Avg. Num.	0.94	0.94

Recommendation w trust/distrust [Ma et al *2009*]

- Experiments
  - Epinion data: a trust network + a rating matrix
  - Metrics: *RMSE* for rating prediction
  - Comparison: PMF, trust-reg, distrust-reg

Recommendation w trust/distrust [Ma et al *2009*]

- Experiments
  - Results

Dataset	Traning Data	Dimensionality	PMF	SoRec	RWD	RWT
Epinions	5%	5D	1.228	1.199	1.186	**1.177**
		10D	1.214	1.198	1.185	**1.176**
	10%	5D	0.990	0.944	0.932	**0.924**
		10D	0.977	0.941	0.931	**0.923**
	20%	5D	0.819	0.788	0.723	**0.721**
		10D	0.818	0.787	0.723	**0.720**

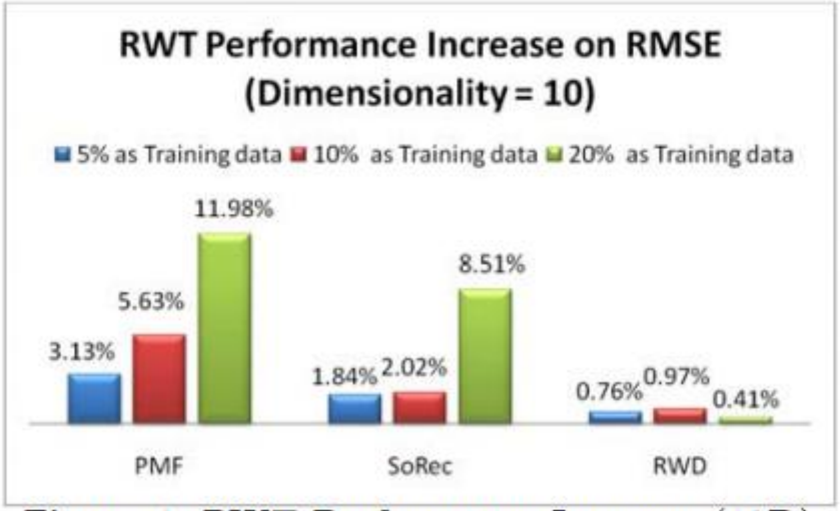

Figure 4: RWT Performance Increase (10D)

Thanks!

Any comments would be appreciated!