Seoul, Korea
WWW 2014
The 23rd International World Wide Web Conference
April 7-11, 2014 coex

Trust in Social Computing

Jiliang Tang and Huan Liu

Data Mining and Machine Learning Lab

Arizona State University

http://www.public.asu.edu/~jtang20/tTrust.htm

April 7, 2014

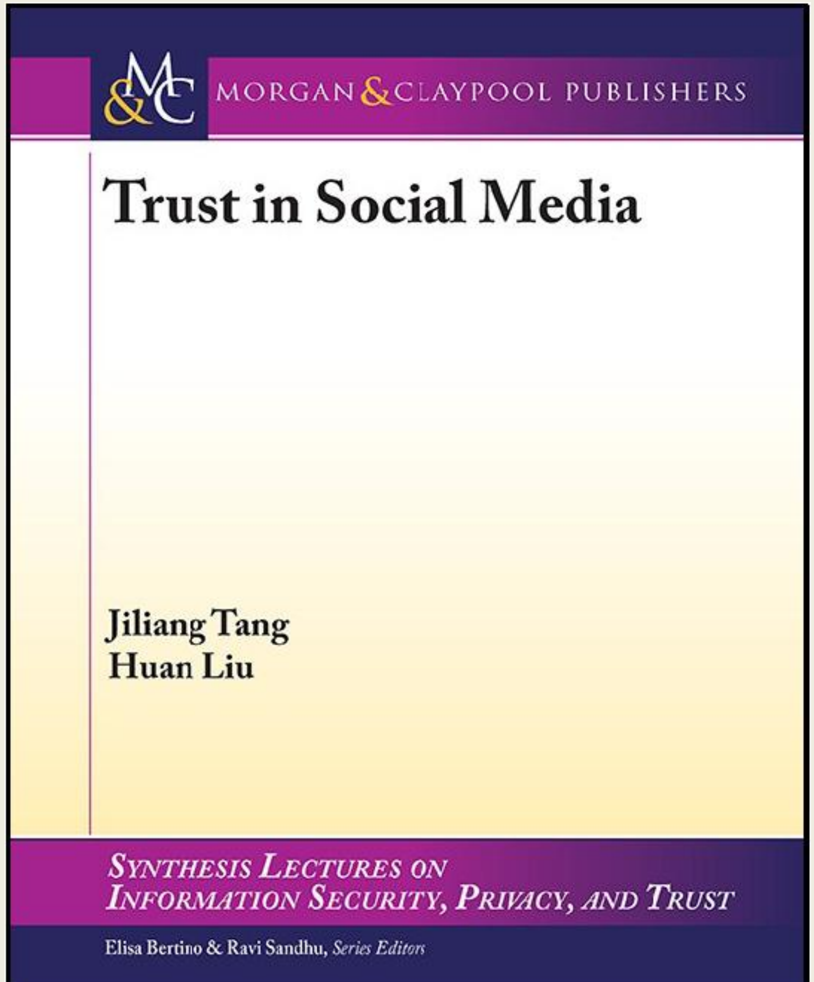
MORGAN & CLAYPOOL PUBLISHERS
Trust in Social Media
Jiliang Tang
Huan Liu
SYNTHESIS LECTURES ON
INFORMATION SECURITY, PRIVACY, AND TRUST
Elisa Bertino & Ravi Sandhu, Series Editors

Outline

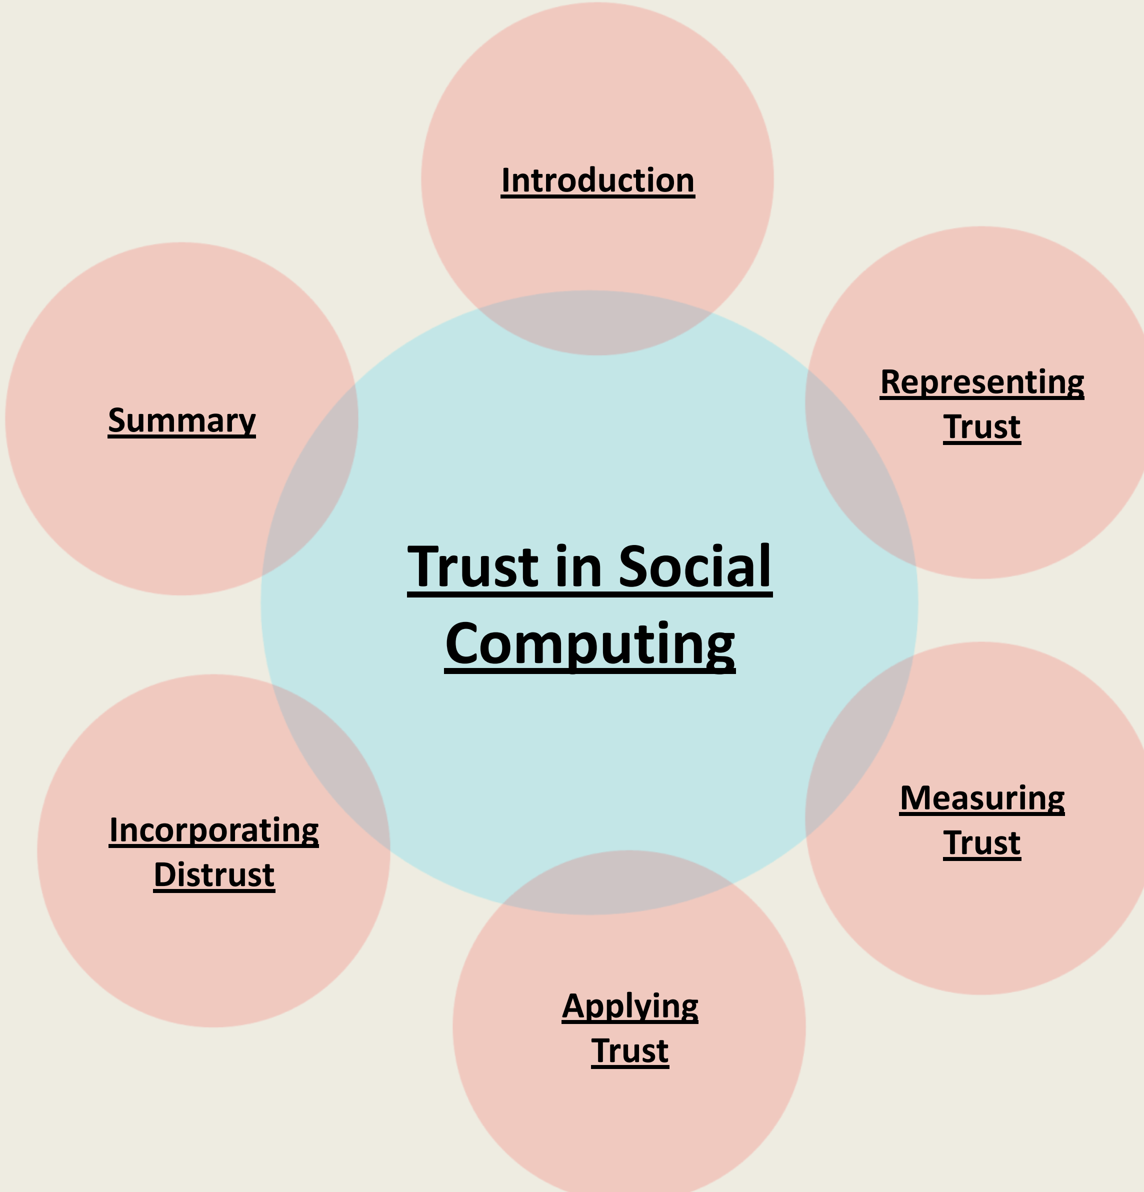
Introduction
Representing Trust
Summary
Trust in Social Computing
Measuring Trust
Incorporating Distrust
Applying Trust

Introduction

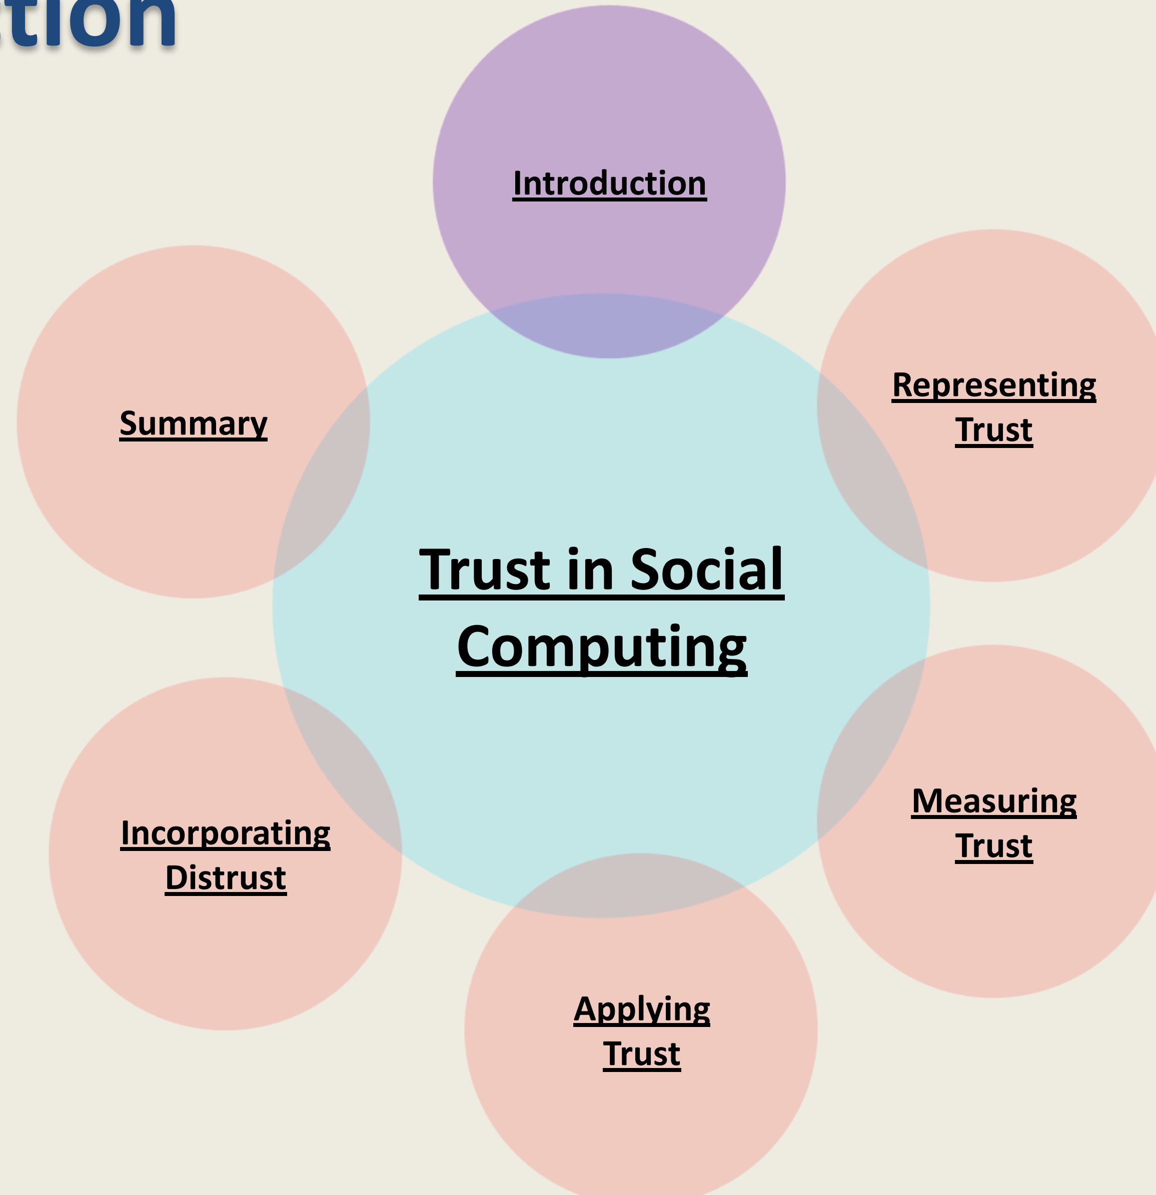

Social Media

- Social media greatly enables people to participate in online activities
  – Networking, tagging and commenting

- It shatters the barrier for online users to create and share information in any place at any time

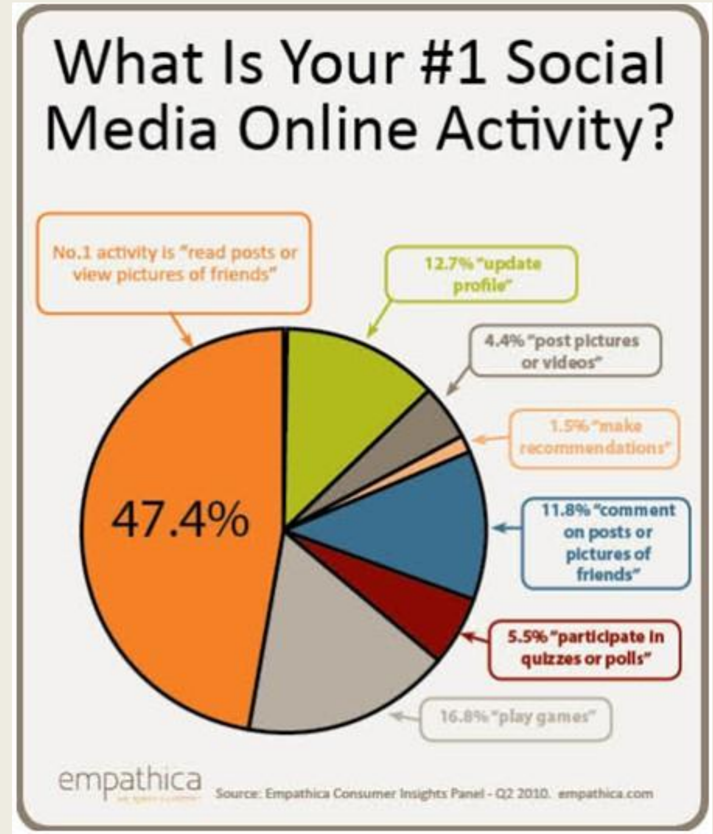

http://www.marketingprofs.com/charts/2010/4101/social-media-brand-followers-hunting-for-deals

Information Overload

- User generated content increases at an unprecedented rate
  – Given the big-data problem, how can we find relevant content?

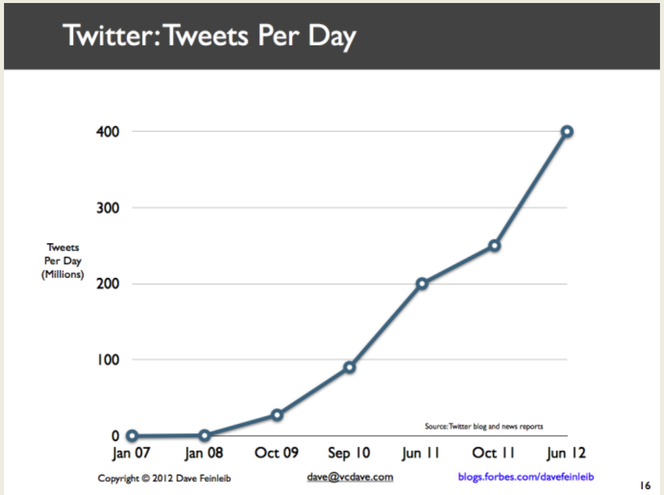

- Anyone can publish content in social media
  – With so many grass-roots authors, from whom I should collect information useful to me

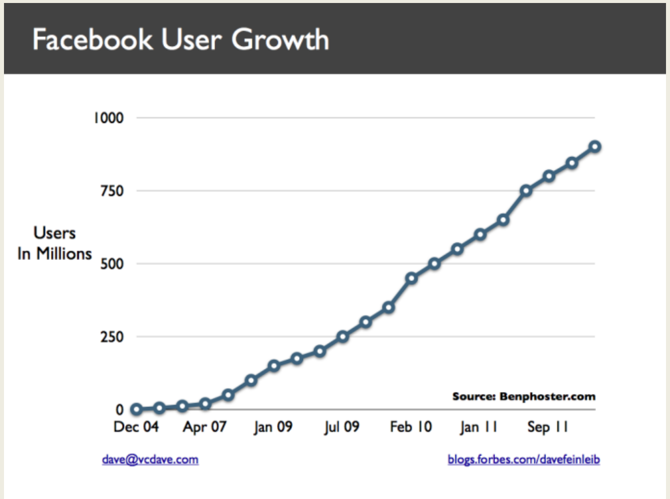

http://www.forbes.com/sites/davefeinleib/2012/07/09/the-3-is-of-big-data/

Information Credibility

- The quality of user generated content varies widely
  – From excellent content to abuse and spam
  – How to find reliable information fast

- Anyone may access my content
  – With whom should I share information?

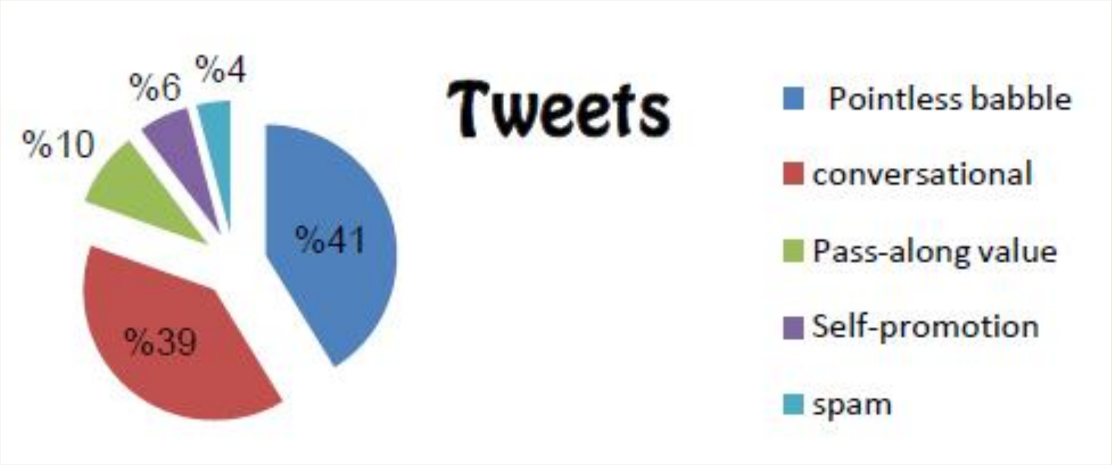

https://infomagnet.wordpress.com/2012/05/04/a-comparison-among-the-top-3-social-networks/

Trust in Social Media

- Trust provides information to find answers like
  – From whom we should collect information
  – With whom we should share information

- It offers a mental shortcut for direct information seeking
  – Without being overwhelmed by excessive information, thus mitigating information overload

  – With credible information due to the trust placed on the information provider, or increasing information credibility

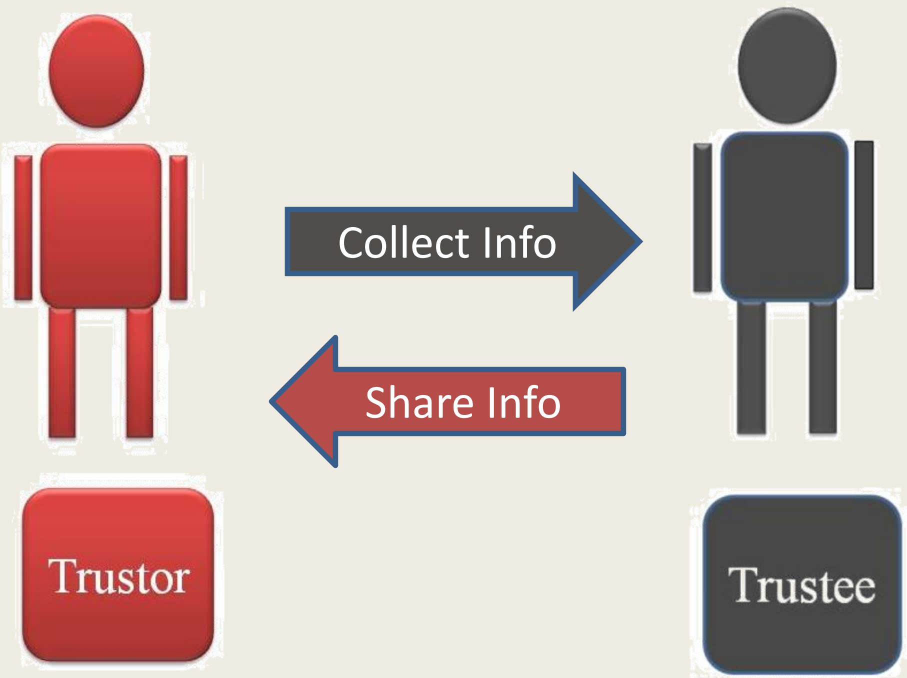

Trust – A Hard to Define Concept [McKnight et al. 2001]

- Trust is a very broad and complex concept
  – Multidisciplinary
  – More than 60 definitions in the literature

- Each discipline has its own perspective of trust
  – Personality in Psychology
  – Social structures in Sociology
  – Rational choice in Economics

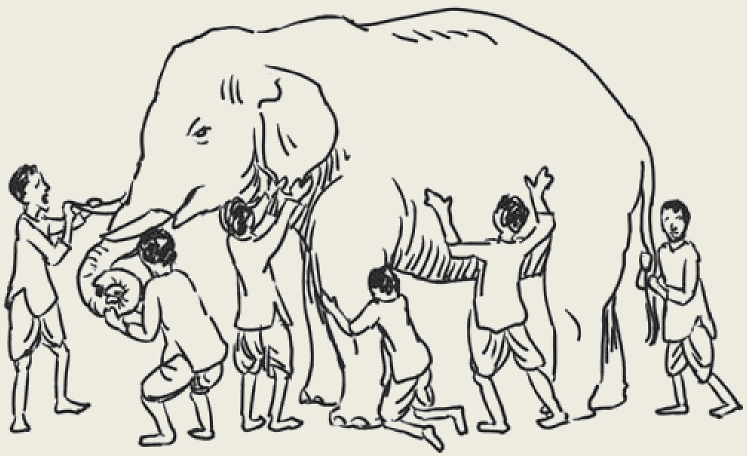

- Let's first look at what happens in social media before we settle on a definition

Online Trust Systems [Massa, 2007]

- E-marketplaces

– Selling and buying items

- Opinion and activity sharing sites

– Sharing opinions

- Business/job networking sites

– Sharing job skills

- Social/entertainment sites

– Networking and sharing UGC

- News sites (e.g., Slashdot)

– Posting news and stories

ebay

- There are two types of users - sellers and buyers
- Buyers assess the trustworthiness of sellers according to the reliability of the services or products they provide

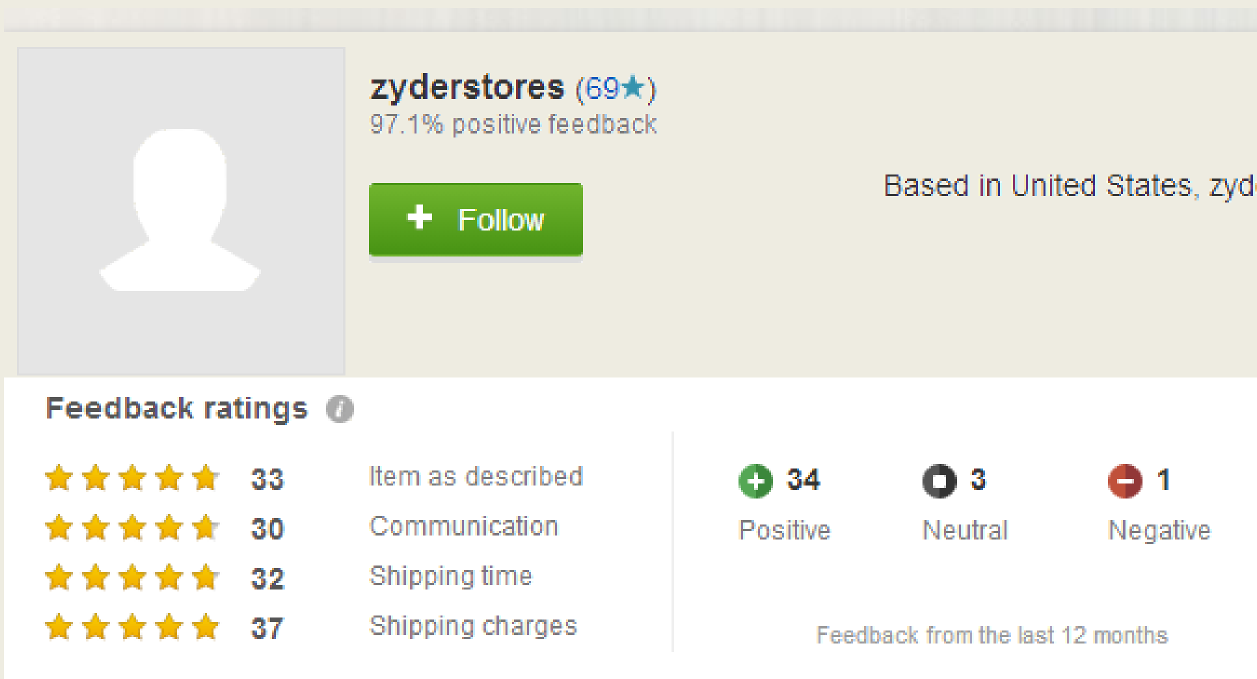

http://www.ebay.com/usr/zyderstores

Epinions.com

- There are two roles of users in Epinions – reviewers (who write reviews) and raters (who rate the helpfulness of reviews)
- Raters add a reviewer into their trust circle if they think her reviews are helpful

http://www.epinions.com/user-jankp

Advogato

- Advogato is a community site of free software developers
- Developers share their developing skills and raters will rate them with three trust levels – Master, Journeyer and Apprentice

Others have certified badvogato as follows:

- dragotown certified badvogato as Master
- beto certified badvogato as Master
- esteve certified badvogato as Journeyer
- aerry13 certified badvogato as Master
- dmitri certified badvogato as Master
- michaelemma certified badvogato as Journeyer
- sashako certified badvogato as Journeyer
- Tofu certified badvogato as Apprentice
- sulaiman certified badvogato as Journeyer
- ekashp certified badvogato as Journeyer
- hereticmessiah certified badvogato as Master
- pencechp certified badvogato as Journeyer
- wardv certified badvogato as Journeyer
- nixnut certified badvogato as Master
- garym certified badvogato as Master
- nikole certified badvogato as Master
- mirwin certified badvogato as Master
- mglazer certified badvogato as Journeyer

http://www.advogato.org/person/badvogato/

A Definition of Trust [Massa, 2007]

The explicit opinion expressed by a user about another user regarding the perceived quality of a certain characteristic of this user

– Inter-personal trust

– Network representation

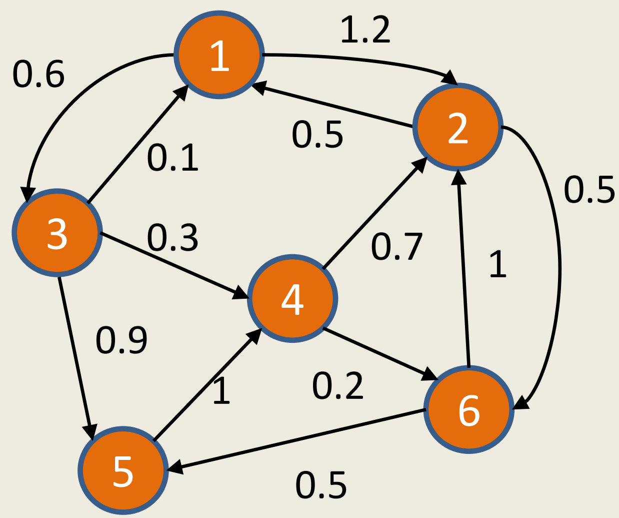

Challenges in Studying Trust in Social Computing

- Challenge 1: Social media data is based on passive observations
  – A large number of online users
  – Lack of some information other disciplines use to study trust
  – Traditional methods require interaction with users (or subjects)
  – Study trust with only passive observation

- Challenge 2: Social media data is social
  – A new type of social data
  – Big, noisy, and incomplete
  – To handle big social media data for trust research, we need effective and efficient computational tools

Computational Tasks for Trust

- The challenges from social media also offer opportunities to study trust from a computational perspective

- Four major computational tasks
  – Representing trust
  – Measuring trust
  – Applying trust
  – Incorporating distrust

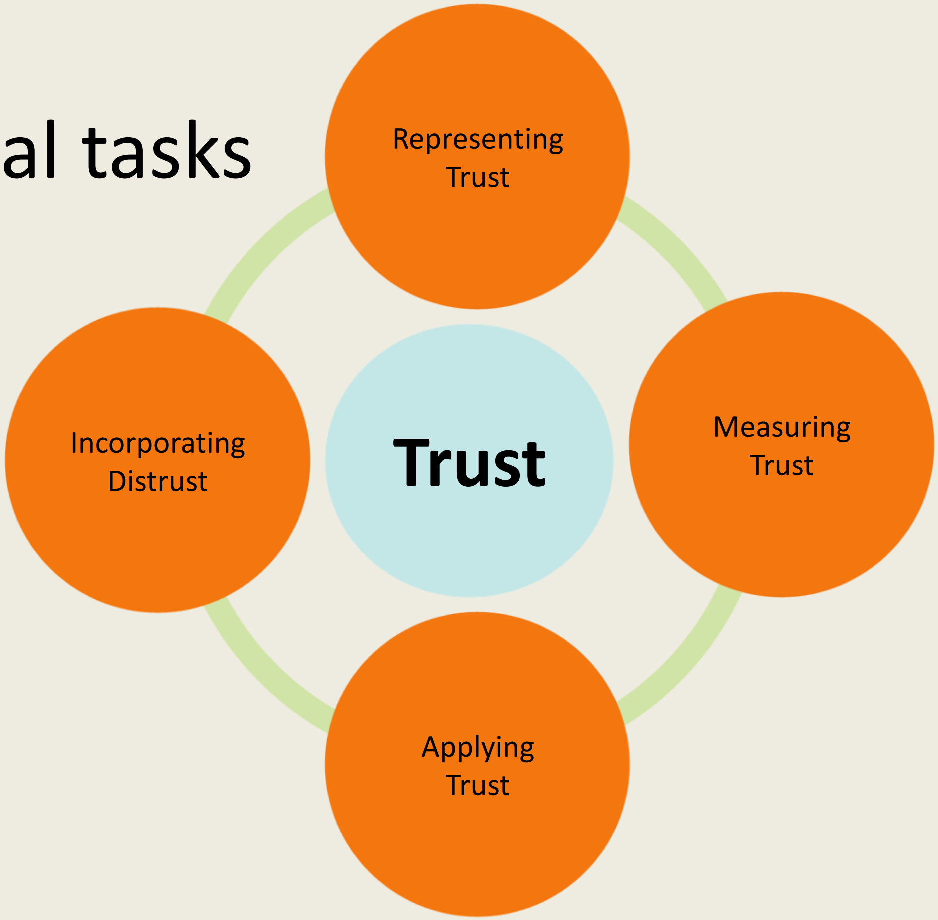

Representing Trust

- It aims to represent trust relations among users
- **Given:** a trust network
- **Output:** a mathematical representation that is computable

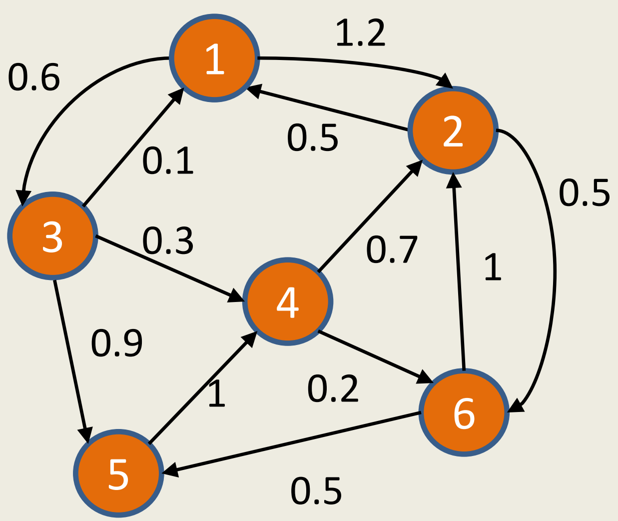

	1	2	3	4	5	6
1	0	1.2	0.6	0	0	0
2	0.5	0	0	0	0	0.5
3	0.1	0	0	0.3	0.9	0
4	0	0.7	0	0	0	0.2
5	0	0	0	1	0	0
6	0	1	0	0	0.5	0

Measuring Trust

- It aims to measure how much a user can be trusted by another user in the same trust network
- **Given:** a trust network and a user pair <u, v, ?>
- **Output:** the missing trust value is found <u, v, t>

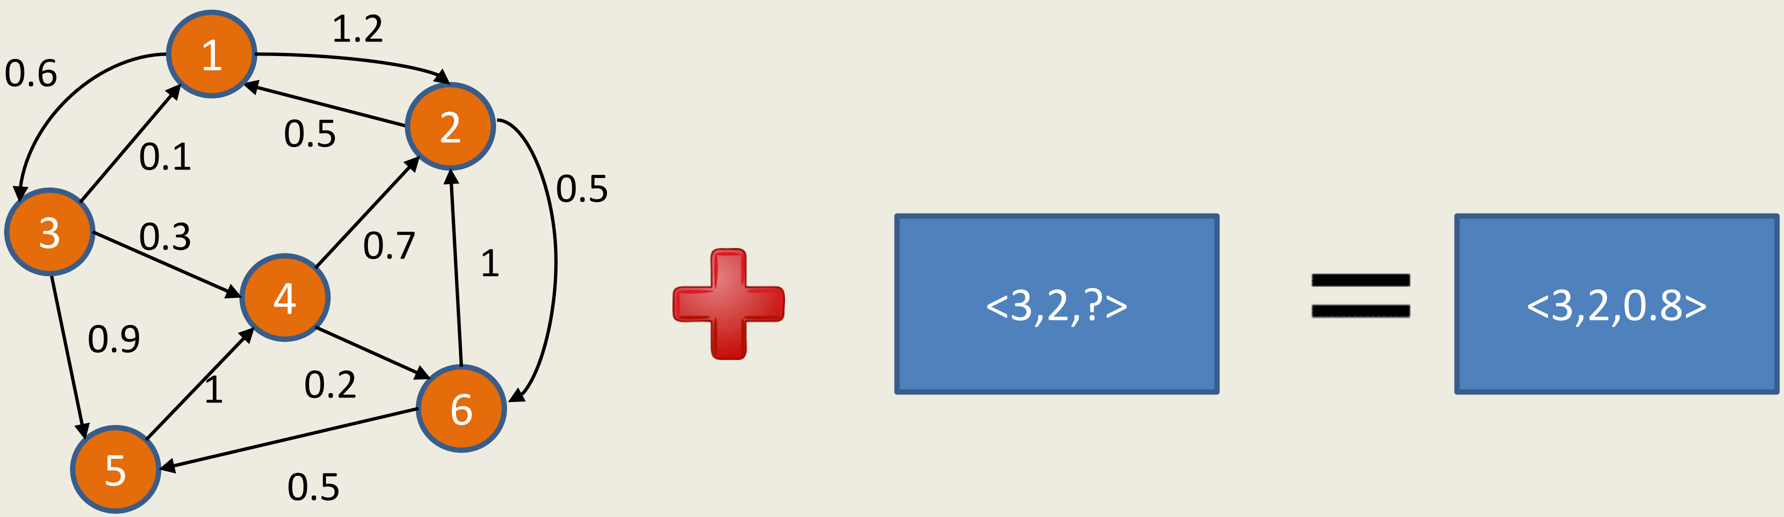

Applying Trust

- It aims to incorporate trust to facilitate online applications such as online recommendation
- Trust-aware recommendation aims to incorporate trust information in traditional recommender systems

– User-user trust information

– User-item rating information

$$\begin{bmatrix} 3 & ? & ? & 5 & ? & 4 \\ ? & 2 & 5 & 4 & 1 & ? \\ ? & ? & ? & ? & ? & ? \\ 2 & 1 & ? & 4 & 2 & ? \\ 3 & 1 & 5 & ? & 2 & ? \\ ? & ? & 4 & ? & 1 & 4 \end{bmatrix}$$

User-item rating information

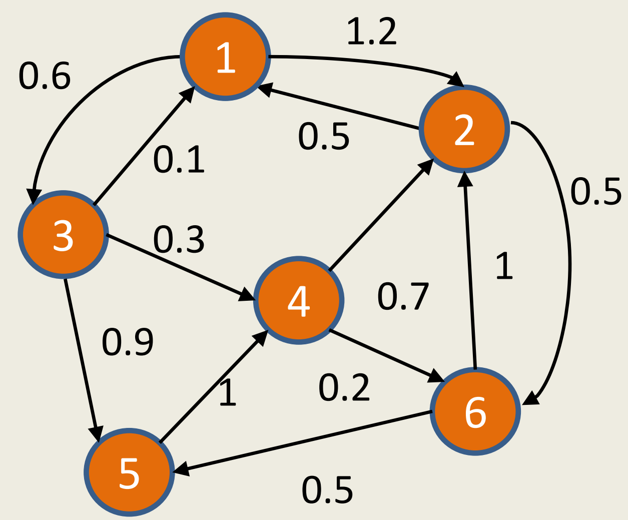

User-user trust information

Incorporating Distrust

It studies how to incorporate distrust in improving trust computation

– Computational understanding of distrust
– Representing distrust with trust
– Measuring distrust with trust
– Applying distrust to improve trust computation

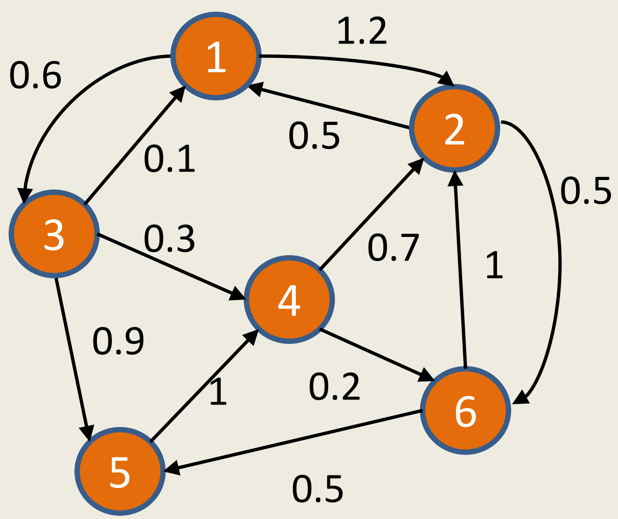

User-user trust information

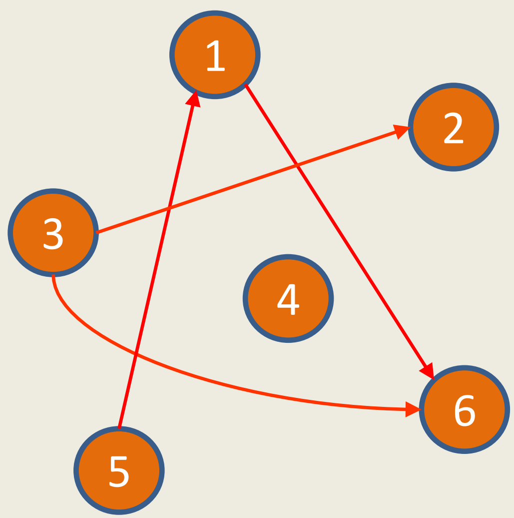

User-user distrust information

A Real-World Dataset for Studying Computational Trust

Epinions

– Trust/distrust relations
– Item ratings
– Helpfulness ratings
– Review content

Web of Trust

sleeper54 trusts:

219 hidden members

sleeper54 is trusted by:

1. brroberts2
2. tbrown
3. csmithesq
4. beeefcake
5. aohcapablanca

View all 568 members who **trust sleeper54**

Review Title	Product / Topic	Product Rating	Review Rating
Mickey and Willie: the Parallel Lives of Baseball's Golden Age	Mickey and Willie : Mantle and Mays, the Parallel Lives of Baseball's Golden Age by Allen Barra (2013, Hardcover) in Books	★★★★☆	Very Helpful

# of Users	30,455
# of Trust	363,773
# of Distrust	46,196
# of Users Receiving Distrust	9,513
# of Users Creating Distrust	5,324
# of Items	89,270
# of Ratings	562,355
Avg of Rating Score	3.9053

Trust and Distrust Distributions

- Distributions follow a power-law-like distribution
  – A typical distribution for networks in social media

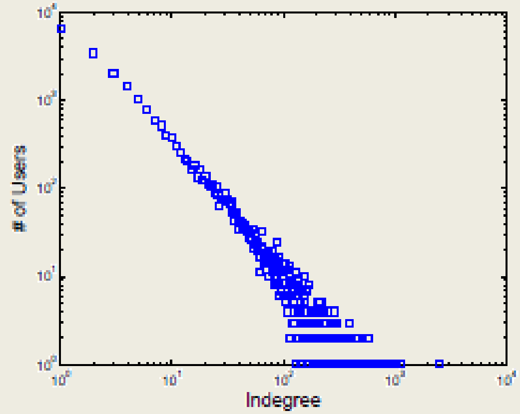

(a) Indegree of Trust

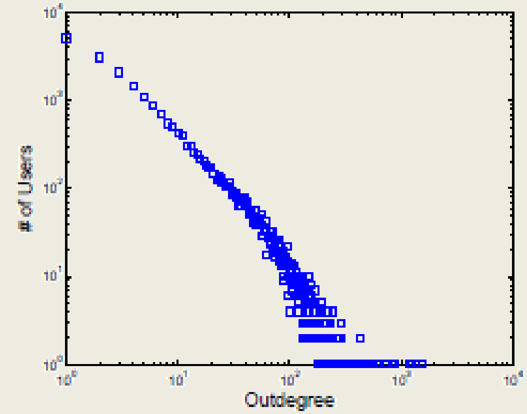

(b) Outdegree of Trust

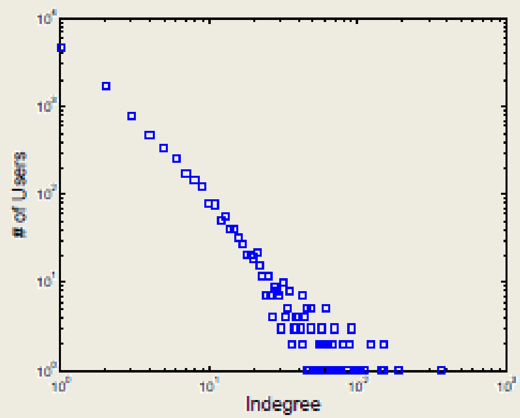

(c) Indegree of Distrust

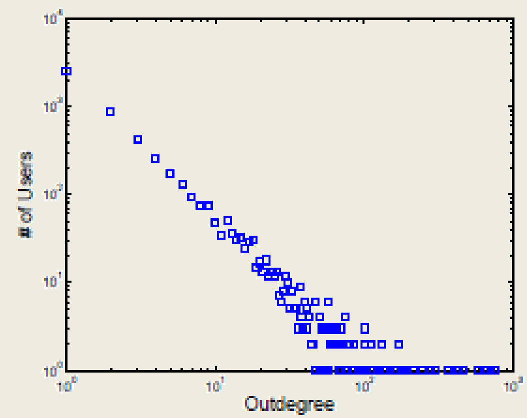

(d) Outdegree of Distrust

Representing Trust

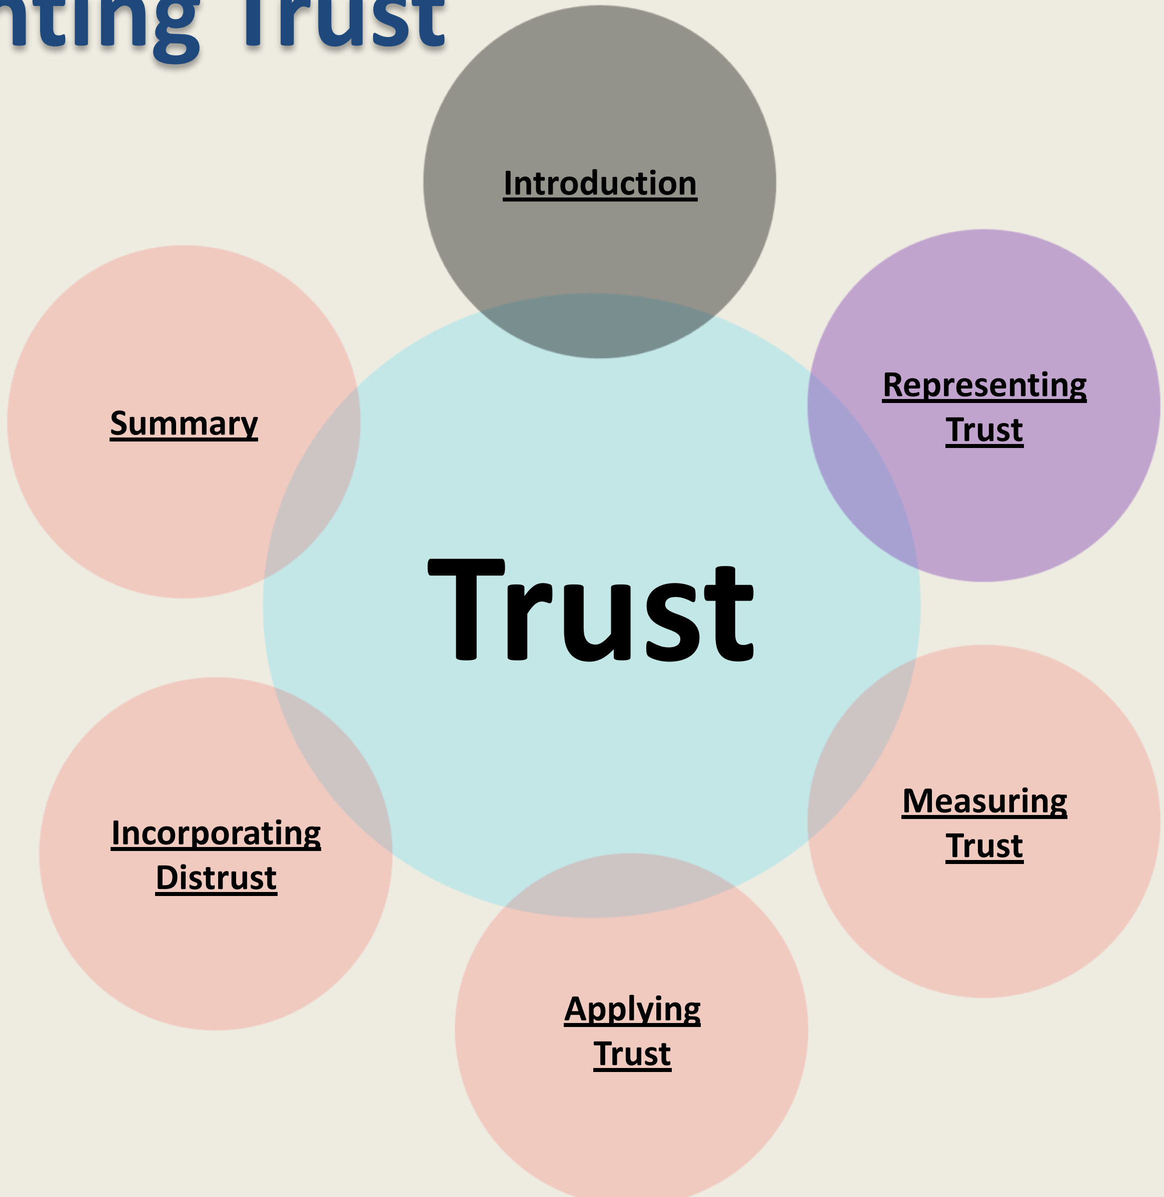

Importance of Representing Trust

- Any algorithms in measuring trust, applying trust, and incorporating distrust are based on certain trust representations

- Representing trust is the first step to make trust computable

- Properties of trust serve as the foundation of trust representations
  – Single vs multi-dimensional trust representations

- Some important properties include
  – Transitivity and composability
  – Asymmetry, and correlation with similarity

Transitivity [Golbeck, 2005]

- Transitivity allows trust to propagate along paths to reach other users
- If *u* trusts *v* and *v* trusts *w*, it can be inferred that *u* might also trust *w* to some extent

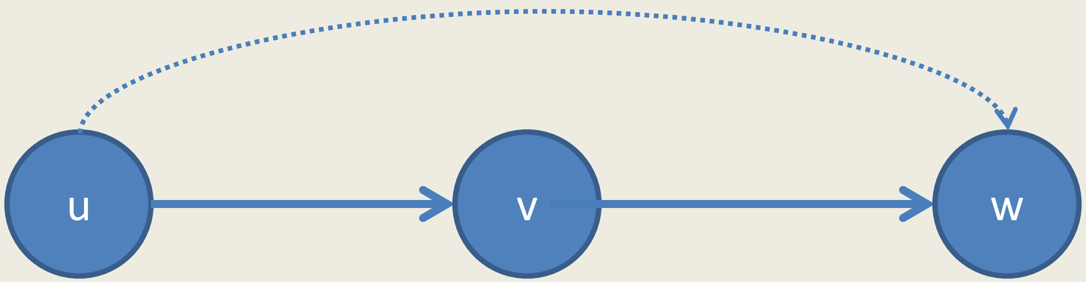

- Trust is not perfectly transitive in the mathematical sense and is conditionally transitive
  – Trust networks in social media are large
  – Users in trust networks are world-widely distributed
  – There are many pairs who do not know each other in trust networks

Transitivity Illustration

- If u+v and v+w, there are 88.34% of u and w without trust relations
- If we can observe relations between u and w, 97.75% of them are trust relations

– P1 = {# (u?w) } / {#(u+w) + #(u-w) + #(u?w)}

– P2 = {#(u+w) } / {#(u+w) + #(u-w)}

Types	Number	P1	P2
$\langle u+v, v+w \rangle$, $u?w$	25,584,525	88.34%	N.A
$\langle u+v, v+w \rangle$, $u+w$	3,320,991	11.46%	97.75%
$\langle u+v, v+w \rangle$, u-w	76,613	0.2%	2.25%

u+v, u-v, and u?v represent u and v with trust, distrust, and missing relations, respectively

Composability [Golbeck, 2005]

- Transitivity describes how trust can be passed through one path
- Composability describes that a user should combine the different trust values received from different paths

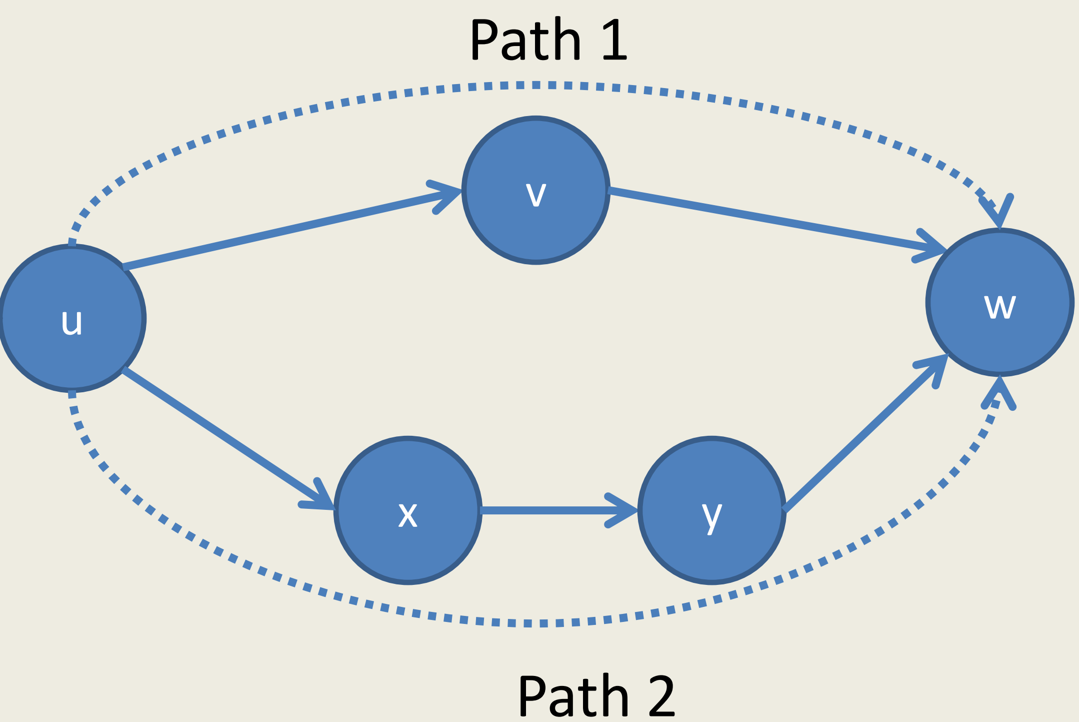

For (u,w), there are two paths (u+v+w and u+x+y+w).

u should compose trust values from both u+v+w and u+x+y+w for w

Asymmetry [Golbeck, 2005]

- For two people involved in a trust relation, trust is not necessarily identical in both directions
- The trust value from user *u* to user *v* is not necessarily equal to that from user *v* to user *u*

– One way trust for binary trust

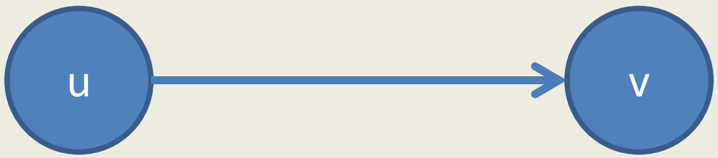

– Different trust values for continuous trust

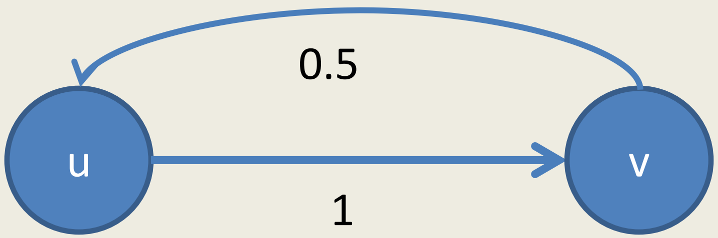

Asymmetry Illustration

- There are 37.61% of pairs of users with mutual trust relations

	$v{+}u(\%)$	$v{-}u(\%)$	$v?u(\%)$
$u{+}v$	136,806(37.61)	967(0.27)	226,000(62.13)

- Trustees who have fewer trustors are more likely to trust their trustors [Tang et al. 2012]

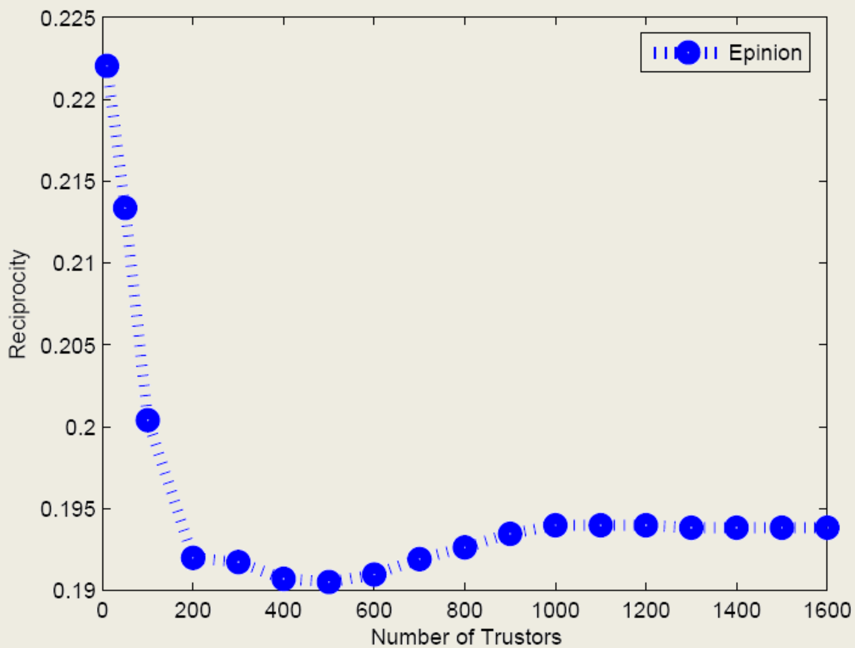

Correlation with Similarity [Ziegler and Golbeck, 2007]

- There is a strong correlation between trust and similarity
  – Users with trust relations are likely to be similar

- Social correlation theories can explain this correlation
  – Homophily, influence and confounding

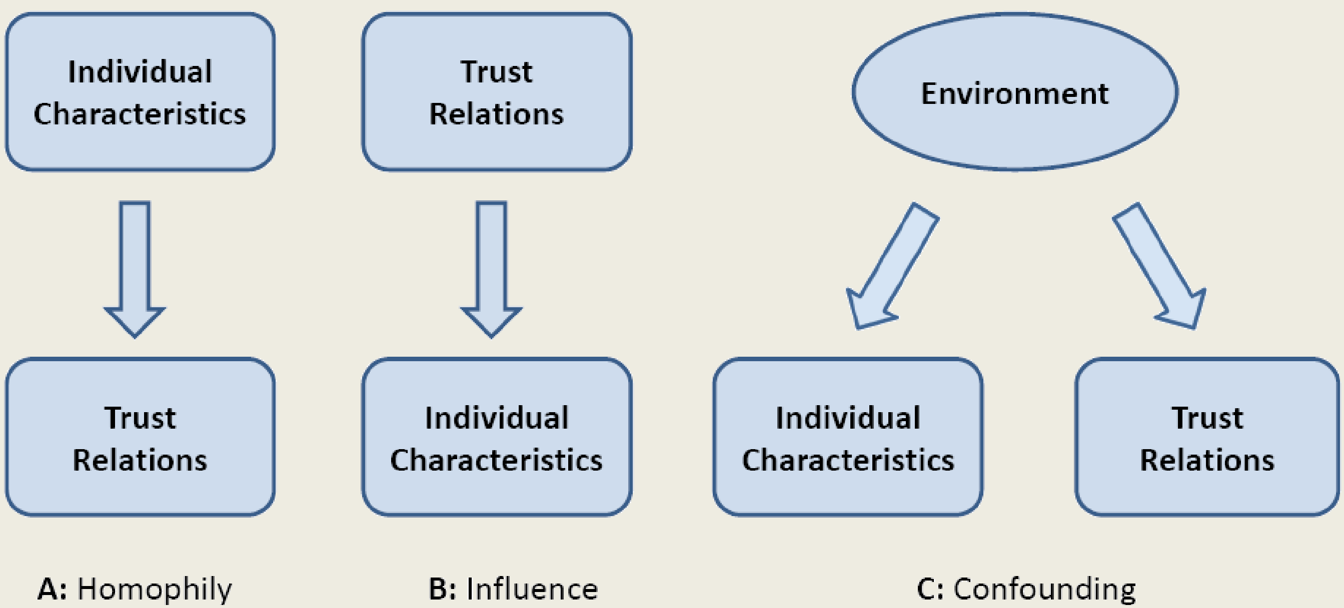

Verifying the Correlation

- For each trust relation, we calculate two similarities
  – Similarity1: trustor and trustee
  – Similarity2: trustor and a randomly chosen user

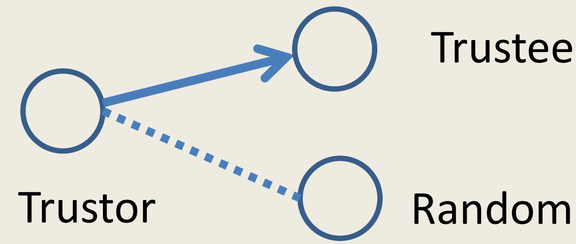

- We define two vectors **s** = { similarity1} and **t** = {similarity2}

- We conduct a two-sample t-test on **s** and **t**

H_0: **s** <= t; H_1: **s** > **t**

The null hypothesis is rejected at significance level 0.01 with p-value of 3.76e−21 in Epinions

Homophily [Tang et al., 2013]

- Similar users are more likely to establish trust relations
- Verification of homophily

– Sort trust relations based on the creation time in chronological order

– Split data into *m* pieces with equal size and time points are $\{t_1, t_2, \ldots, t_m\}$

– Divide pairs of users without trust relations until time t_i into two equal groups – high-similarity group H and low-similarity group L

– Compute the numbers of pairs creating trust relations h_i and l_i at time t_{i+1} for H and L, respectively

– $\mathbf{h} = \{h_i\}$ and $\mathbf{l} = \{l_i\}$

– A two sample t-test is conducted on **h** and **l**

$$H_0: \mathbf{h} <= \mathbf{l}; \qquad H_1: \mathbf{h} > \mathbf{l}$$

The null hypothesis is rejected at significance level 0.01 with p-value of 7.51e−64 in Epinions

Influence [Yeung and Iwata, 2011]

- Users tend to follow the behaviors of trusted users and users with trust relations are likely to exhibit similar behaviors
- Changes of users' cosine similarity 100 days before and after trust relations were established
  – Similarity increases before they trust each other, and continues to increase after that

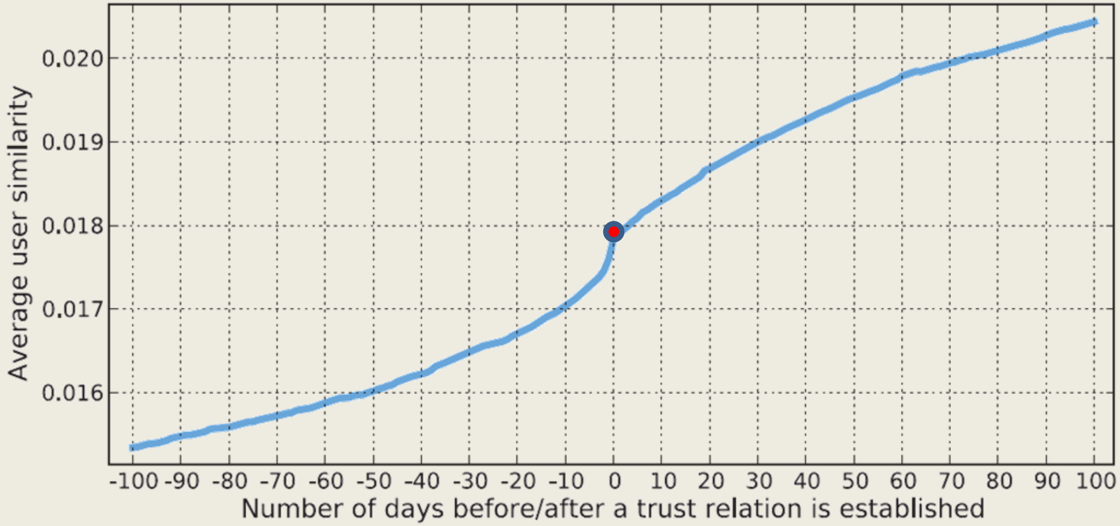

(a) Average user cosine similarity

Trust Representation Classifications

- Trust representations can be classified from different perspectives

- Probabilistic vs. gradual trust representations
  – From an interpretation perspective

Probabilistic Representations [Victor et al., 2011]

- Probabilistic representations use probabilities to indicate how much trust is placed by a user to another
  - Stronger trust corresponds to a higher probability
  - (u,v,p) represents the probability p of u trusting v
  - p = 1 represents full trust while p = 0 indicates no trust

- Weights in the trust network represent the probability
  - (u,v,0) is represented no link between u and v in the trust network
  - Probabilities usually follow a certain distribution such as beta distribution

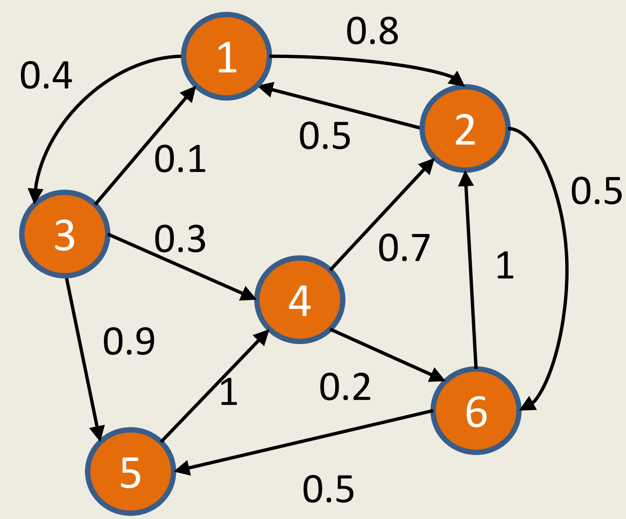

Illustrations of Probabilistic Representations

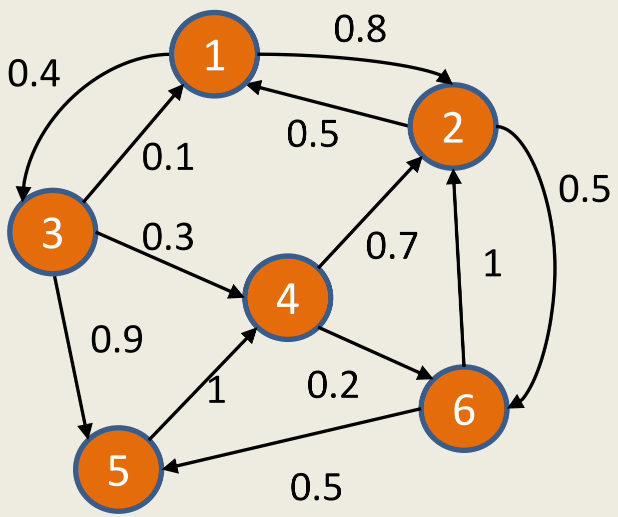

	1	2	3	4	5	6
1	0	0.8	0.4	0	0	0
2	0.5	0	0	0	0	0.5
3	0.1	0	0	0.3	0.9	0
4	0	0.7	0	0	0	0.2
5	0	0	0	1	0	0
6	0	1	0	0	0.5	0

Probabilistic Matrix

Gradual Representations [Victor et al., 2011]

- Trust is often interpreted as a gradual phenomenon in real life
  – Trusting someone ``very much’’, “ more or less”, “little”…
  – Gradual representations become increasing popular

- Gradual representations use continuous values to represent trust
  – The values can be any values so they cannot be explained as probabilities
  – The values directly indicate trust strengths
  – (u,v,t) denotes that the trust value from u to v is t
  – Weights in the trust network denote trust values
  – t = 0 indicates no trust and there is no link in the trust network

Illustrations of Gradual Models

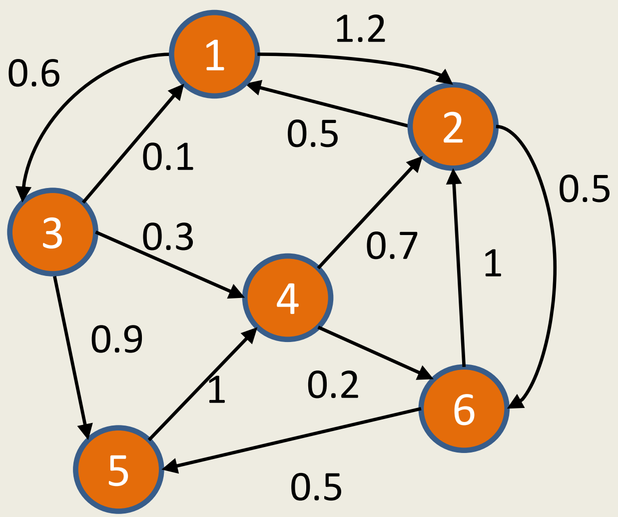

	1	2	3	4	5	6
1	0	1.2	0.6	0	0	0
2	0.5	0	0	0	0	0.5
3	0.1	0	0	0.3	0.9	0
4	0	0.7	0	0	0	0.2
5	0	0	0	1	0	0
6	0	1	0	0	0.5	0

Trust Value Matrix

Trust Representation Classifications

- Trust representations can be classified from different perspectives

- Probabilistic vs. gradual trust representations
– From an interpretation perspective

- <u>Single</u> vs. <u>multi-dimensional</u> trust representations
– From a dimension perspective

Multi-dimensional trust representations

- Trust is a complex concept with multiple dimensions
  – Multi-faceted trust
  – Trust evolution

- We need to extend single trust representations to multi-dimensional trust representations

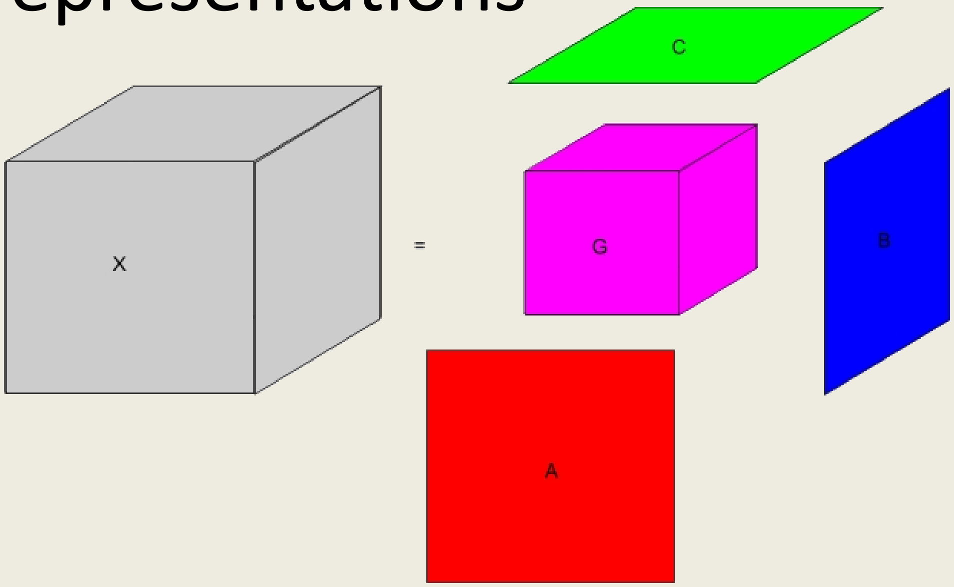

Multi-Faceted Trust [Tang et al., 2012]

- Trust is context dependent

– Trusting someone on one topic does not necessarily mean he will be trusted on others

- An illustrative example using the Epinions dataset

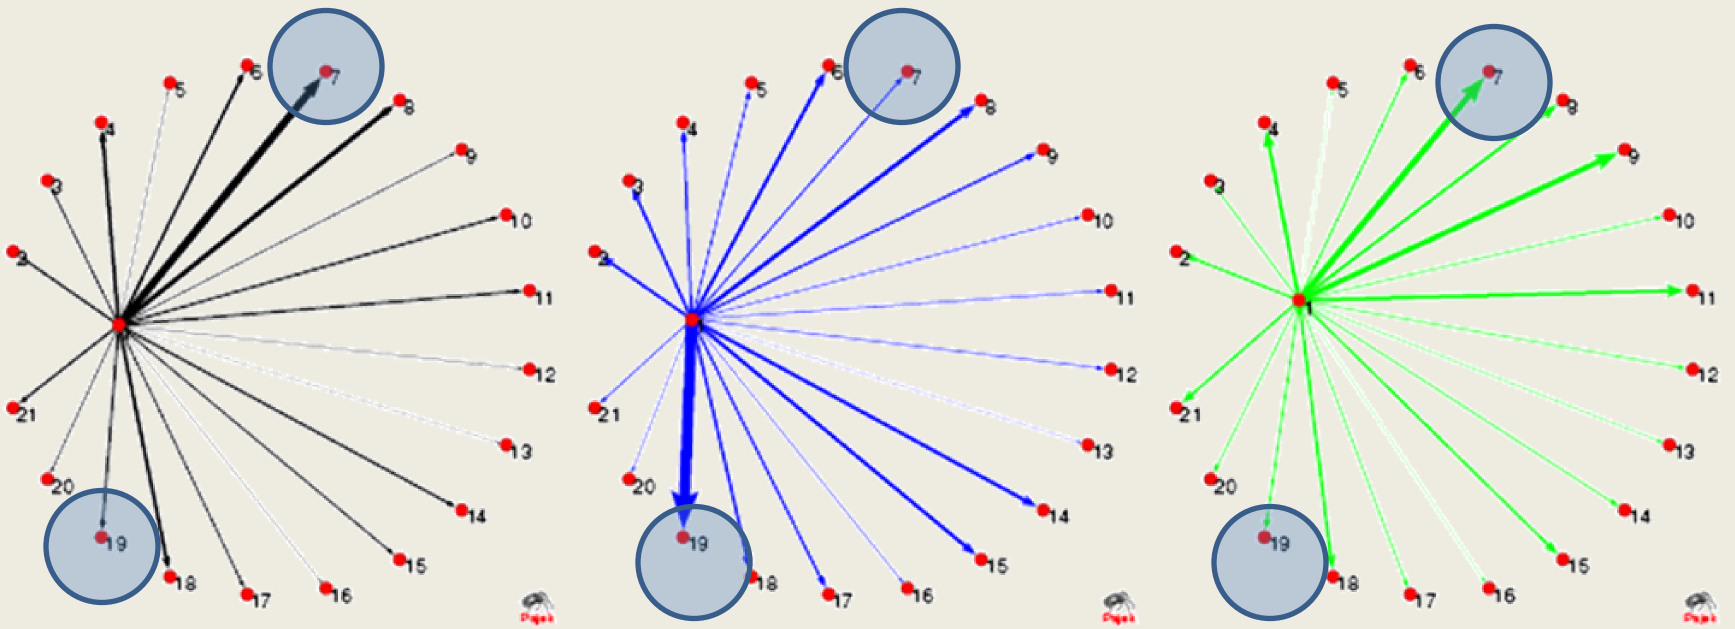

(a) Single Trust (b) Trust in Home & Garden (c) Trust in Restaurants

Multi-faceted Trust Representation

A matrix representation can be extended to a tensor representation for multi-faceted trust

– <u, v, f, p>
– For probabilistic models, u trusts v with probability p in the facet f
– For gradual models, the trust value between u and v in the facet f is p

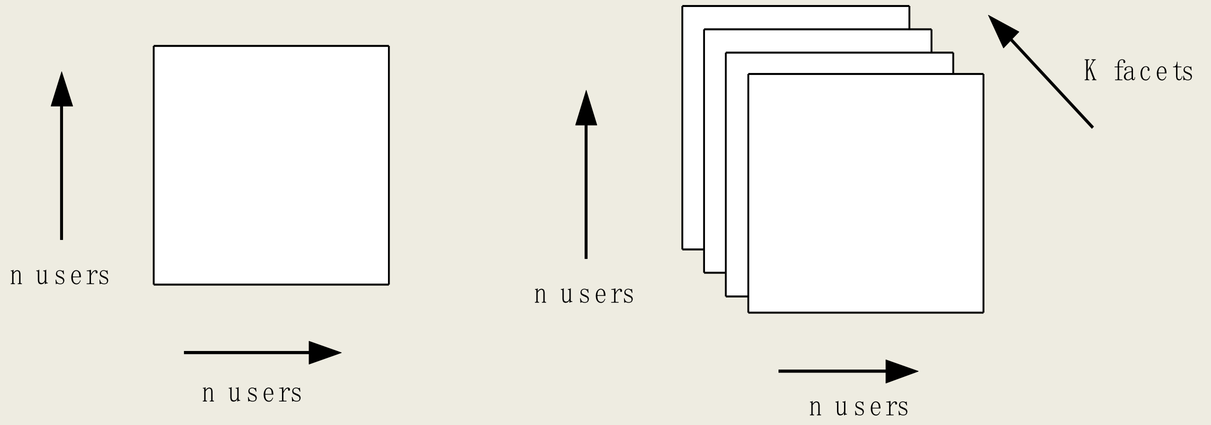

Findings from Multi-faceted Trust Representation

- Heterogeneous transitive trust

– User i trusts user j in f_1 and user j trusts user k in f_2

– 22.3% transitive trust relations are heterogeneous

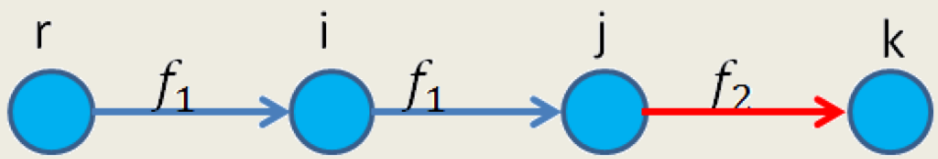

- Heterogeneous reciprocal trust

– User i trusts user k in f_1 and user k trusts i in f_2

– 23.5% of reciprocal trust relations are heterogeneous

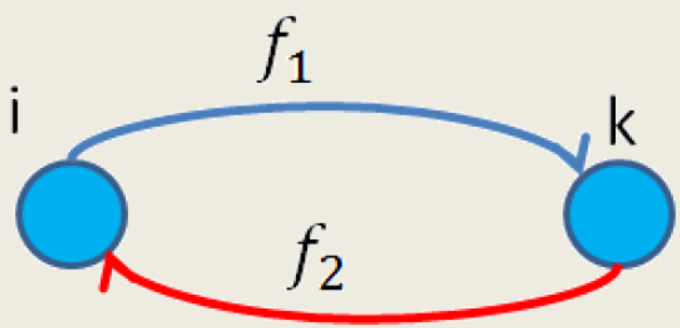

Trust Evolution [Tang et al. 2012]

- Social sciences suggest that trust evolves as humans interact
- An example from an online rating system Epinions

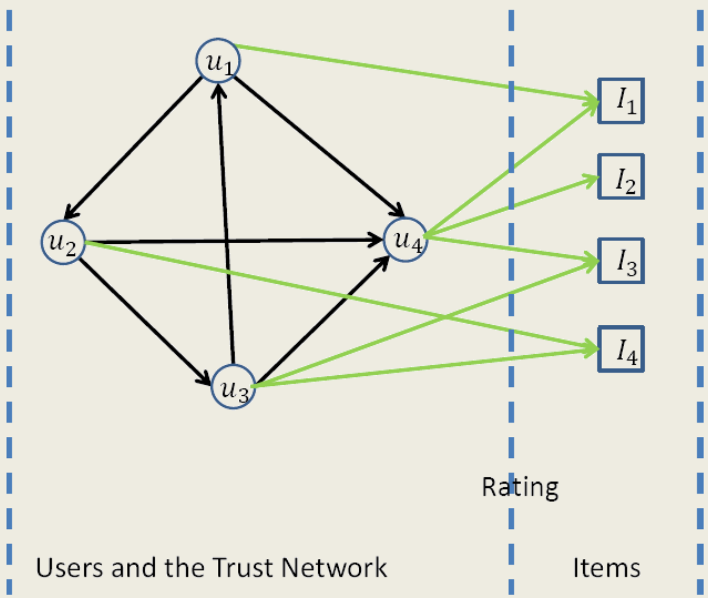

(a) Rating System at T_1

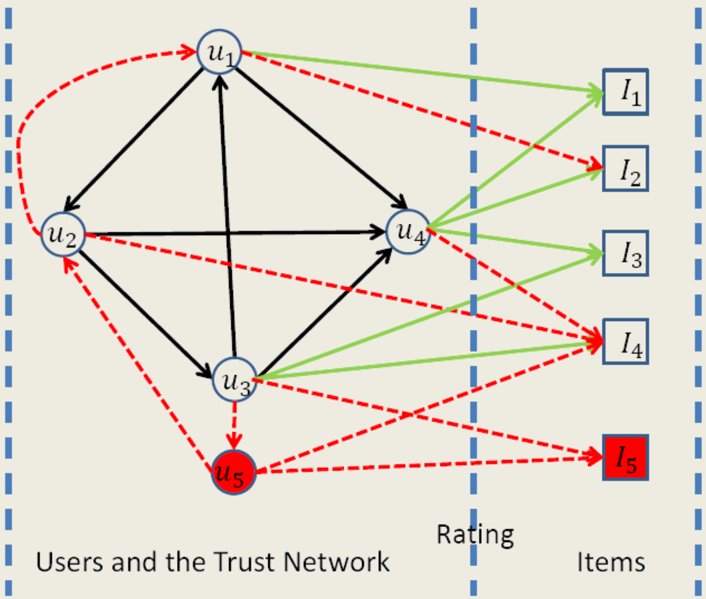

(b) Rating System at T_2

Trust Evolution Representations

- 3-order tensor representations for trust evolution

– <u, v, T, p>

– For probabilistic models, u trusts v with probability p at time T

– For gradual models, the trust value between u and v at time T is p

- 4-order tensor representations for multi-faceted and evolved trust

– <u, v, f, T, p>

– For probabilistic models, u trusts v with probability p in the facet f at time T

– For gradual models, the trust value between u and v in the facet f at time T is p

Findings of Trust Evolution

- Trust strength in an open triad evolves faster than that in a close triad

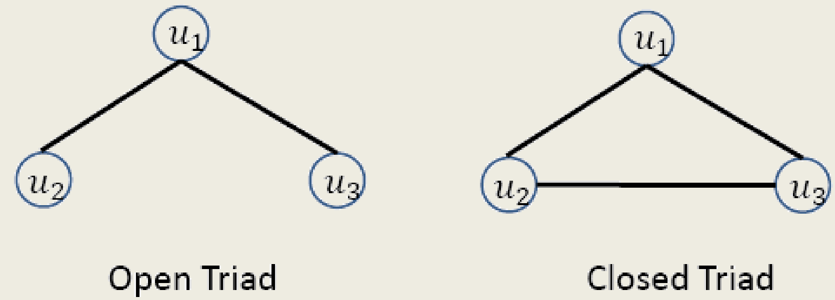

- Trust evolves over time with the changes of user preferences

- Trust evolves differently in different facets

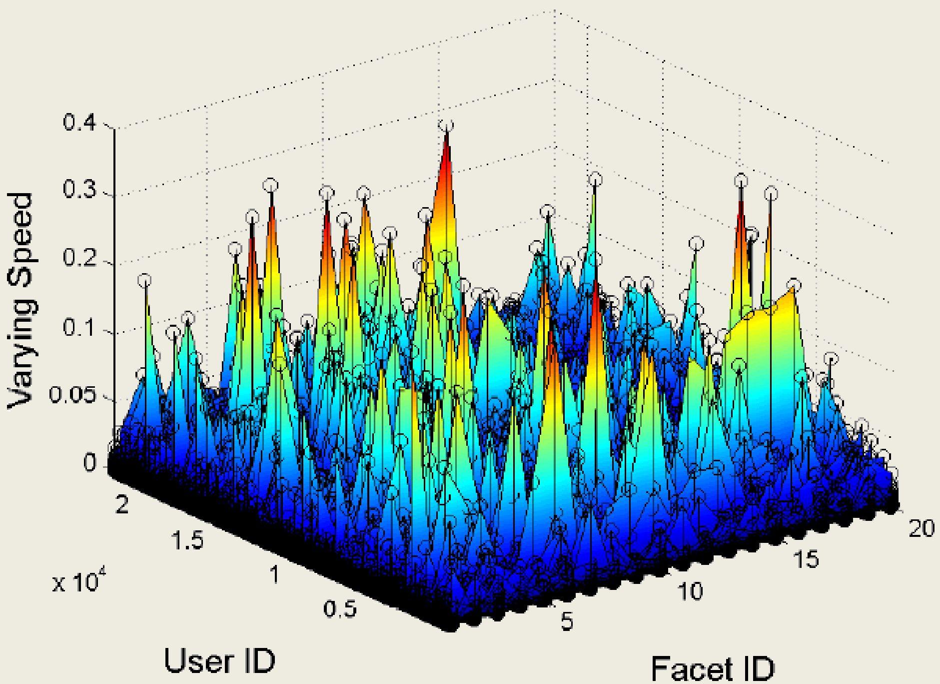

Trust Representation Classifications

- Trust representations can be classified from different perspectives

- Probabilistic vs. gradual trust representations
  – From a interpretation perspective

- Single vs. multi-dimensional trust representations
  – From a dimension perspective

- <u>Trust</u> vs. <u>trust and distrust</u> representations
  – From a network perspective

Measuring Trust

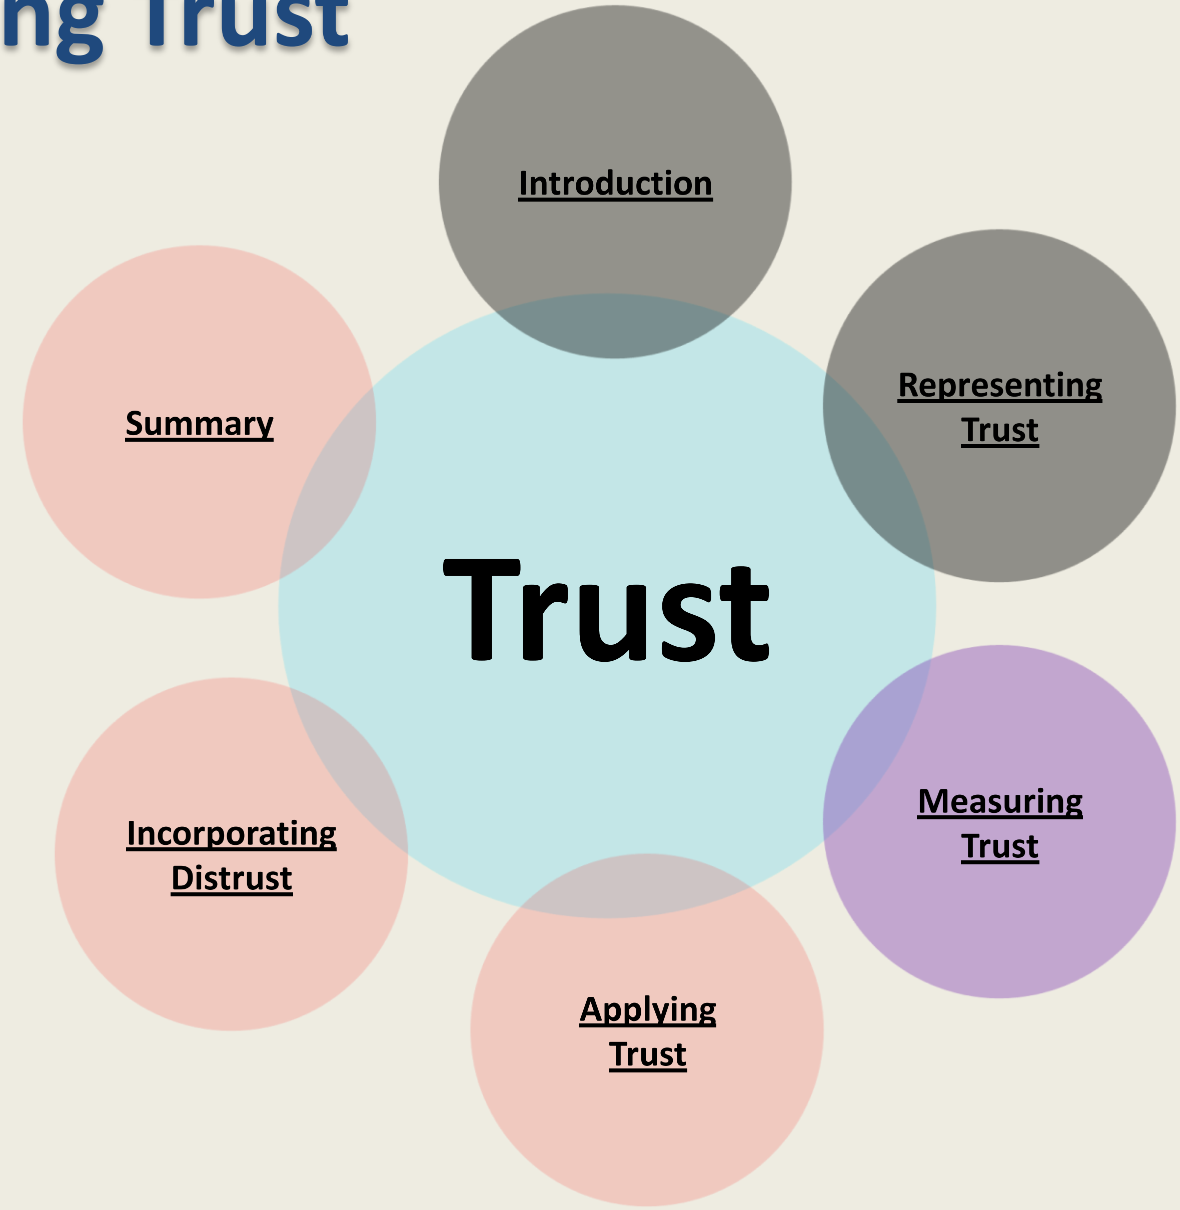

Definitions

- A trust metric measures how much a certain user can be trusted by the other users for the community
  – Measuring, inferring and predicting trust

- Propagation is assumed in most trust metrics
  – We trust our trustees more than a stranger
  – A trustee of our trustee is possibly more trustworthy than a random stranger

Classifications [Ziegler and Lausen, 2005]

- Trust metrics can be classified from different perspectives

- Global and local trust metrics
  – From a personalization perspective

- Supervised vs unsupervised trust metrics
  – From a methodology perspective

- Binary or continuous trust metrics
  – From a network perspective

Global and Local Metrics

- Global metrics predict the same trust of a given user for all users
  - Each user with a global trust value
  - E.g., Reputation systems

- Local metrics provide a personalized trust score that depends on the point of view of the evaluating user
  - Each pair of users with a trust score
  - Personalized trust
  - Users may have completely different opinions about the same user

PageRank [Page et al., 1999]

- PageRank is a global metric

- Pagerank is from trustees' perspective

- The trustworthiness of a trustee is aggregated from her trustors

–Trustors' trustworthiness

–Trust values

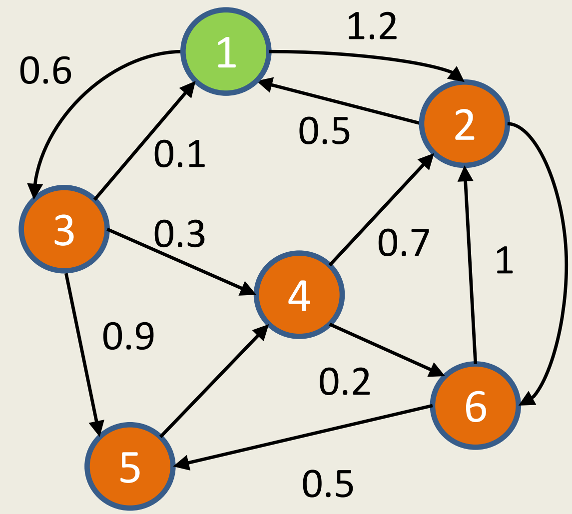

User 1 will give more her trustworthiness to user 2 compared to user 3

$$x_i = \alpha \sum_j \boxed{x_j} \frac{\boxed{T_{ij}}}{\sum_j T_{ij}} + (1-\alpha)\beta$$

EigenTrust [Kamvar et al., 2003]

- Asking your trustees and aggregating trust for trustees

$$t_{ik} = \sum_j c_{ij} c_{jk}$$

What your trustee trust k

Ask your trusted users j

Weight each trustee's opinions by how much you trust him

- Iterative Aggregating
  - Ask your trustees: $\mathbf{t} = C^t c_i$
  - Ask trustees' trustees: $\mathbf{t} = (C^T)^2 c_i$
  - Keep asking until $\mathbf{t}$ converges: $\mathbf{t} = (C^T)^n c_i$

$$C = [c_{ij}]$$

$$c_{i1} + c_{i2} + \ldots + c_{in} = 1$$

- When n is large, **t** converges to the same vector for every user
  - **t** is the eigenvector of C

$$t^{(k+1)} = C^T t^{(k)}$$

An Illustration Example of EigenTrust

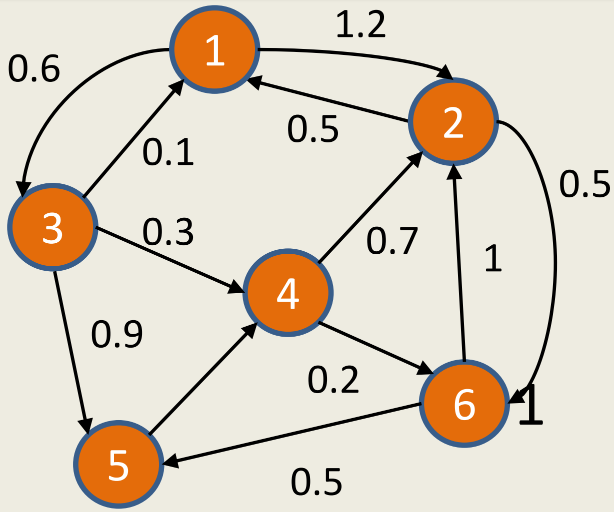

Normalization

$$C = \begin{matrix} 0 & 0.67 & 0.33 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0 & 0 & 0.5 \\ 0.08 & 0 & 0 & 0.23 & 0.69 & 0 \\ 0 & 0.78 & 0 & 0 & 0 & 0.22 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0.67 & 0 & 0 & 0.33 & 0 \end{matrix}$$

$$t^0 = \begin{matrix} 0.1667 \\ 0.1667 \\ 0.1667 \\ 0.1667 \\ 0.1667 \\ 0.1667 \end{matrix} \qquad t^1 = C^T t^0 = \begin{matrix} 0.0967 \\ 0.3534 \\ 0.0550 \\ 0.2050 \\ 0.1700 \\ 0.1200 \end{matrix} \qquad t^2 = C^T t^1 = \begin{matrix} 0.1811 \\ 0.3051 \\ 0.0319 \\ 0.1827 \\ 0.0776 \\ 0.2218 \end{matrix} \quad \cdots\cdots\cdots \quad \begin{matrix} 0.1764 \\ 0.3434 \\ 0.0582 \\ 0.1188 \\ 0.1055 \\ 0.1979 \end{matrix}$$

TidalTrust [Golbeck, 2006]

- TidalTrust is guided by two observations
  – Shorter propagation paths produce more accurate trust estimates
  – Paths with higher trust values create better results

- For a pair of users *i* and *s* who are not directly connected, a trust value is aggregated from the trust value from *i*'s direct neighbors to *s*, weighted by the direct trust values from *i* to her direct neighbors

Neighbors

Higher trust values

$$t_{is} = \frac{\displaystyle\sum_{j \in adj(j) \mid t_{ij} \geq max} t_{ij} t_{js}}{\displaystyle\sum_{j \in adj(j) \mid t_{ij} \geq max} t_{ij}}$$

The length of a path is determined by the number of edges the source must traverse before reaching the sink

An Illustration Example of TidalTrust

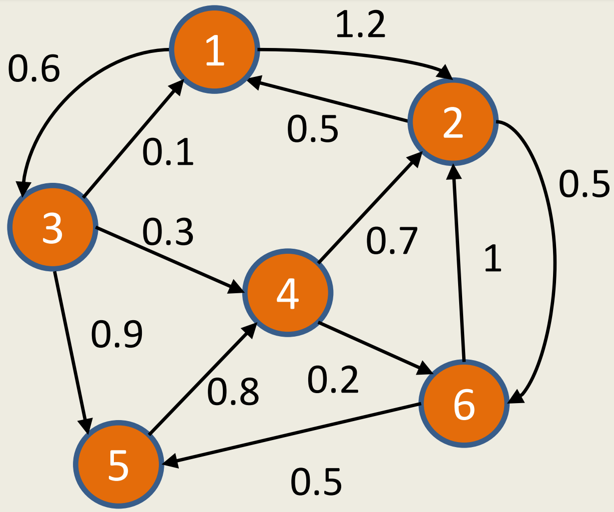

t_{36} = ?

max = 0.15

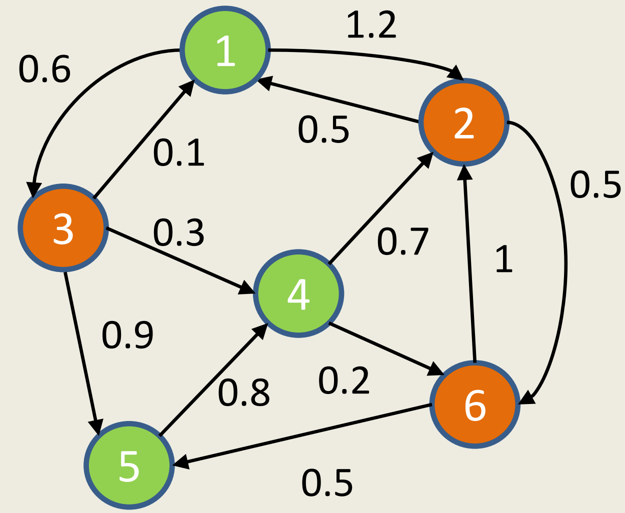

t_{31} < max

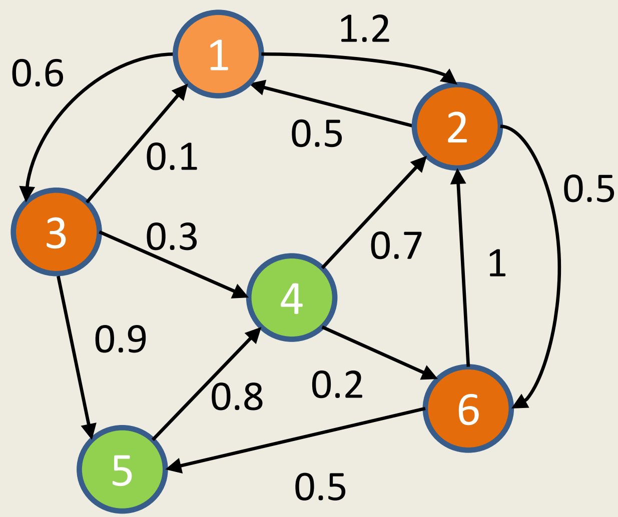

$$t_{36} = \frac{t_{34} * t_{46} + t_{35} * t_{56}}{t_{34} + t_{35}} = \frac{0.3*0.2 + 0.9*t_{56}}{1.2}$$

$$t_{56} = \frac{t_{54} * t_{46}}{t_{54}} = 0.2$$

$$t_{36} = 0.2$$

MoleTrust [Massa and Avesani, 2005]

- Cycles in a trust network are removed
  – The removal reduces the number of trust propagations
  – The trust network is transformed into a directed acyclic graph

- Trust values are calculated based on the directed acyclic graph by performing a simple graph random walk
  – The trust of the users at 1-hop away is computed
  – The trust of the users at 2-hop away, etc.

Comparison between Global and Local Metrics

- Global metrics
  – For n users, we only need to compute n values
  – Efficient to compute and maintain
  – Providing a global view about a user's reputation

- Local metrics
  – For n users, we have n*n pairs of users
  – Providing personalized trust values
  – Applying to controversial users

Supervised Metrics

- Supervised metrics consider trust prediction as a classification problem
- Training data preparation
  – Trust as the positive label
  – Not trust as the negative label
- Feature extraction
  – Extracting a set of features from available sources to represent pairs of users
- Different supervised metrics use different feature sets and classifiers

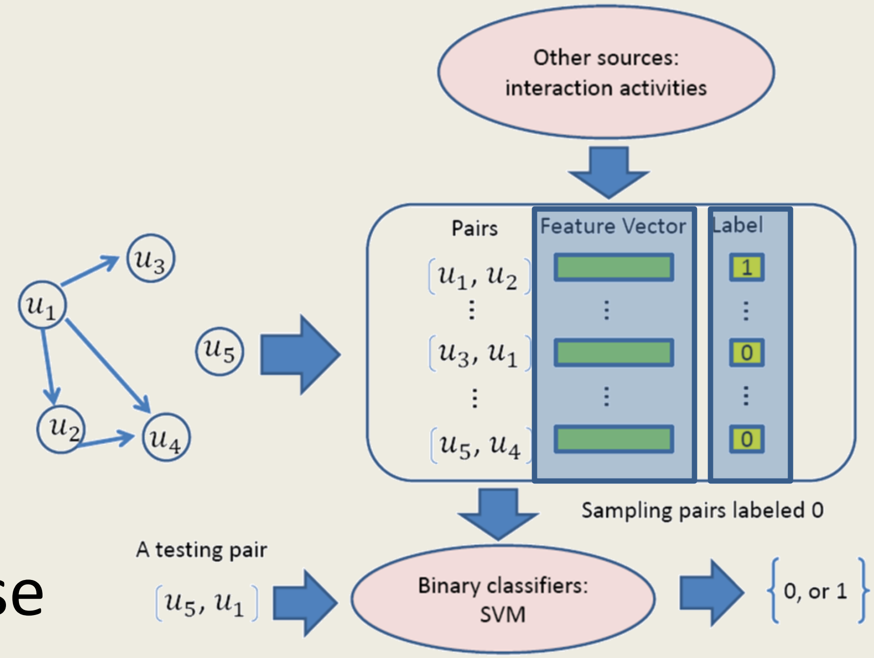

Method by Liu et al. [Liu et al., 2008]

- It assumes there are two reasons a user trusts another user
  – The trustee has good reputation
  – There have been good personal interactions between the two users

- Each reason is captured by a set of features
  – Features based on user factors
  – Features based on interaction factors

- SVM and NB classifiers are trained to predict trust

- Interaction factors have greater impact on trust decisions than user factors
  – Trust is highly relevant to user interactions

Method by Nguyen et al. [Nguyen et al., 2009]

- It is based on trust antecedent framework in management science
  – Ability, benevolence and integrity as key factors that leads to trust on a trustee
  – Trust propensity is a factor that determines how easy a trustor trusts someone
  – Once a trust is formed, the trustor is more willing to take more risk
  – The outcome of risk taking will serve as feedback to modify the perception about trustee's ability, benevolence and integrity

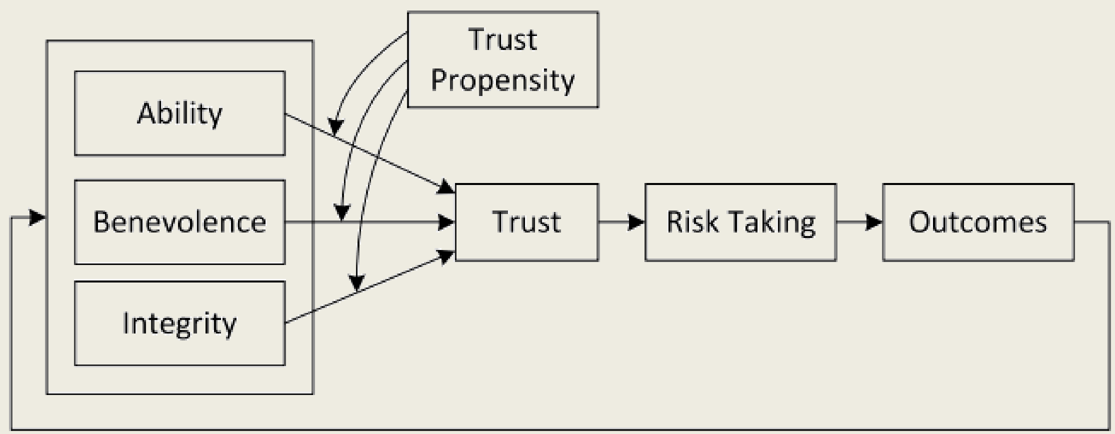

- Each factor is approximated through a set of quantitative features
  – Features for integrity: the number of trust statements the user receives
  – Features for ability: the number of reviews rated by the rater

Unsupervised Metrics

- Unsupervised methods are usually based on the connectivity of users in trust networks
  - Trust Propagation

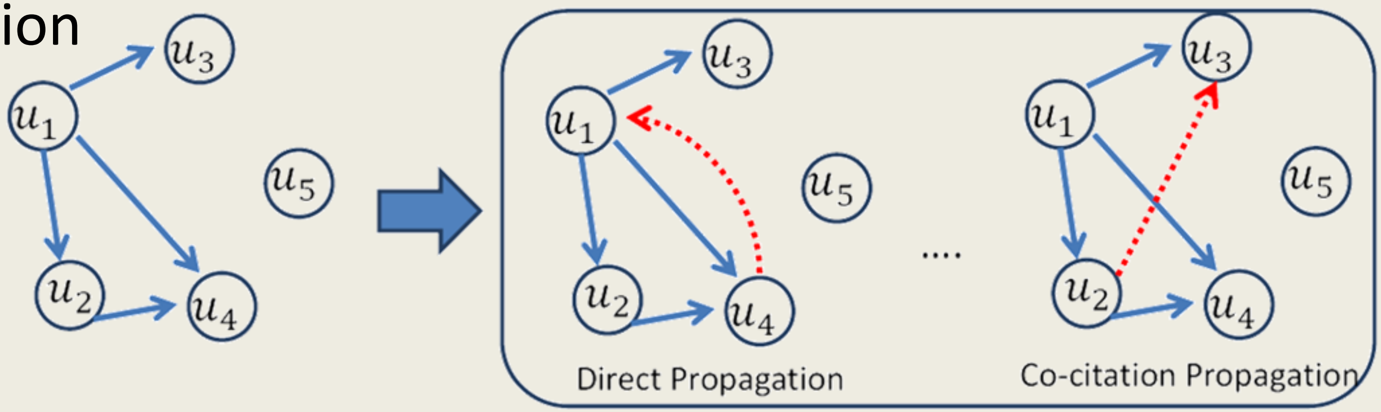

  - Low-rank representations of trust networks

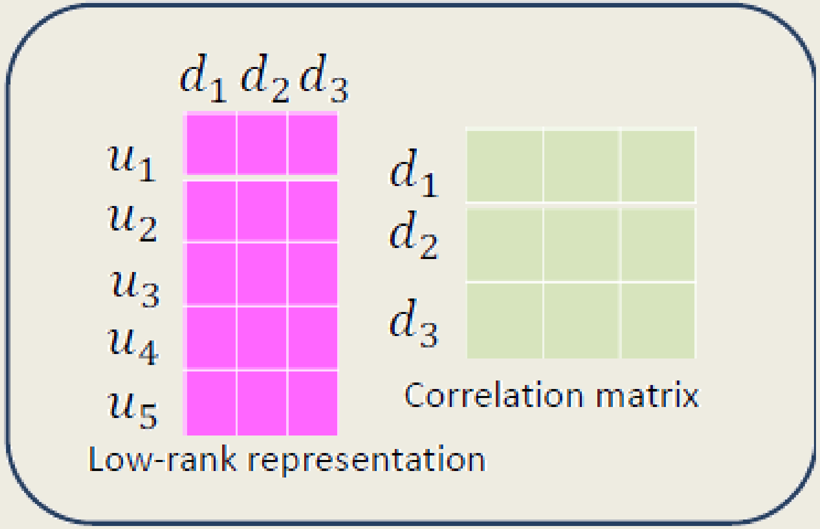

Trust Propagation [Guha et al., 2004]

Four types of atom trust propagations

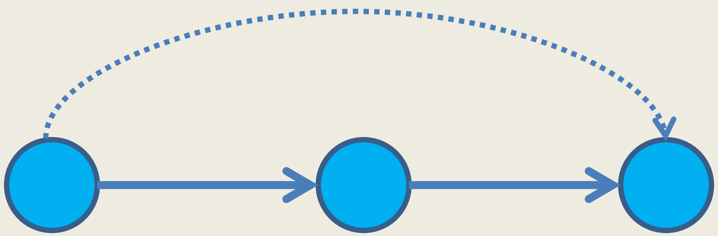

Direct propagation

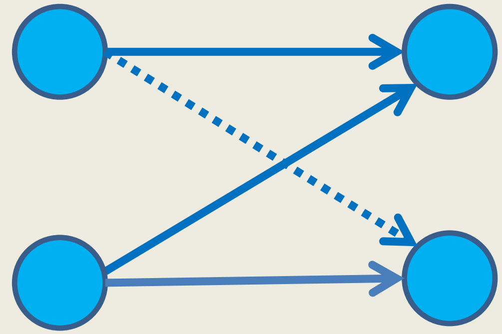

Co-citation

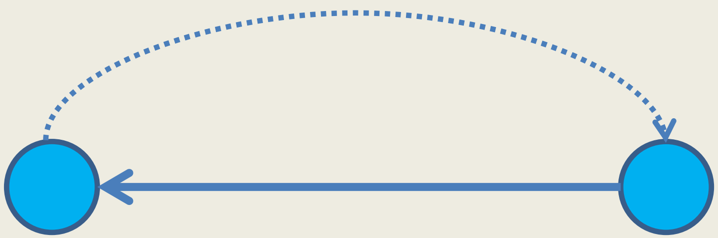

Transpose trust

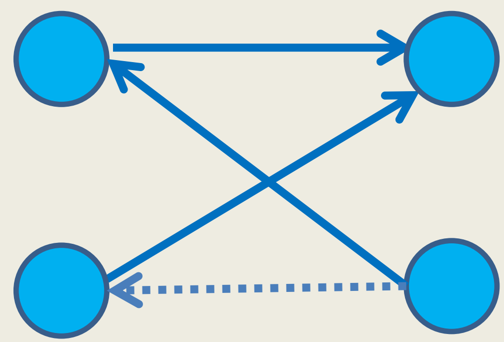

Trust Coupling

Propagation Operations

Atomic Propagation	Operator	Description
Direct propagation	T	If A trusts B, and B trusts C, then A may trust C
Transpose trust	T^T	A's trust of B causes B to develop some level of trust towards A
Co-citation	$T^T T$	If A trusts B and C, D trusting B implies D should trust C
Trust Coupling	TT^T	If A and B trust C, trusting A should imply trusting B

Trust Propagation Aggregation

- A combination of four types of propagations

$$\mathrm{C} = a\mathrm{T} + b\mathrm{T}^{\mathrm{T}}\mathrm{T} + c\mathrm{T}^{\mathrm{T}} + d\,\mathrm{T}\mathrm{T}^{\mathrm{T}}$$

a, b, c, and d are the weights for these four types of propagation, respectively

- Aggregation after K-step propagation

$$\widehat{\mathrm{T}} = \sum_{k=1}^{K} r_k C_k$$

r_k is the aggregation weight for the k-th propagation

An Illustration of Trust Propagation

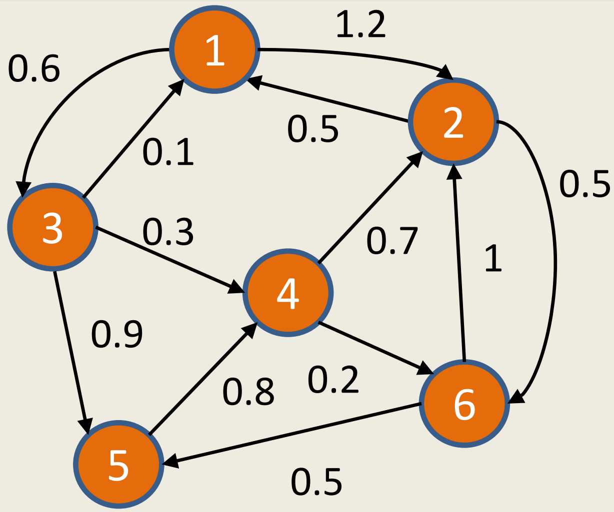

a=b=c=d=0.25

$r_1=r_2=0.5$

$$T = \begin{matrix} 0 & 1.2 & 0.6 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0 & 0 & 0.5 \\ 0.1 & 0 & 0 & 0.3 & 0.9 & 0 \\ 0 & 0.7 & 0 & 0 & 0 & 0.2 \\ 0 & 0 & 0 & 0.8 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0.5 & 0 \end{matrix}$$

$$C_1 = \begin{matrix} 0.51 & 0.42 & 0.17 & 0.22 & 0.02 & 0.36 \\ 0.42 & 0.86 & 0.19 & 0.20 & 0.12 & 0.41 \\ 0.17 & 0.19 & 0.32 & 0.07 & 0.28 & 0.11 \\ 0.22 & 0.20 & 0.07 & 0.31 & 0.27 & 0.22 \\ 0.02 & 0.12 & 0.28 & 0.27 & 0.42 & 0.12 \\ 0.36 & 0.41 & 0.11 & 0.22 & 0.12 & 0.38 \end{matrix}$$

$$C_2 = \begin{matrix} 0.58 & 0.62 & 0.23 & 0.29 & 0.12 & 0.47 \\ 0.62 & 1.02 & 0.29 & 0.33 & 0.28 & 0.58 \\ 0.23 & 0.29 & 0.29 & 0.15 & 0.28 & 0.19 \\ 0.29 & 0.33 & 0.15 & 0.31 & 0.27 & 0.29 \\ 0.12 & 0.23 & 0.28 & 0.27 & 0.39 & 0.19 \\ 0.47 & 0.58 & 0.19 & 0.29 & 0.19 & 0.45 \end{matrix}$$

An Illustration Example of Trust Propagation

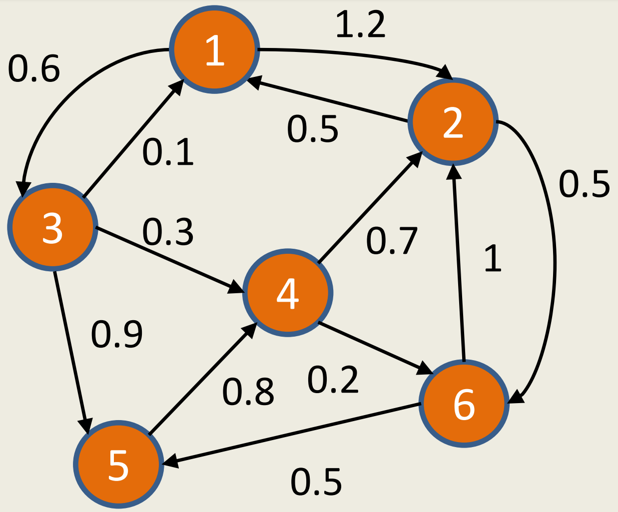

$$T = \begin{matrix} 0 & 1.2 & 0.6 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0 & 0 & 0.5 \\ 0.1 & 0 & 0 & 0.3 & 0.9 & 0 \\ 0 & 0.7 & 0 & 0 & 0 & 0.2 \\ 0 & 0 & 0 & 0.8 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0.5 & 0 \end{matrix}$$

a=b=c=d=0.25

$r_1=r_2=0.5$

$$\hat{T} = 0.5C_1 + 0.5C_2 = \begin{matrix} 0.28 & 0.55 & 0.25 & 0.11 & 0.01 & 0.18 \\ 0.46 & 0.49 & 0.10 & 0.10 & 0.06 & 0.45 \\ 0.13 & 0.10 & 0.16 & 0.15 & 0.49 & 0.05 \\ 0.11 & 0.49 & 0.04 & 0.16 & 0.13 & 0.22 \\ 0.01 & 0.06 & 0.14 & 0.63 & 0.21 & 0.06 \\ 0.18 & 0.54 & 0.05 & 0.11 & 0.23 & 0.19 \end{matrix}$$

Low-rank Matrix Factorization [Tang-etal.,2013]

- Trust is multi-faceted and is correlated to user preferences
  - Assume **U**$_i$ is the k-dimensional preference vector of the user i
  - **V** is a K x K correlation matrix
  - A trust relation between user i and user j can be modeled as the interactions between their preferences by **V** as

$$T_{ij} = U_i V U_j$$

- Low-rank matrix factorization model can capture the major properties of trust,
  - Multi-faceted
  - Correlation with user preferences
  - Transitivity and asymmetry

An Illustration of Low-rank Matrix Factorization

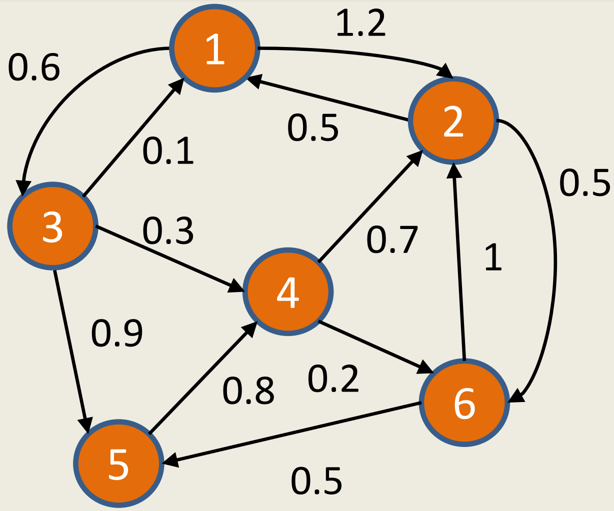

$$T = \begin{matrix} 0 & 1.2 & 0.6 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0 & 0 & 0.5 \\ 0.1 & 0 & 0 & 0.3 & 0.9 & 0 \\ 0 & 0.7 & 0 & 0 & 0 & 0.2 \\ 0 & 0 & 0 & 0.8 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0.5 & 0 \end{matrix}$$

K=1

$$U = \begin{matrix} 0.4969 \\ 0.6635 \\ 0.2616 \\ 0.3229 \\ 0.2730 \\ 0.4534 \end{matrix} \qquad V = 1.1511 \qquad \hat{T} = UVU^T = \begin{matrix} 0.28 & 0.38 & 0.15 & 0.18 & 0.16 & 0.26 \\ 0.38 & 0.51 & 0.20 & 0.25 & 0.21 & 0.34 \\ 0.15 & 0.20 & 0.08 & 0.10 & 0.08 & 0.13 \\ 0.18 & 0.24 & 0.10 & 0.12 & 0.10 & 0.17 \\ 0.16 & 0.21 & 0.08 & 0.10 & 0.08 & 0.14 \\ 0.26 & 0.34 & 0.14 & 0.17 & 0.14 & 0.23 \end{matrix}$$

Comparison of Supervised and Unsupervised Metrics

- Supervised metrics
  - The number of pairs without trust is much larger than those with
  - The classification problem is highly imbalanced
  - They need extra sources to extract features
  - They usually outperform unsupervised metrics
  - They can be applied to users with few trust relations

- Unsupervised metrics
  - They only depend on the structure of trust networks
  - They may fail for users with few trust relations
  - They can be applied to both binary and weighted trust networks

Binary and Continuous Metrics

- Binary metrics are used to predict whether users are trusted or not trusted
  – {1:trust, 0: not trust}
  – For supervised metrics, trust and not trust are positive and negative labels to predict by learnt classifiers
  – For unsupervised metrics, trust values in a certain region are treated as trust, otherwise as not trust

- Continuous metrics are to infer trust values for pairs of users
  – Nonnegative real number
  – Continuous metrics are usually unsupervised methods

Evaluations of Trust Metrics

- There are many evaluation metrics proposed to assess the performance of trust metrics

- Ranking-based evaluation

- RMSE evaluation

- Leave-one-out cross-validation evaluation

- F-measure evaluation

Ranking-based Evaluation

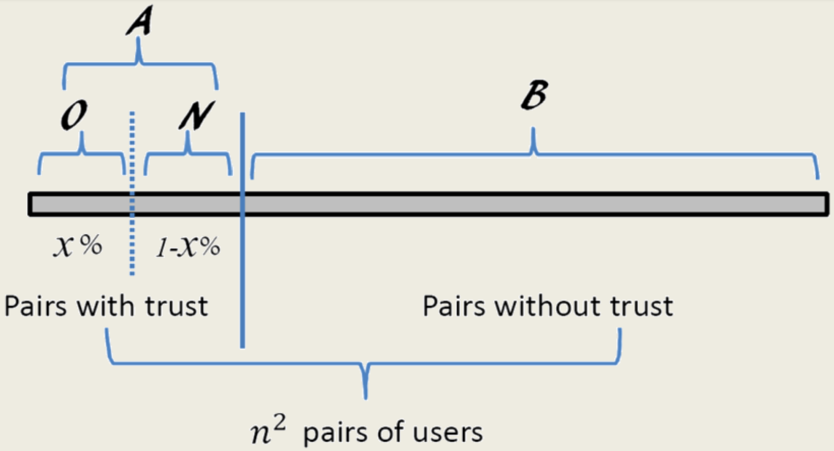

- x% as old trust relations
- 1 - x% as new trust relations
- Ranking pairs of users in *N* and *B*
- Choosing top-|N| ranked pairs as C as predicted trust relations
- Calculating the prediction quality as

$$PA = \frac{|N \cap C|}{|N|}$$

N

1	2
2	5
3	5
4	6
1	5

C

1	2
2	3
2	5
4	5
1	6

$$PA = \frac{1}{5} = 0.2$$

The value of PA is usually small and to demonstrate the significance of performance, randomly guessing predictor is usually used as a baseline method

RMSE Evaluation

- x% as old trust relations O
- 1 - x% as new trust relations N
- Computing trust values for pairs of users in *N*
- Calculating RMSE as

$$RMSE = \sqrt{\frac{\sum_{<u_i, u_j> \in N} (\hat{T}_{ij} - T_{ij})^2}{|N|}}$$

T_{ij}	$\widehat{T}_{ij}$
0.2	**0.2**
0.4	**0.5**
0.5	**0.5**
0.7	**0.6**
0.4	**0.5**

$$RMSE = \sqrt{\frac{(0.2-0.2)^2 + (0.5-0.4)^2 + (0.5-0.5)^2 + (0.6-0.7)^2 + (0.5-0.4)^2}{5}} = 0.0775$$

Leave-one-out Cross-validation

- Step 1: Given a full network, randomly hide one of the trust relations
- Step 2: Predict the existence of the hidden trust relation which has been suppressed
- Step3: Repeat **Step 1** and **Step 2** *n* times
- Assume that the predictor infers these hidden relations correctly *m* of *n* times, then the accuracy is *m/n*

F-measure Evaluation

- Trust and not trust are treated as the positive and negative labels, respectively
- Precision, recall and F-measure are defined as

$$\text{precision} = \frac{tp}{tp + fp}$$

$$\text{recall} = \frac{tp}{tp + fn}$$

$$\text{F1} = 2.\frac{\text{precision}.\text{recall}}{\text{precision} + \text{recall}}$$

	p' (Predicted)	n' (Predicted)
p (Actual)	True Positive	False Negative
n (Actual)	False Positive	True Negative

Evaluation Metrics for Measuring Trust

- Ranking-based evaluation
  - Binary trust metrics
  - Unsupervised metrics
- RMSE evaluation
  - Continuous metrics
- Leave-one-out cross-validation evaluation
  - Supervised metrics
- F-measure evaluation
  - Supervised metrics

Applying Trust

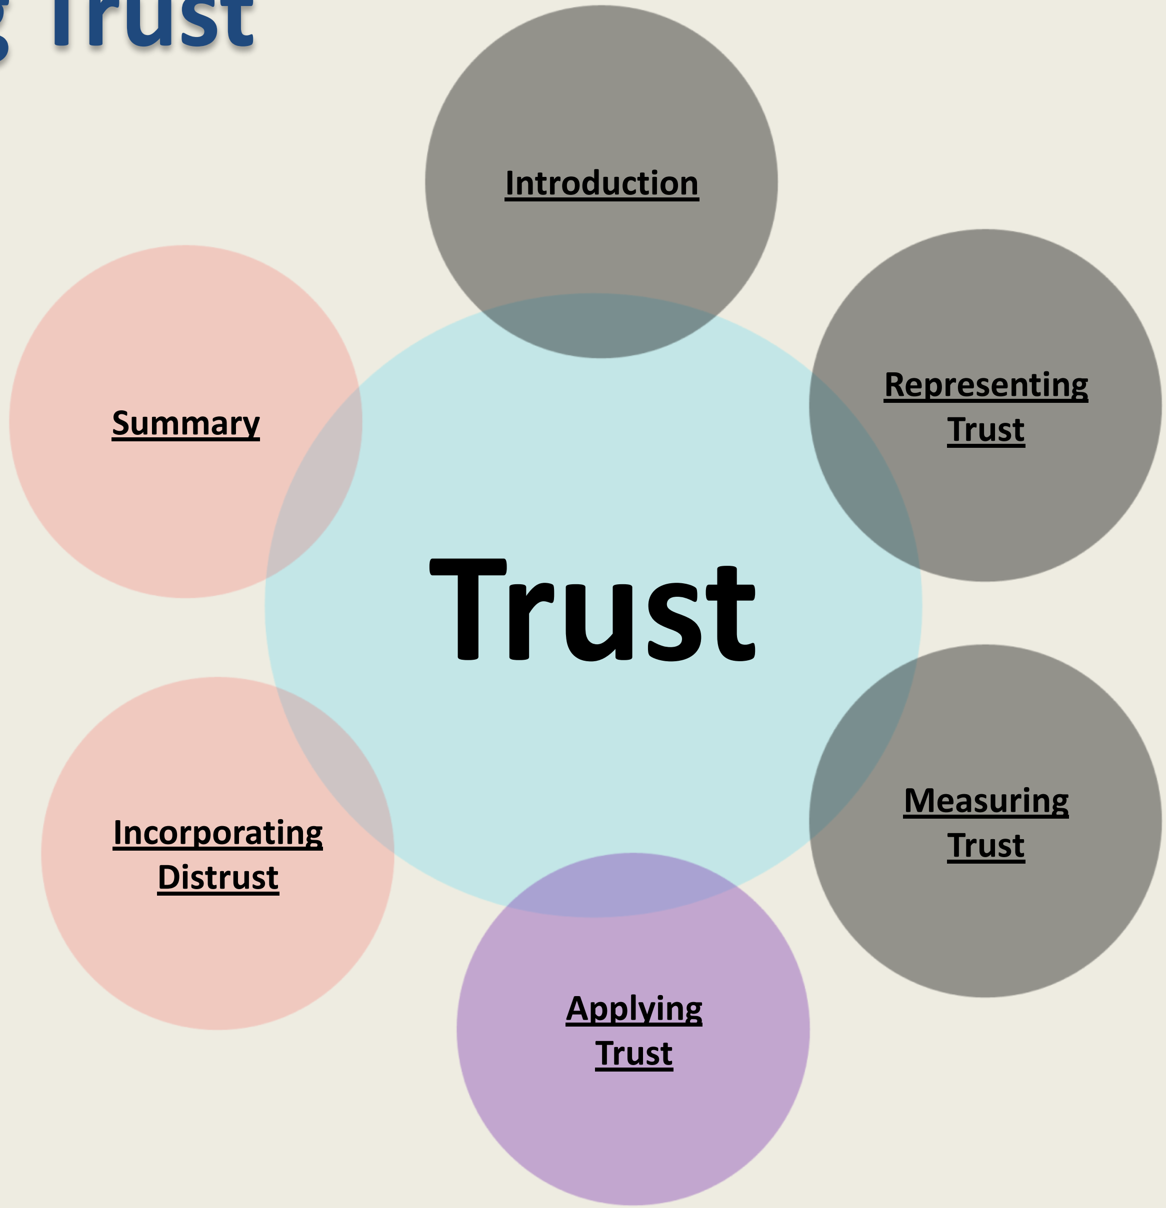

Trust-aware Recommender Systems

- In this physical world, people seek recommendations from their trusted friends

- A user's preference is more likely to be similar to those of her trust network than to those of randomly chosen users
  - Trust information may provide preference context of a user
  - Homophily and influence

- Trust-aware recommender systems augment traditional recommender systems with trust information
  - User-item rating matrix **R**
  - User-user trust matrix **T**

Traditional Recommender Systems

- Content-based recommender systems

– Recommend items similar to the ones that the user has preferred in the past

- Collaborative filtering (CF) -based recommender systems

– Using the user's past behavior to uncover user preferences and recommend items that match their preferences

– Only depending on users' past behaviors

– Memory-based CF and Model-based CF

Memory-based Collaborative Filtering

- It uses either the whole user-item matrix or a sample to generate a prediction
  – Needing memory to store the user-item rating matrix **R**
- User-oriented collaborative filtering
  – Calculating user-user similarity
  – Aggregating ratings from similar users
- Item-oriented collaborative filtering
  – Computing item-item similarity
  – Aggregating ratings from similar items

An Illustration of User-oriented Collaborative Filtering

	A	B	C	D	E
1	5	3	4	?	?
2	?	3	4	4	?
3	1	?	2	2	5

2 and 3 are similar users to 1

A, B, C, D, and E are items

R(1,D) = ?

- Cosine similarity calculation

$$S(1,2) = \frac{3*3+4*4}{\sqrt{3*3+4*4}\sqrt{3*3+4*4}} = 1$$

$$S(1,3) = \frac{4*1+4*2}{\sqrt{1*1+2*2}\sqrt{4*4+4*4}} = 0.9487$$

- Aggregating ratings

$$R(1,D) = \frac{4*1+2*0.9487}{1+0.9487} = 3.03$$

Model-based Collaborative Filtering

- It assumes a model to generate the ratings and learns the parameters of the model
  - Storing only parameters instead of the rating matrix
  - Using the assumed model with parameters to do prediction

- Matrix factorization methods are very competitive and are widely adopted to build recommender systems
  - U_i is the k-dimensional user preference vector of user i
  - V_j is the k-dimensional item characteristic vector for item j
  - A rating from user i to item j is modeled as

$$R_{ij} = U_i V_j^T$$

An Illustration of Matrix Factorization based CF

	A	B	C	D	E
1	5	3	4	?	?
2	?	3	4	4	?
3	1	?	2	2	5

1,2 and 3 are users
A, B, C, D, and E are items
R(1,D) = ?
The latent dimension k =1

- Learning U and V, and reconstructing the rating matrix

$$U = \begin{matrix} 2.6308 \\ 2.4109 \\ 1.4706 \end{matrix} \quad V = \begin{matrix} 1.6252 \\ 1.2182 \\ 1.5740 \\ 1.5990 \\ 2.5716 \end{matrix} \quad \hat{R} = UV^T = \begin{matrix} 4.2756 & 3.2047 & 4.1408 & \boxed{4.2066} & 6.7654 \\ 3.9182 & 2.9368 & 3.7946 & 3.8550 & 6.1999 \\ 2.3901 & 1.7915 & 2.3147 & 2.3515 & 3.7819 \end{matrix}$$

- R(1,D) is predicted as 4.2066

Challenges of Traditional Recommender Systems

- Data sparsity problem

– Social media data is big data but the available data for most individuals is very limited

– The user-item rating matrix is extremely sparse with less than 1% observed ratings

- Cold-start problem

– The number of ratings for users follows a power-law distribution

– There are many users with no or very few ratings

– Existing recommender systems may fail to make recommendations for cold-start users

Opportunities from Trust Information

- Trust provides an additional source for recommendation improvement
  – Overlap between one's similar users and trusted users is low (less than 10%)
  – Mitigating data sparsity problem

- Since a user has her trust network, we can do recommendation based on her trusted users
  – Users' preferences are similar to their trust networks
  – Reducing significantly the number of cold-start users

Categorization[Tang et al., 2013]

- Most existing trust-aware recommender systems are CF-based methods

Trust–aware CF model = A basic CF model + Trust model

- We can categorize existing trust-aware systems based on their basic CF models
  – Memory-based trust-aware recommender systems
  – Model-based trust-aware recommender systems

Memory-based Trust-aware Recommendation

- It uses memory-based CF methods, especially user-oriented methods, as basic models

- It usually consists of two steps

– Step 1: obtaining the trusted users N_i for a given user i,

– Step 2: aggregating ratings from the trusted users obtained by the first step to predict ratings for user i

- Step 2 is the same as that in traditional memory-based methods, and different methods in this category provide different ways to obtain trusted users in Step 1

TidalTrust vs MoleTrust

- TidalTrust only considers raters at the shortest distance
  - Trusted users N_i is the set of users at the shortest distance
  - Efficient
  - High precision
  - Low recall

- MoleTrust considers raters up to a maximum-depth *d*
  - Trusted users N_i is the set of users within maximum-depth
  - Trade-off between precision and recall

$$\hat{\mathbf{R}}_{ij} = \bar{\mathbf{R}}_i + \frac{\sum_{u_k \in \mathcal{N}_i} \mathbf{S}_{ik}(\mathbf{R}_{kj} - \bar{\mathbf{R}}_k)}{\sum_{u_k \in \mathcal{N}_i} \mathbf{S}_{ik}}$$

An Illustration of MoleTrust for Recommendation

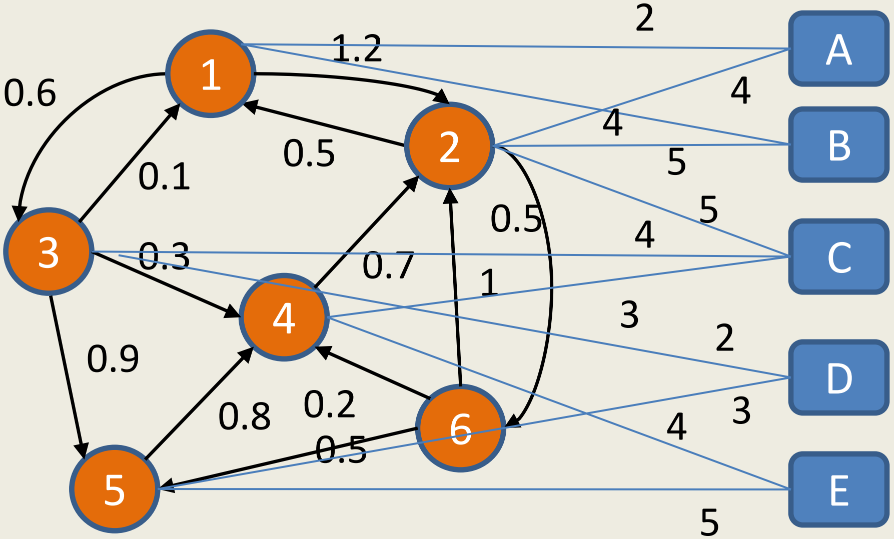

$$T = \begin{matrix} 0 & 1.2 & 0.6 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0 & 0 & 0.5 \\ 0.1 & 0 & 0 & 0.3 & 0.9 & 0 \\ 0 & 0.7 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.8 & 0 & 0 \\ 0 & 1 & 0 & 0.2 & 0.5 & 0 \end{matrix} \qquad R = \begin{matrix} 2 & 4 & ? & ? & ? \\ 4 & 5 & 5 & ? & ? \\ ? & ? & 4 & 2 & ? \\ ? & ? & 3 & ? & 4 \\ ? & ? & ? & 3 & 5 \\ ? & ? & ? & ? & ? \end{matrix}$$

Setting d = 2 and predicting R(6,E) = ?

– The set of raters in the paths with length of 2 is {5, 4, 2, 1}

– Only {5,4} rate item E

$$R(6,E) = \frac{T_{65}R(5,E) + T_{64}R(4,E)}{T_{65} + T_{64}} = \frac{0.5*5 + 0.2*4}{0.5 + 0.2} = 4.714$$

User 6 does not have any ratings and MoleTrust can do recommendations based on her trust network

TrustWalker [Jamali and Ester, 2009]

- Instead of distant neighbors who have rated the target item, it uses near neighbors who have rated similar items
  - Trusted friends on similar items
  - Distant users on the exact target item

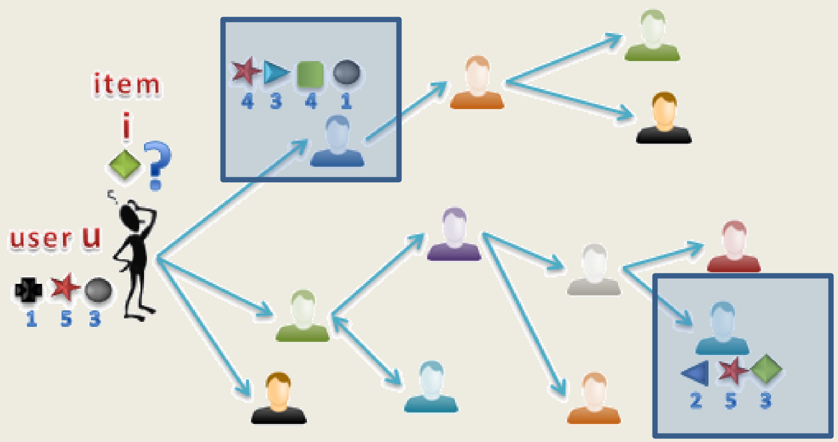

- It combines item-based recommendation and trust-based recommendation via random walk
- Each random walk returns a rating of the target item or a similar item

TrustWalker

- Each random walk starts from a target user u to seek rating score for item i

- In step *k* at node *v*:

– If u has rated i, return R_{vi}

– With the probability $\phi_{v,i,k}$, stop random walk, select a similar item j rated by u and return R_{vj}

– With the probability 1 - $\phi_{v,i,k}$, continue the random walk to a direct neighbor of v

An illustration of TrustWalker

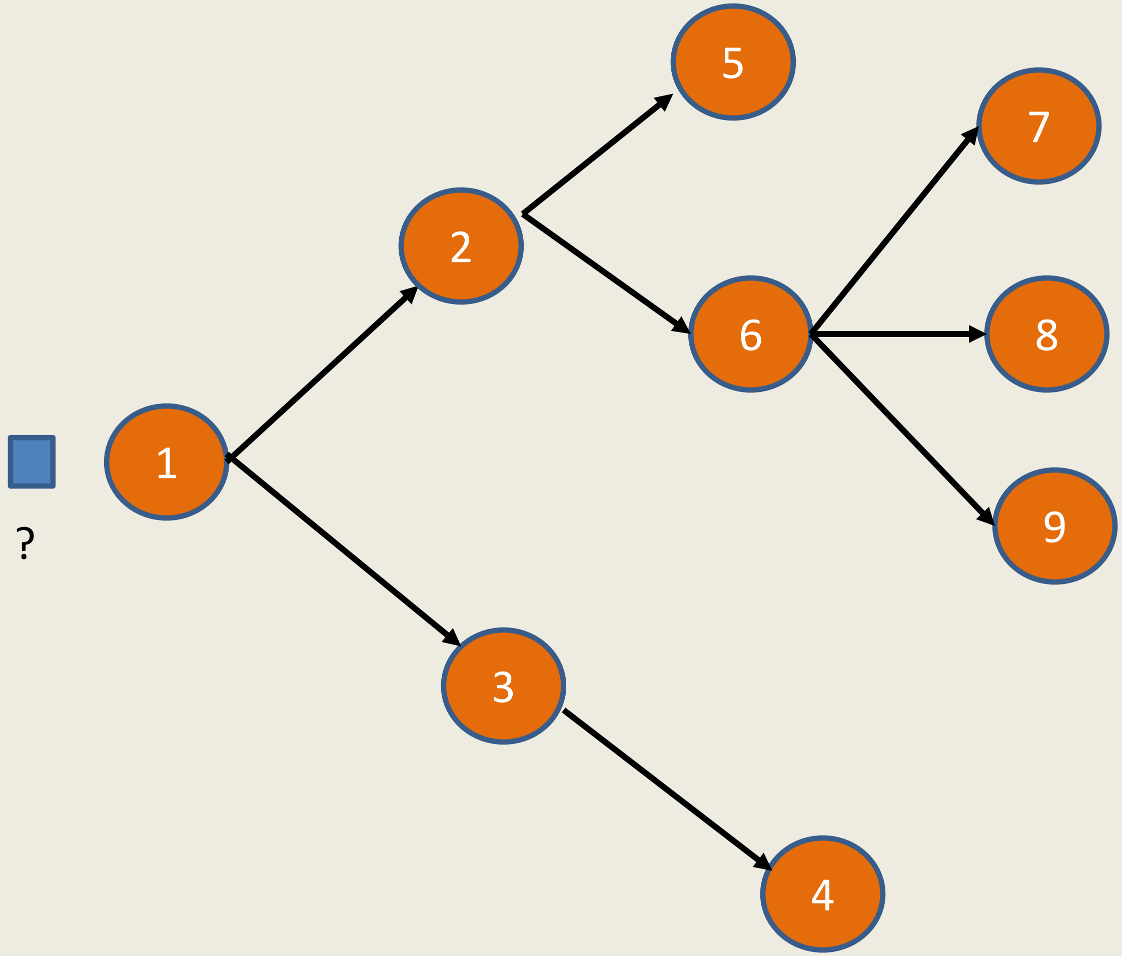

An illustration of TrustWalker

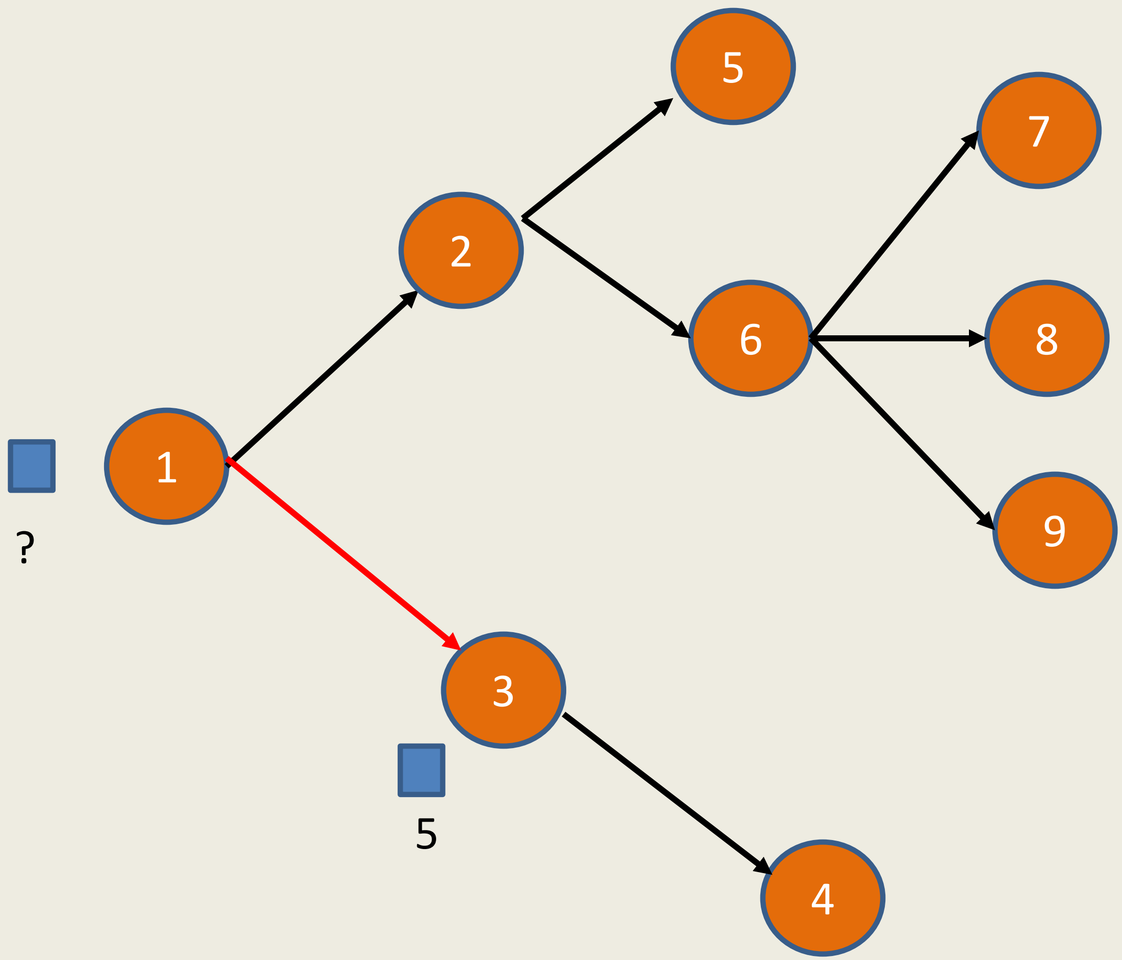

R1=5

An illustration of TrustWalker

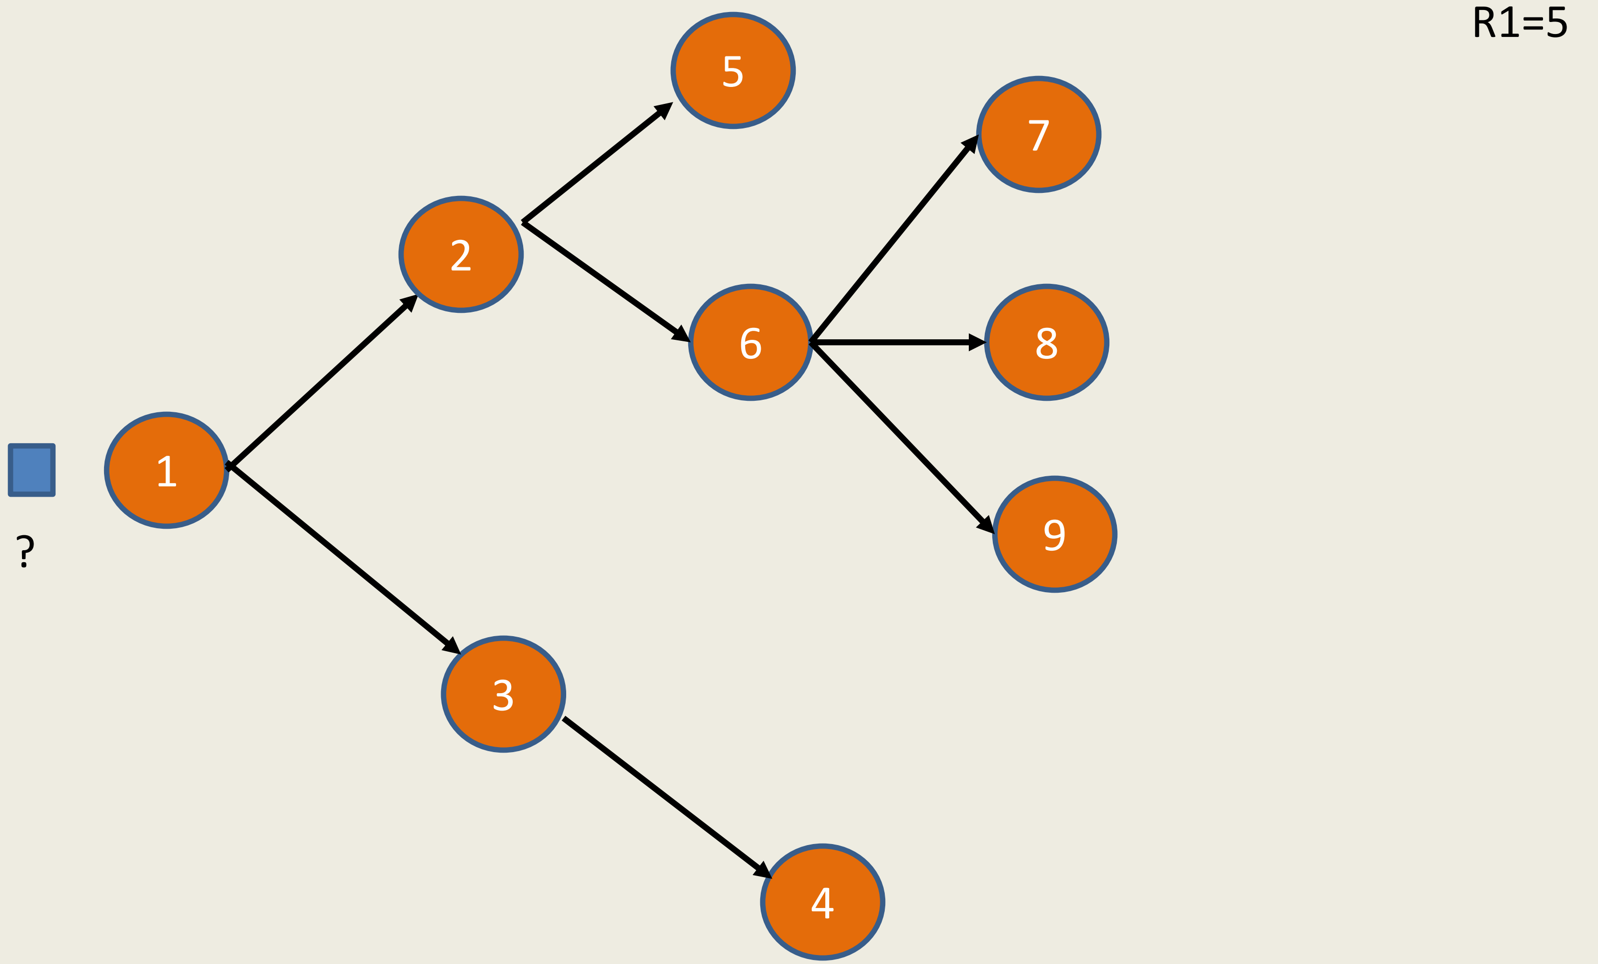

An illustration of TrustWalker

Continue?

Yes

R1=5

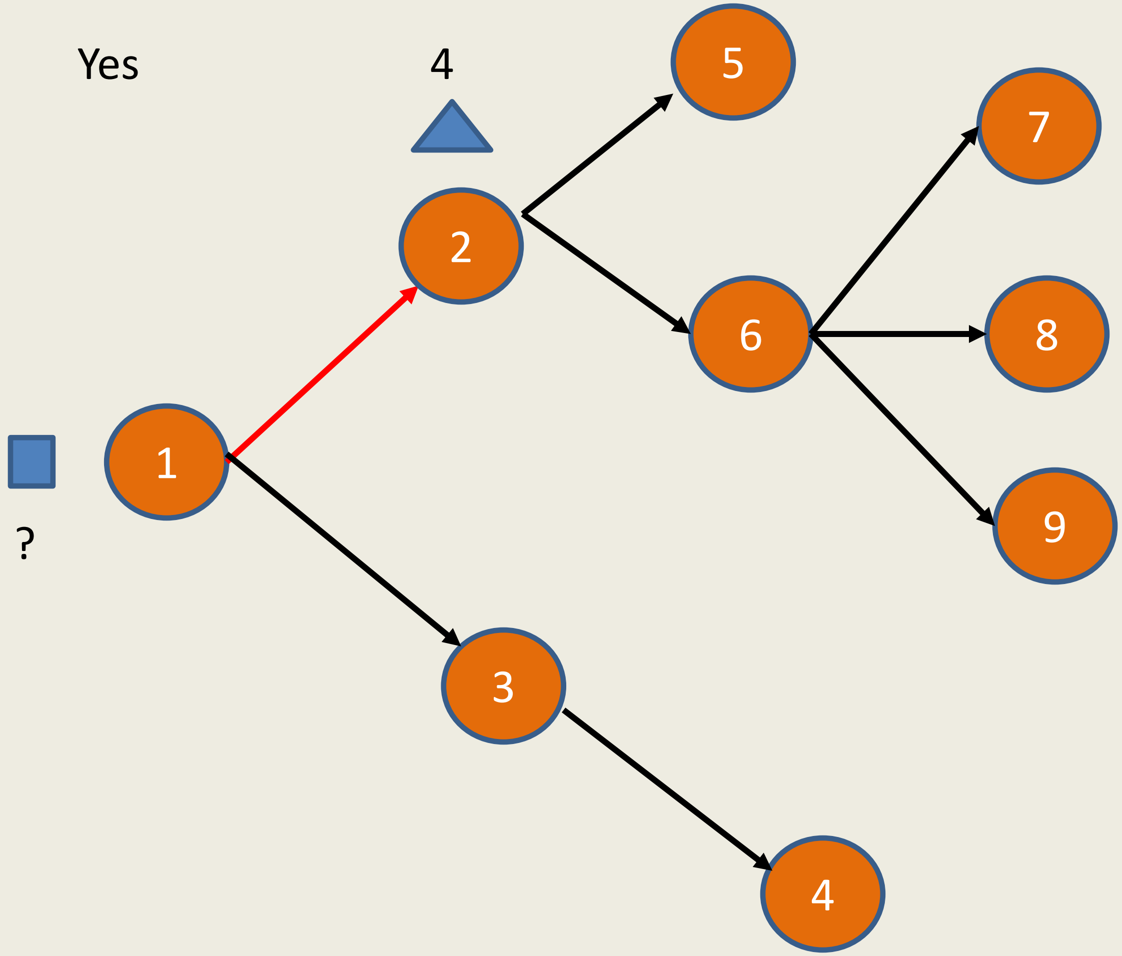

An illustration of TrustWalker

Continue?

Yes

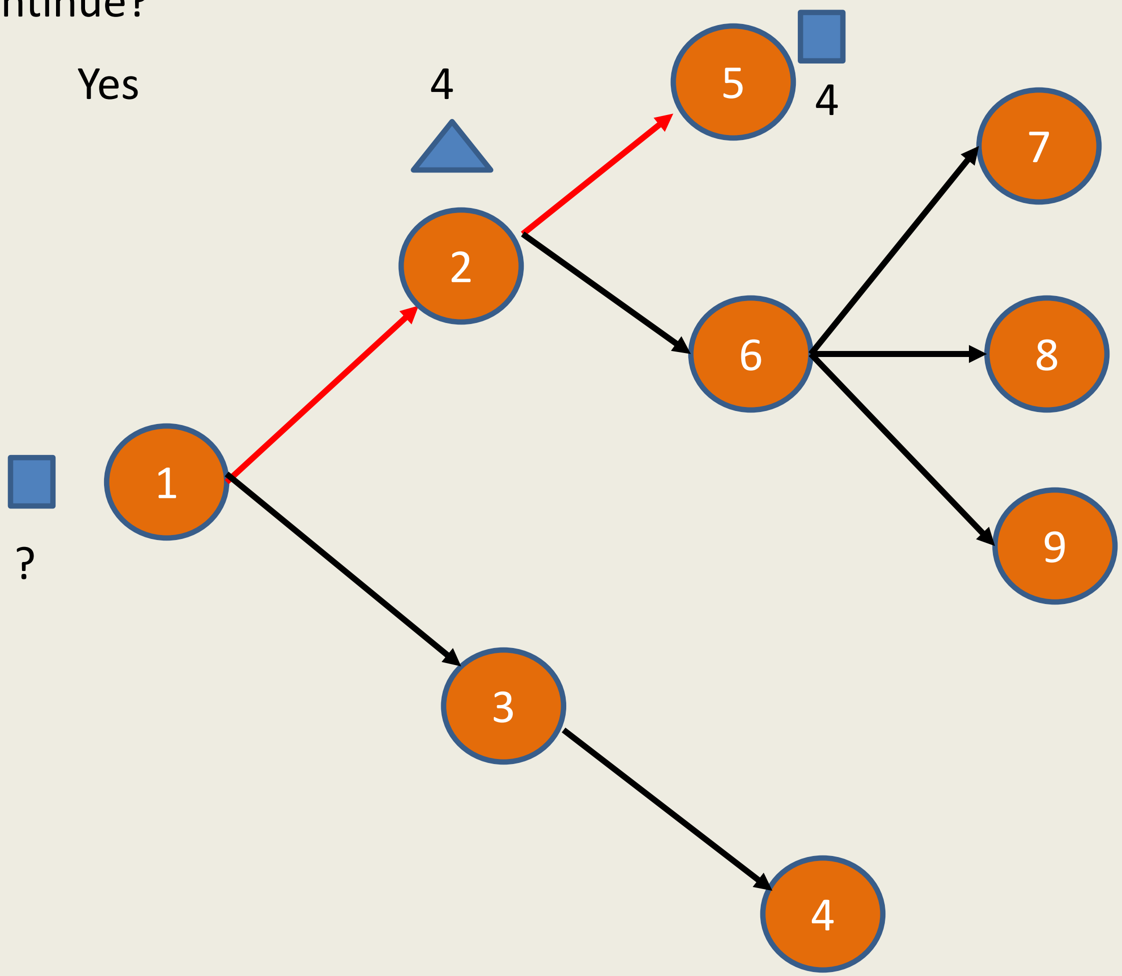

R1=5

R2=4

An illustration of TrustWalker

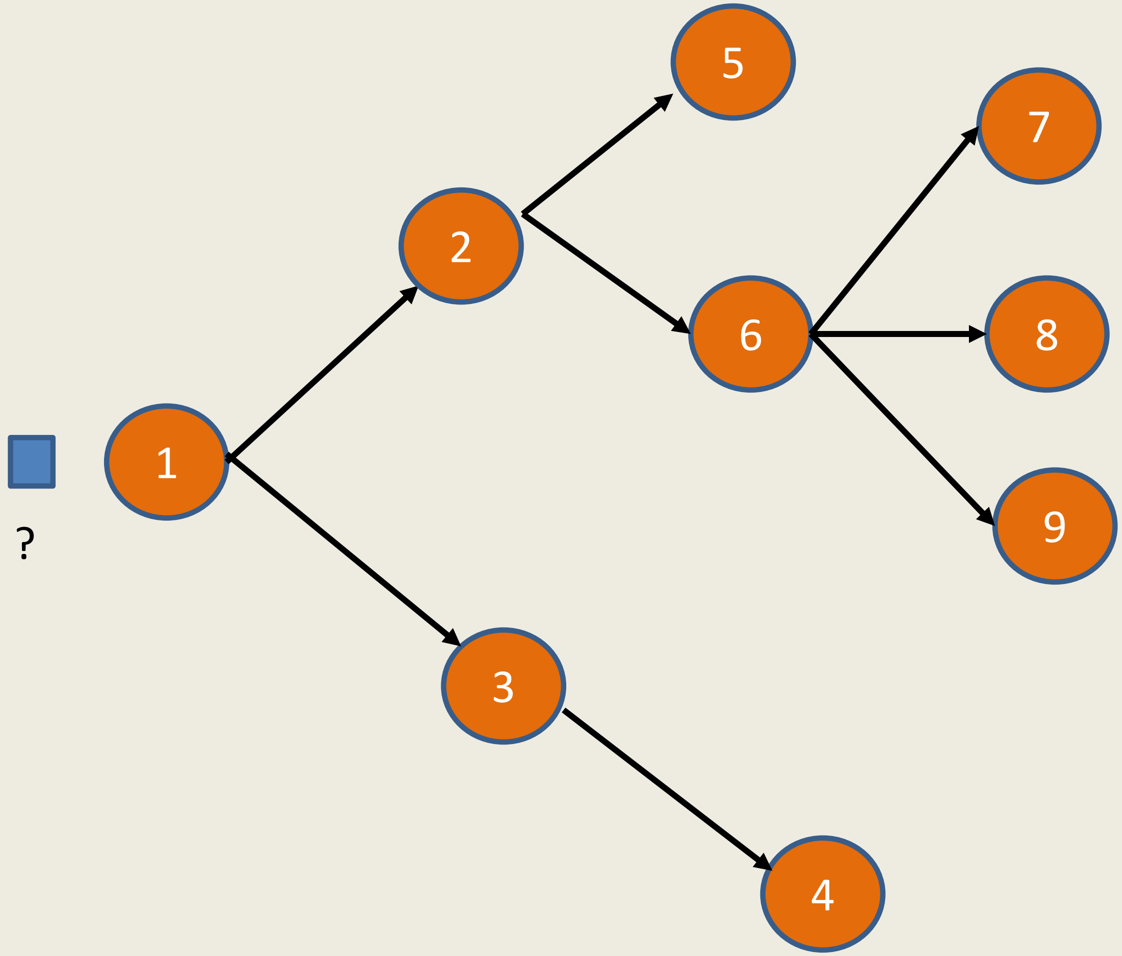

R1=5

R2=4

An illustration of TrustWalker

Continue?

Yes

R1=5

R2=4

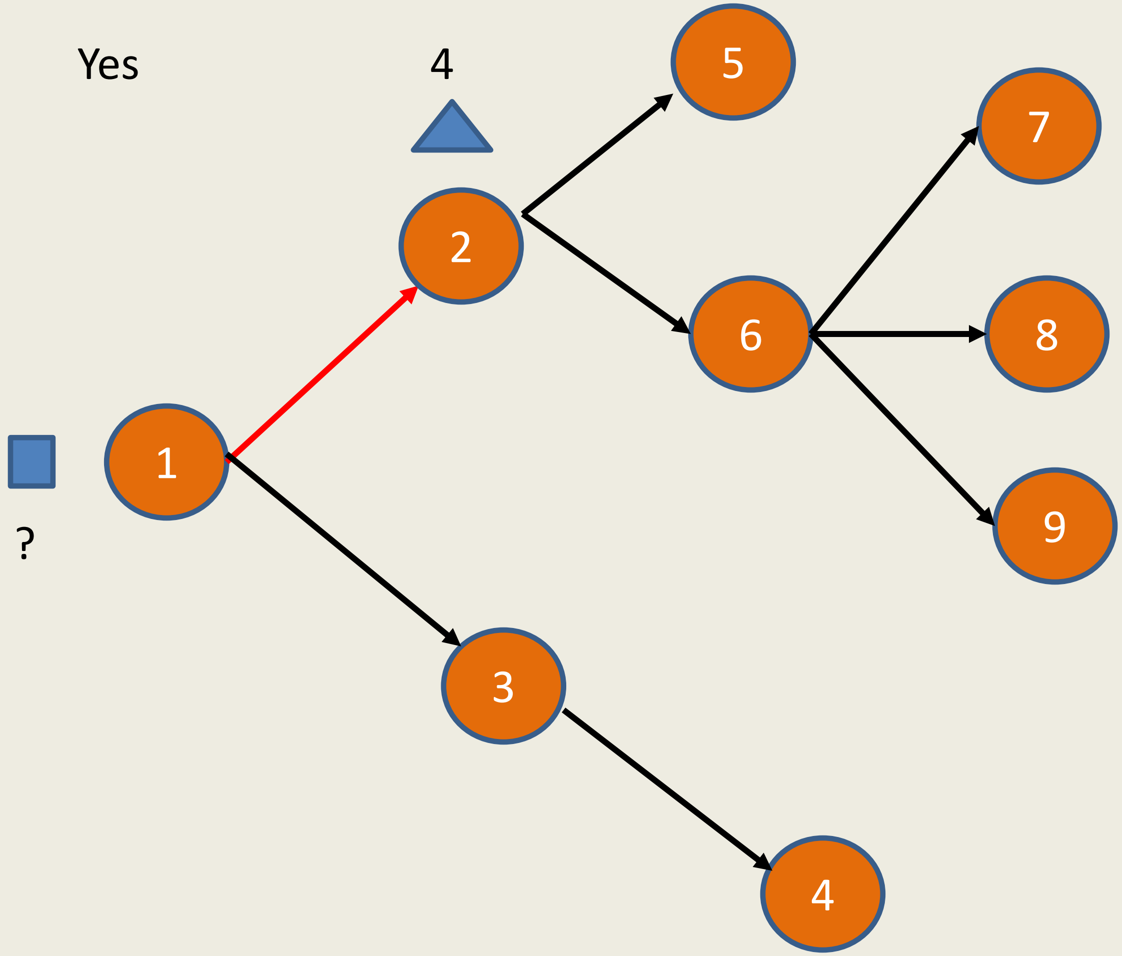

An illustration of TrustWalker

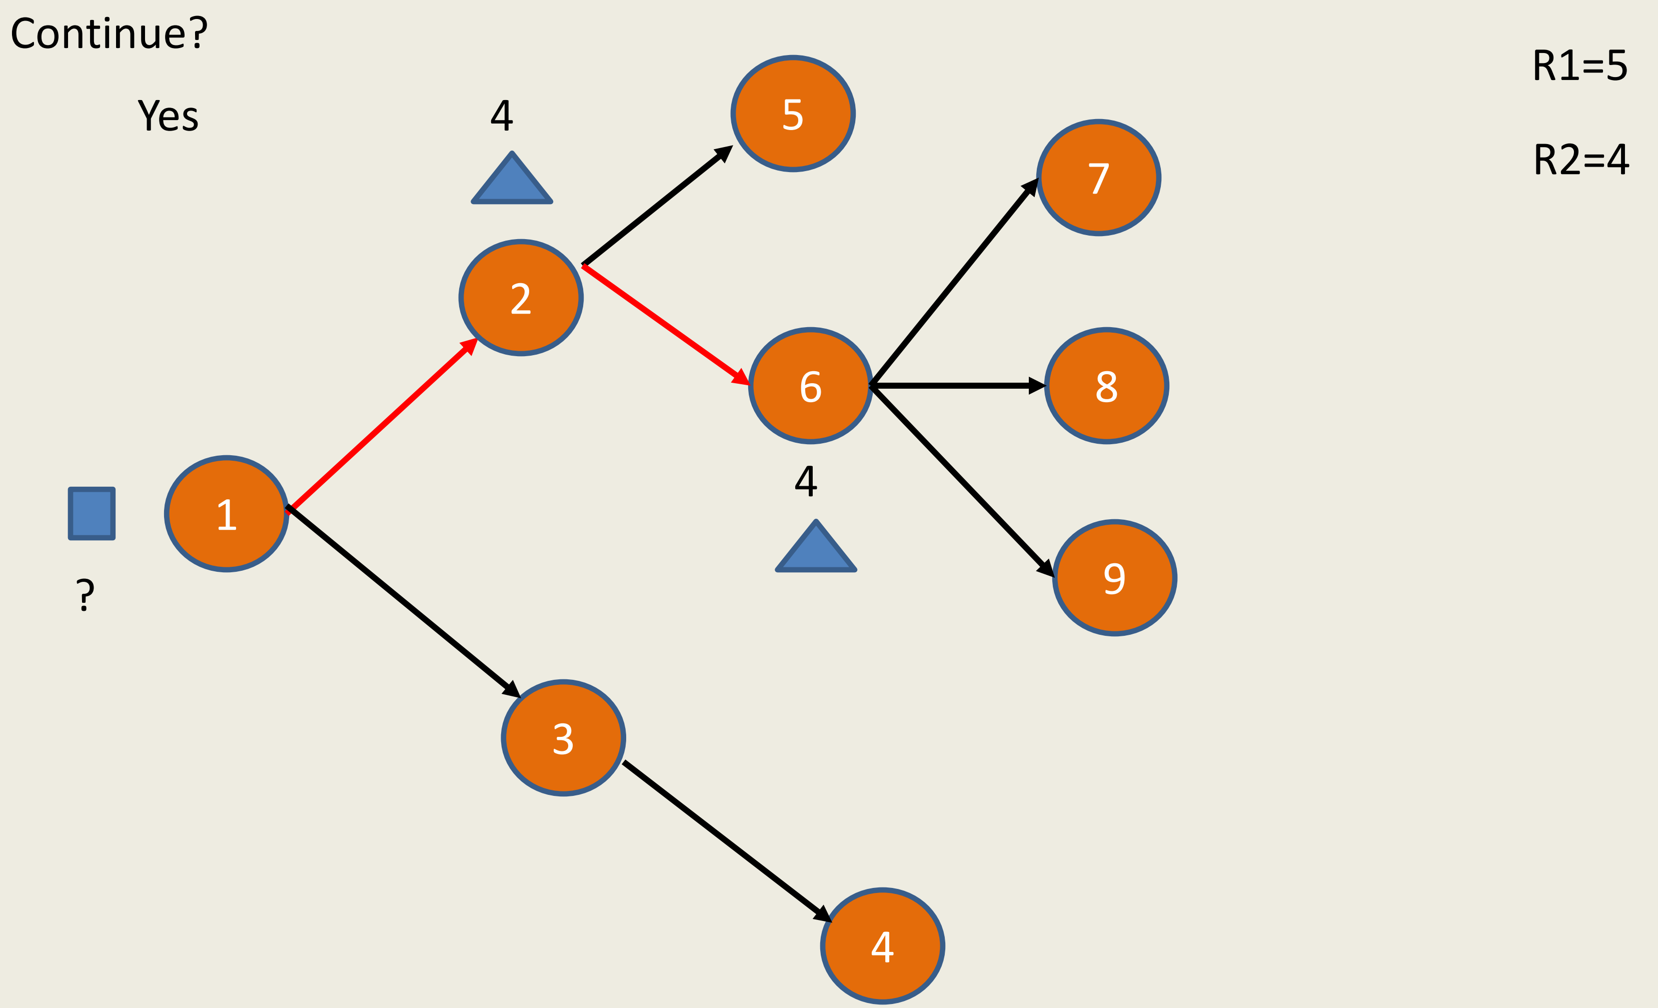

An illustration of TrustWalker

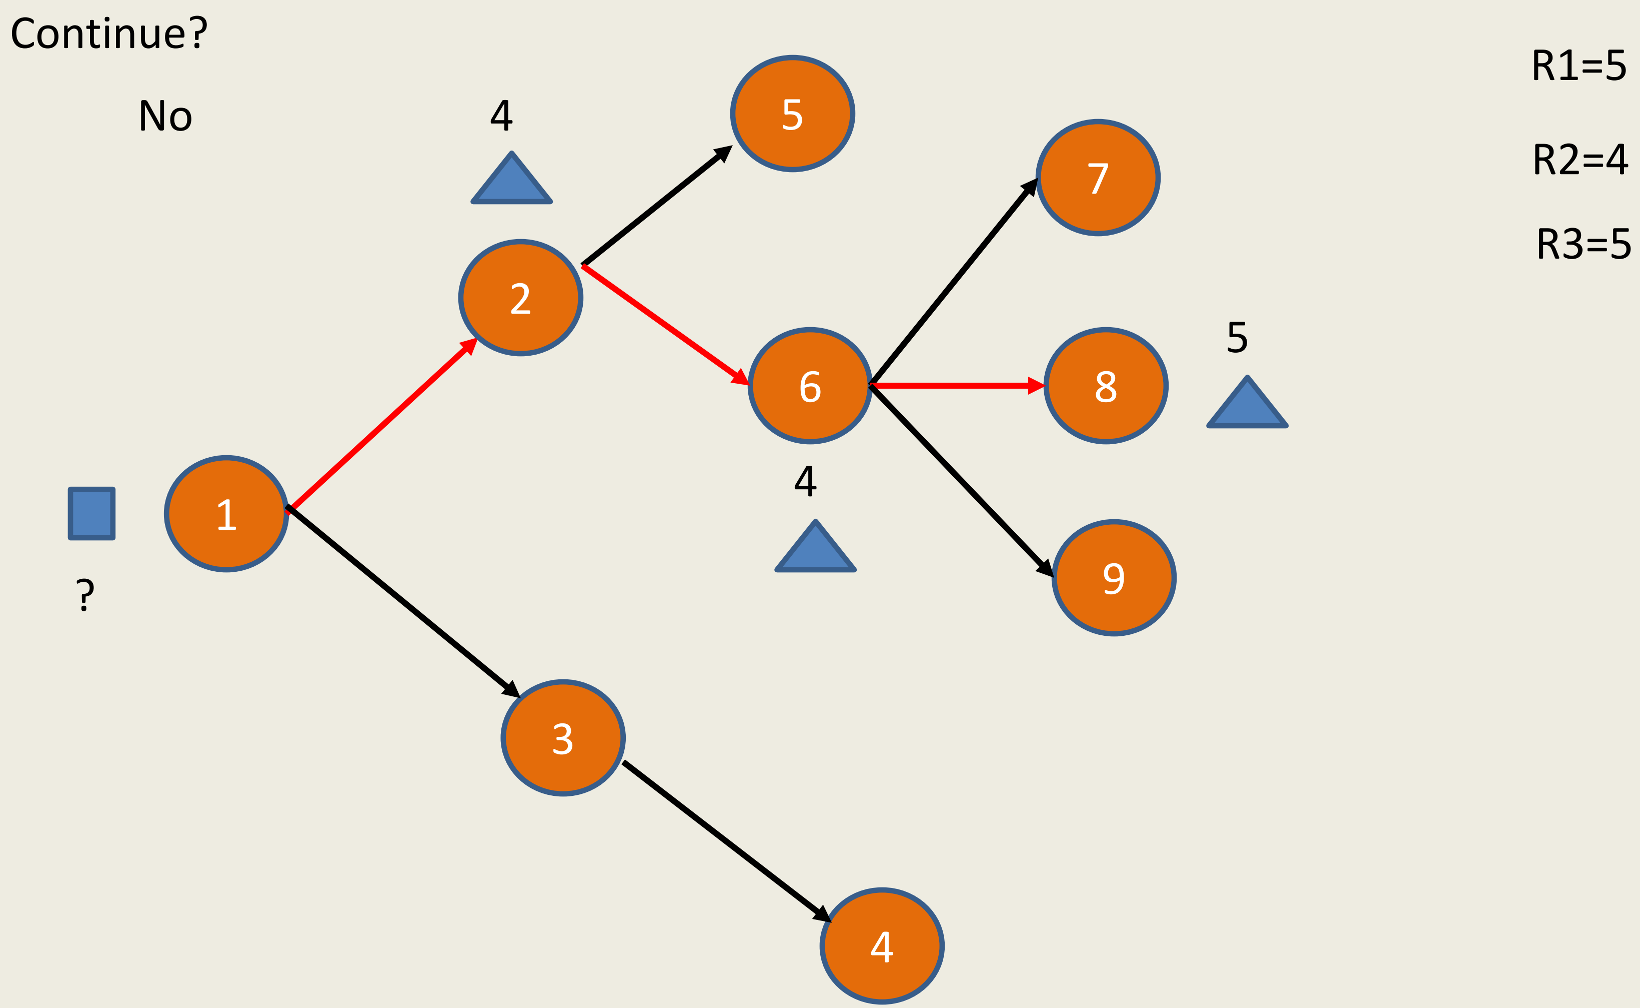

An illustration of TrustWalker

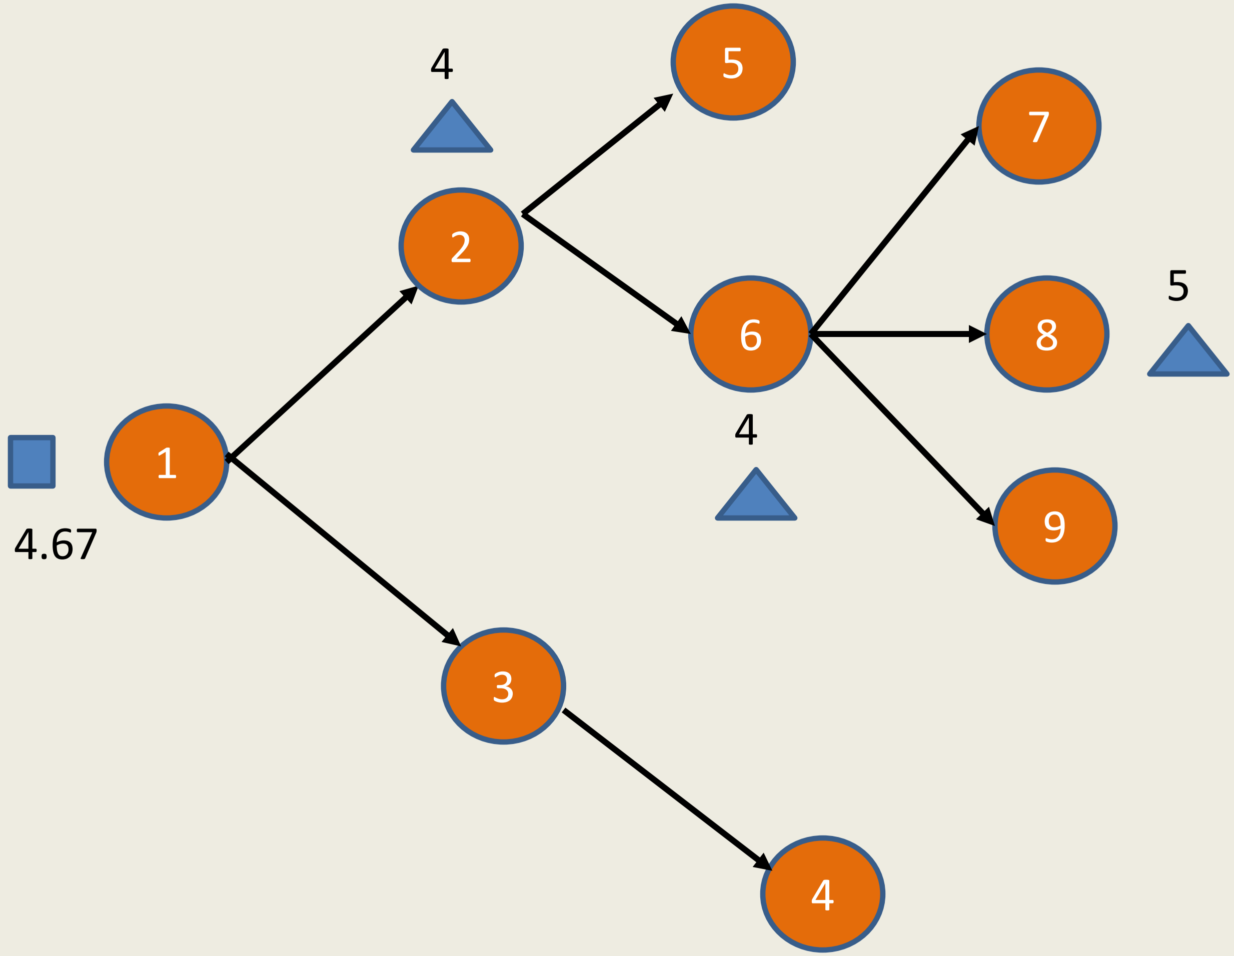

R1=5

R2=4

R3=5

Model-based Trust-aware Recommendation

- Model-based trust-aware recommender systems choose model-based CF methods as their basic models
  – Matrix factorization is widely chosen as the basic model

- There are three common ways to integrate trust information under the matrix factorization framework
  – Co-factorization methods
  – Ensemble methods
  – Regularization methods

Co-factorization Methods

A user shares the same user preference factor in the rating space (rating information) and the social space (social information)

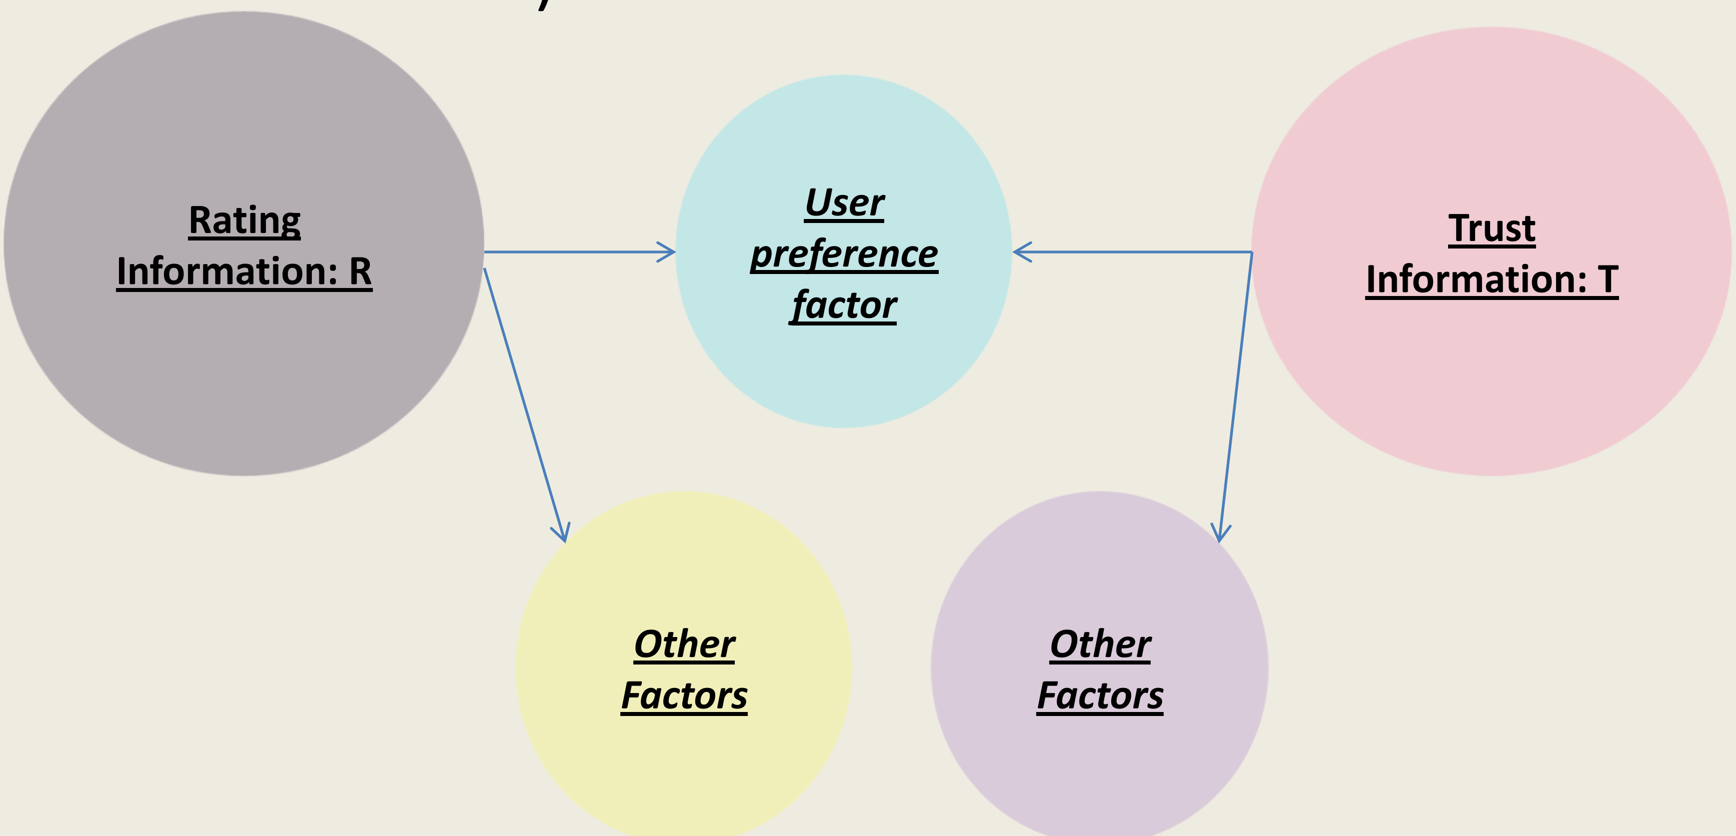

Representative Systems

- SoRec [Ma et al., 2008]

$$R_{ij} = U_i V_j^T$$
$$T_{ij} = U_i Z_j^T$$

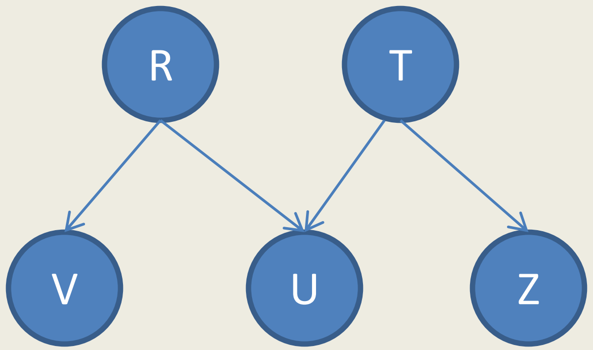

- LOCABAL [Tang et al., 2013]

$$R_{ij} = U_i V_j^T$$
$$T_{ij} = U_i H U_j^T$$

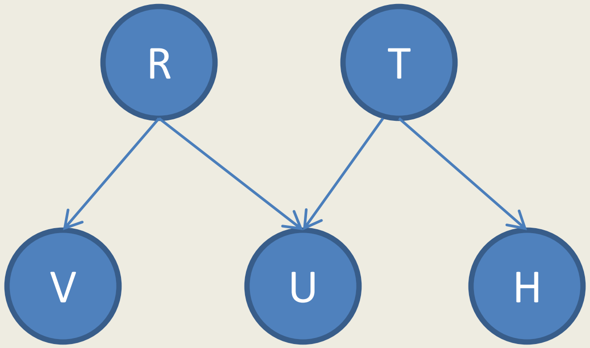

Co-factorization methods can jointly predict missing ratings and trust relations.

Ensemble Methods

- Users and their social networks should have similar ratings on items

- A missing rating for a given user is predicted as a combination of ratings from the user and her trust network

A missing rating for u = u + *a* u's trust network

Representative Systems

- STE - Ensemble of predicted ratings [Ma et al., 2009a]

$$R_{ij} = U_i V_j^T + \boxed{\alpha \sum_{u_k \in N(u_i)} T_{ik} U_k V_j^T}$$

- mTrust – Ensemble of predicted ratings and observed ratings from her trust network [Tang et al., 2012]

$$R_{ij} = U_i V_j^T + \boxed{\alpha \frac{\sum_{u_k \in N(u_i)} T_{ik} R_{jk}}{\sum_{u_k \in N(u_i)} T_{ik}}}$$

- T_{ik} is the trust strength between the i-th user and k-th user

Regularization Methods

- Regularization methods focus on a user’s preference and assume that a user’s preference should be similar to that of her social network.

- Regularization methods add a regularization term to force users’ preferences to be close to those of trust networks.

Representative systems

- SocialMF – a user's preference should be close to that of her social network [Jamali and Ester, 2010]

$$\min \sum_i \| U_i - \boxed{\sum_{u_k \in N(u_i)} T_{ik} U_k} \|$$

- SoReg – two connected users should have similar preferences [Ma et al., 2011]

$$\min \sum_i \sum_{u_k \in N(u_i)} T_{ik} \| U_i - \boxed{U_k} \|$$

Prediction Accuracy Evaluation

- Root Mean Squared Error (RMSE)

$$RMSE = \sqrt{\frac{\sum_{<u_i, v_j> \in O} (\hat{R}_{ij} - R_{ij})^2}{|O|}}$$

- Mean Absolution Error (MAE)

$$MAE = \frac{\sum_{<u_i, v_j> \in O} |\hat{R}_{ij} - R_{ij}|}{|O|}$$

– Small improvement in RMSE or MAE terms can have a significant impact on the quality of the top-few recommendation [Koren, 2008]

Ranking Accuracy Evaluation

- Recall
  - –How many of acquired items were recommended
  - –Recall@N: how many top-N acquired items are recommended
- Precision
  - –How many recommended items are acquired
  - –Precision@N: how many top-N recommended items are acquired

Long recommendation lists typically improve recall while reducing precision

Coverage Evaluation

- Item coverage

–The proportion of items that the recommendation system can recommend

- User coverage

–The proportion of items that the recommendation system can recommend

Incorporating Distrust

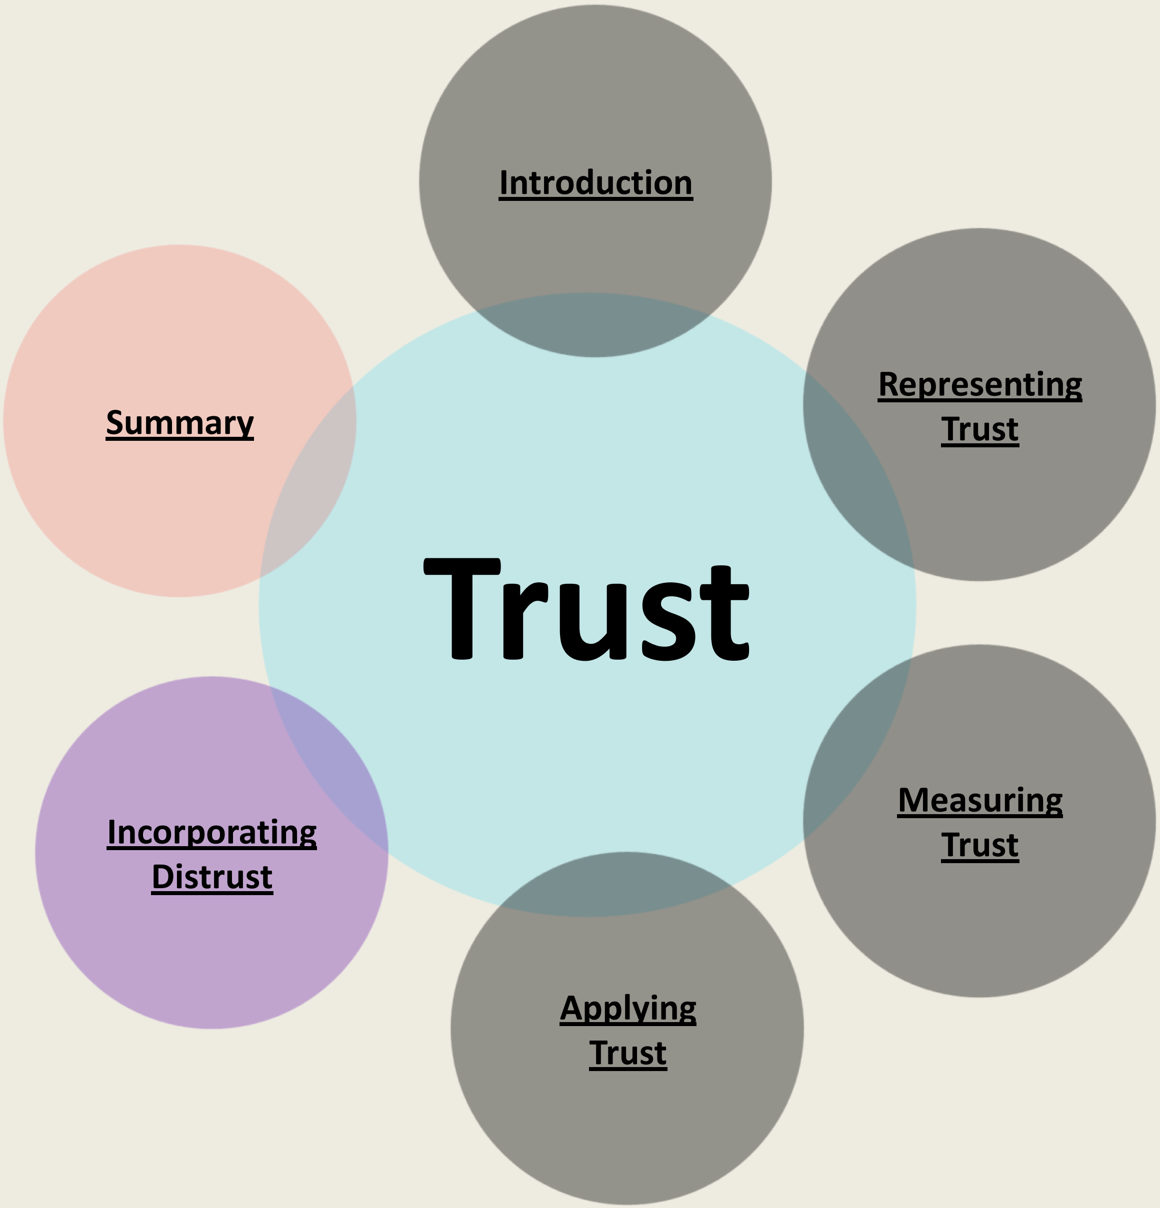

Distrust in Social Sciences

- Distrust can be as important as trust
- Both trust and distrust help a decision maker reduce the uncertainty and vulnerability associated with decision consequences
- Distrust may exert an equally important, if not more, critical role as trust in consumer decisions

Understandings of Distrust from Social Sciences

- Distrust is the negation of trust [Jøsang et al.,2003]
  - Low trust is equivalent to high distrust
  - The absence of distrust means high trust
  - Lack of the studying of distrust matters little

- Distrust is a new dimension of trust [Lewicki et al., 1998]
  - Trust and distrust are two separate concepts
  - Trust and distrust can co-exist
  - A study ignoring distrust would yield an incomplete estimate of the effect of trust

Distrust in Social Media

- Distrust is rarely studied in social media
  - Social media data is based on passive observations
  - Lack of some information social sciences use to study distrust
  - Lack of computational understanding of distrust with social media data
- Let us first examine the properties of distrust before going to the computational understanding of distrust

Examining Properties of Distrust

- Properties of trust are systematically and extensively studied
  - Transitivity, asymmetry, and homophily

- Properties of distrust are rarely studied with social media data

- Can we simply or conversely extend the properties of trust to those of distrust ?
  - We study the properties of distrust in parallel to those of trust

Transitivity [Tang and Liu, 2014]

- Trust is transitive (1st table); how about distrust?
- For distrust (2nd table), #u+w is comparable to #u-w
- Transitivity may not be applicable to distrust

Trust		
Types	Number	Percentage
$\langle u{+}v, v{+}w \rangle$, $u{+}w$	3,320,991	97.75%
$\langle u{+}v, v{+}w \rangle$, u-w	76,613	2.25%

Distrust		
Types	Number	Percentage
$\langle u$-v, v-$w \rangle$, $u{+}w$	38,729	59.73%
$\langle u$-v, v-$w \rangle$, u-w	26,114	40.27%

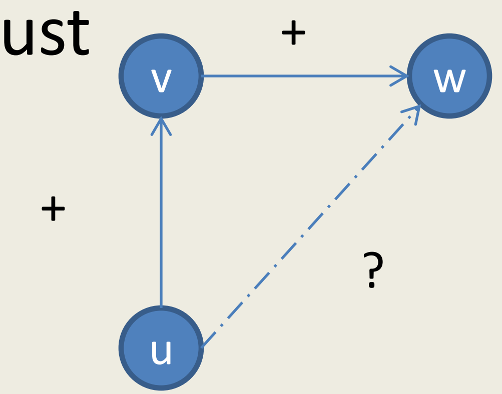

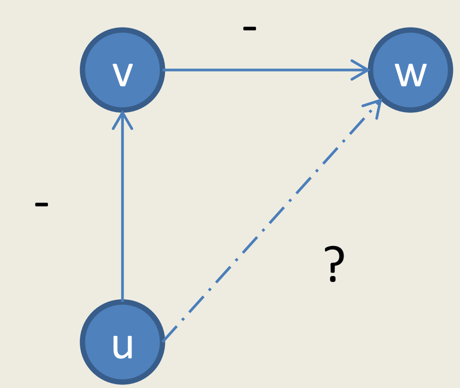

u+v and u-v represent trust and distrust relations between u and v, respectively

Understanding (-,-,+) and (-,-,-) [Guha et al., 2004]

- User u disagrees with the statement of user v
  – <u-v, v-w> → <u+w>
  – My enemy’s enemy is my friend
  – Structural balance

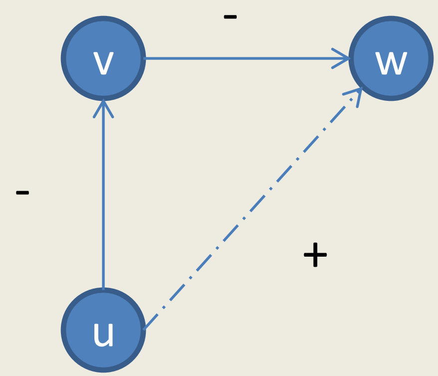

- User u thinks user v ’s judgments are inferior to her own
  – <u-v, v-w> → <u-w>
  – Status theory

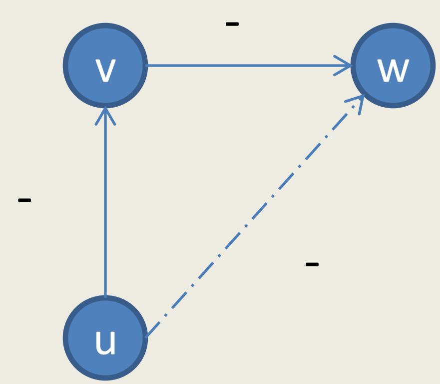

- So, it is more complicated than it appears and further research is needed

Asymmetry [Tang and Liu, 2014]

- Trust is asymmetric
  - 37.61% of relations are mutual-trust
  - Hence, u+v is not equivalent to v+u, or v+u $\not\equiv$ u+v

- Distrust is even more skewed
  - Only 5.86% of relations are mutual-distrust
  - That is, we are more confident that v-u $\not\equiv$ u-v

	v+u(%)	v-u(%)	v?u(%)
u+v	136,806(37.61)	967(0.27)	226,000(62.13)
u-v	967(2.09)	2,623(5.86)	42,606(92.23)

Homophily [Tang and Liu, 2014]

- Users with distrust relations are more likely to be similar than two randomly chosen users

	CI	COSINE	COSINE-CI
Distrust ($\mathbf{s}_d$)	0.4994	0.0105	0.0142
Trust $\mathbf{s}_t$	0.6792	0.0157	0.0166
Random Pairs ($\mathbf{s}_r$)	0.1247	0.0027	0.0032
P1	9.57e-87	1.19e-120	4.88e-45
P2	1.71e-132	5.83e-157	3.72e-108
P3	7.84e-23	1.99e-19	9.32e-17

CI: Commonly-rated Items

COSINE: Rating-cosine similarity

COSINE-CI: Rating-cosine similarity of commonly rated items

P1- P-values: H_0: $\mathbf{s}_d \le \mathbf{s}_r$; H_1: $\mathbf{s}_d > \mathbf{s}_r$

P2- P-values: H_0: $\mathbf{s}_t \le \mathbf{s}_r$; H_1: $\mathbf{s}_t > \mathbf{s}_r$

P3- P-values: H_0: $\mathbf{s}_t \le \mathbf{s}_d$; H_1: $\mathbf{s}_t > \mathbf{s}_d$

- Distrust is not a dissimilarity measurement

Computational Understanding of Distrust

- We leverage machine learning and data mining techniques to design computational tasks to help us understand distrust with passively observed social media data

- Task 1: Is distrust the negation of trust?
  – If distrust is the negation of trust, distrust should be predictable from only trust

- Task 2: Can we predict trust better with distrust?
  – If distrust is a new dimension of trust, distrust should have added value on trust and can improve trust prediction

- The first step to understand distrust is to make distrust computable by incorporating distrust in trust models

Distrust in Trust Representations

There are three major ways to incorporate distrust in trust representation

- Considering low trust as distrust
- Extending negative values in trust representations
- Adding a dimension in trust representations

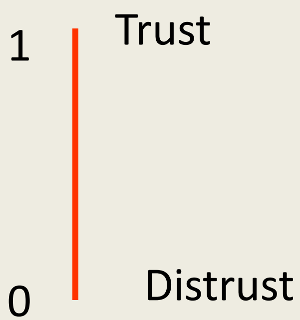

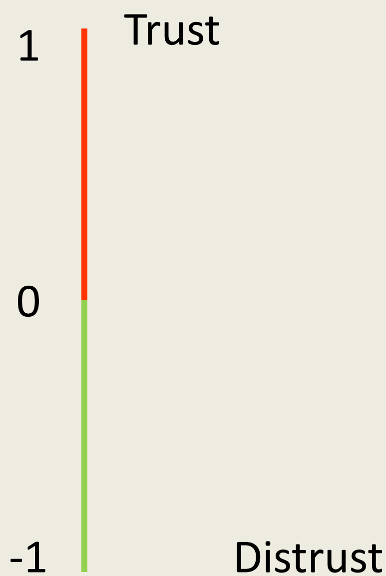

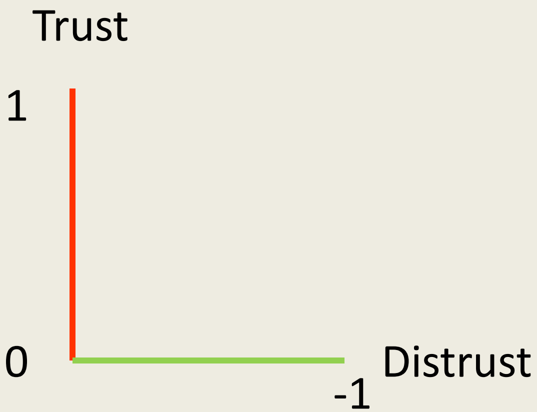

A Network Illustration of Distrust in Trust Representations

- Considering low trust as distrust
  – Weighted unsigned network

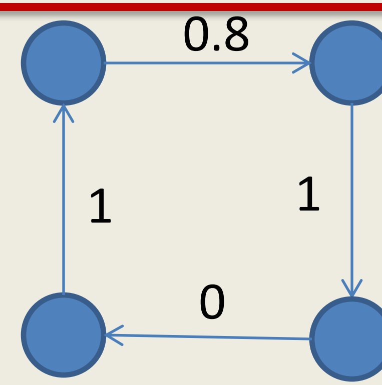

- Extending negative values in trust representations
  – Weighted signed network

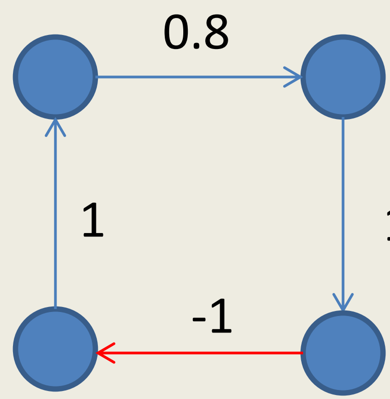

- Adding a dimension in trust representations
  – Two-dimensional unsigned network

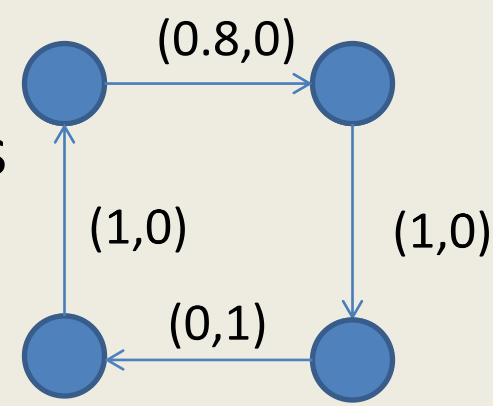

Task 1: Is Distrust the Negation of Trust?

- If distrust is the negation of trust, low trust is equivalent to distrust and distrust should be predictable from trust

IF	**Distrust**	≡	**Low Trust**
THEN	**Predicting Distrust**	≡	**Predicting Low Trust**

- Given the transitivity of trust, we resort to trust prediction algorithms to compute trust scores for pairs of users in the same trust network

Algorithm for Task 1

- A trust predictor is chosen to calculate trust scores for pairs of users without trust relations

- Pairs of users with low trust scores are suggested as distrust

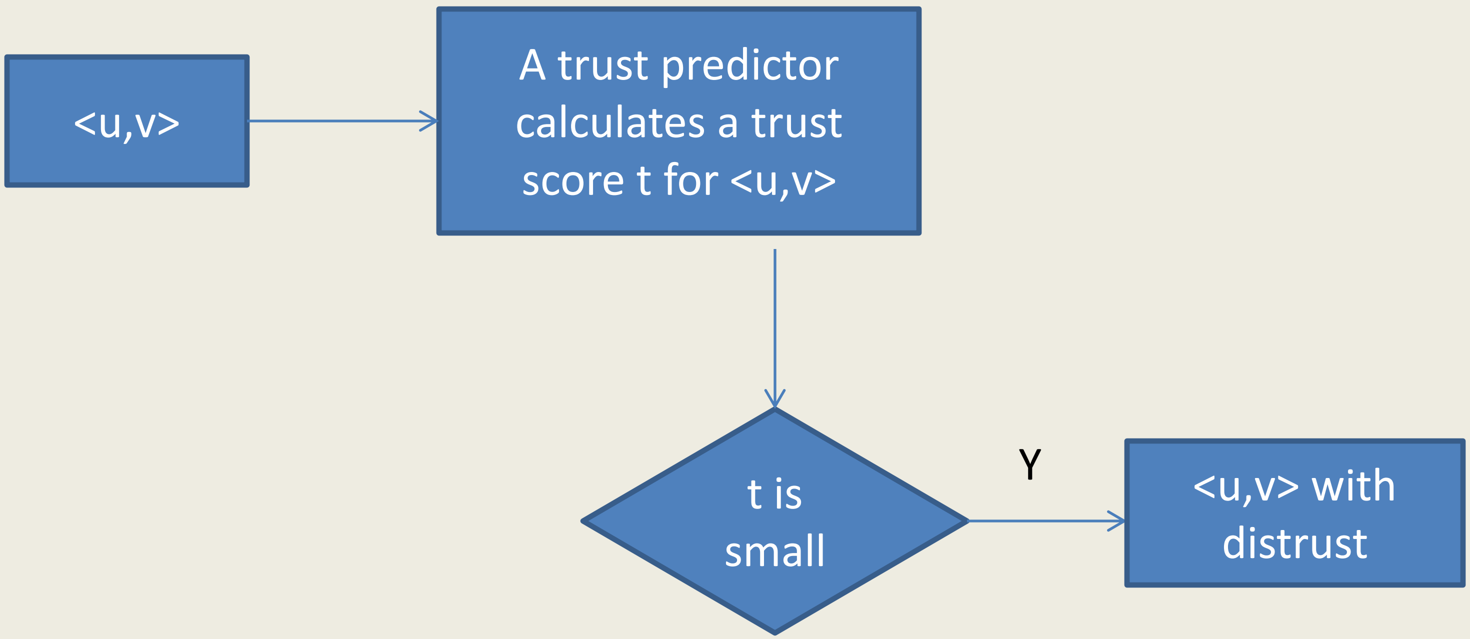

Experimental Settings for Task 1

- Each time we choose x% of pairs of users with trust relations as A_T^x

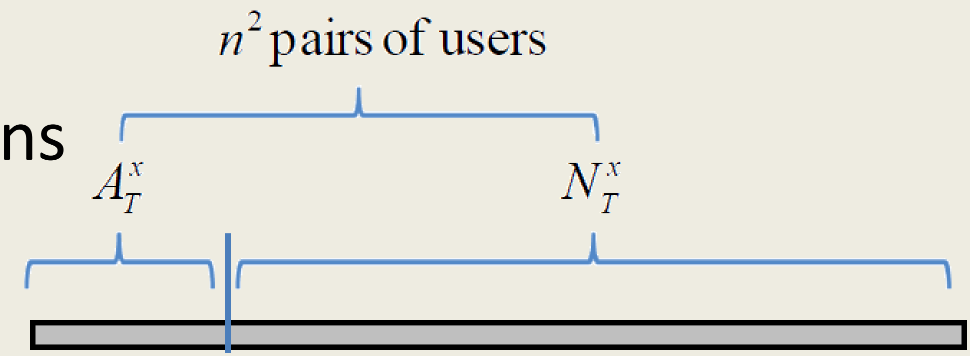

- Through Task 1 with A_T^x, we predict $|D|$ pairs of users with low trust P from N_T^x as distrust
  – D is the set of pairs with distrust as ground truth in the data set

- The performance is computed as

$$PA = \frac{|D \cap P|}{|D|}$$

Evaluation of Task 1

- The performance of using low trust to predict distrust is consistently worse than randomly guessing
- Task 1 fails to predict distrust with only trust and distrust is not the negation of trust

x (%)	dTP ($\times 10^{-5}$)	dMF($\times 10^{-5}$)	dTP-MF($\times 10^{-5}$)	Random($\times 10^{-5}$)
50	4.8941	4.8941	4.8941	5.6824
55	5.6236	5.6236	5.6236	8.1182
60	7.1885	7.1885	7.1885	15.814
65	11.985	11.985	11.985	19.717
70	13.532	13.532	13.532	18.826
80	10.844	10.844	10.844	16.266
90	12.720	12.720	12.720	25.457
100	14.237	14.237	14.237	29.904

dTP: It uses trust propagation to calculate trust scores for pairs of users

dMF: It uses the matrix factorization based predictor to compute trust scores for pairs of users

dTP-MF: It is the combination of dTP and dMF using OR

Task 2: Can we predict Trust better with Distrust

- If distrust is a new dimension of trust, distrust should provide additional information about users, and could have added value beyond trust

- We seek answer to whether using both trust and distrust information can help achieve better performance than using only trust information

- Task 2 is to incorporate distrust into trust measurements

Distrust in Trust Metrics

- Most trust and distrust metrics are based on the model by extending negative values to trust models to incorporate distrust

- The introduction of distrust in trust networks converts unsigned trust networks to signed trust and distrust networks

– Social theories for signed network such as balance theory and status theory can be used to understand trust and distrust

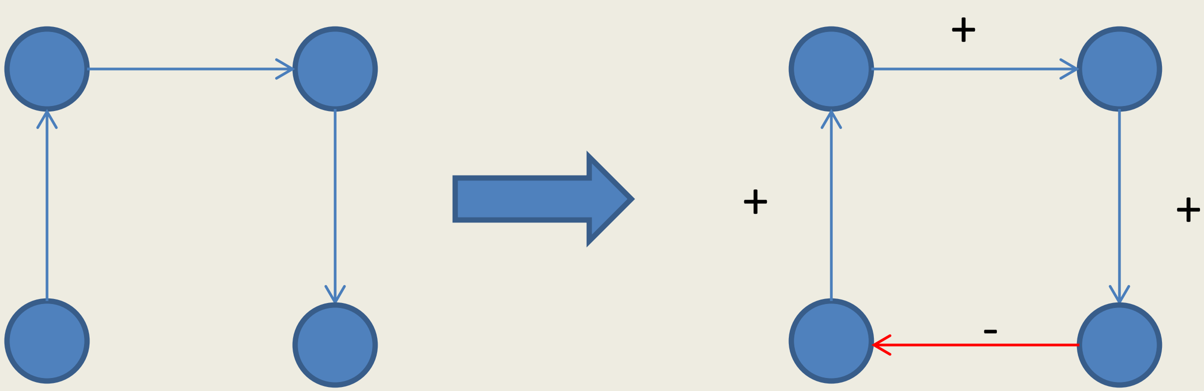

Balance Theory [Heider, 1946]

- Balance theory suggests that "the friend of my friend is my friend" and "the enemy of my enemy is my friend"

- For a triad, there are four possible sign combinations A(+,+,+), B(+,+,-) C(+,-,-) and D(-,-,-), but only A(+,+,+) and C(+,-,-) are balanced

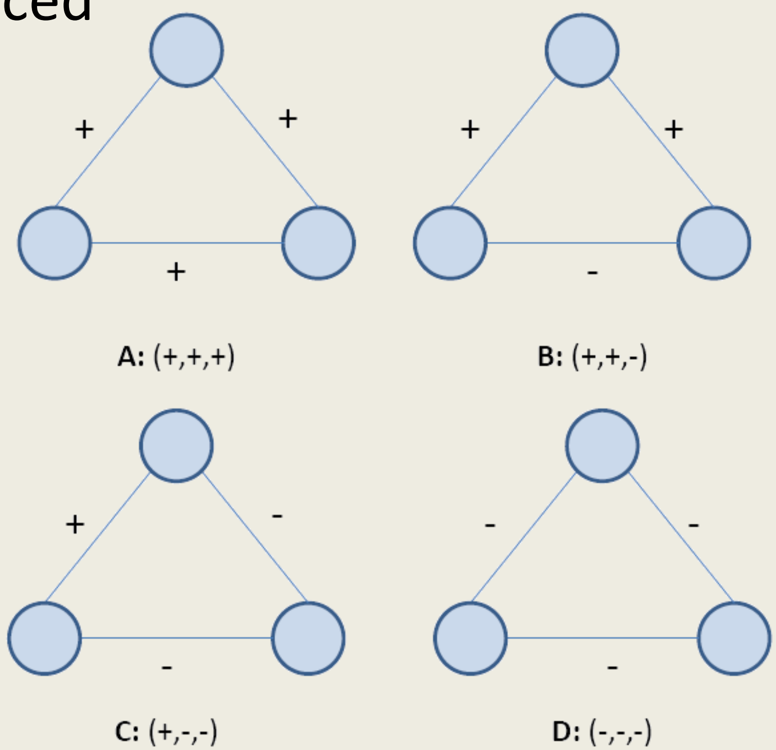

Balance theory is developed for undirected networks and can be applied to directed networks by ignoring their directions.

Status Theory [Leskovec et al., 2010a]

- Status theory is developed for directed networks
- A positive link from u to v indicates that u has a higher status than v
- A negative link from u to v indicates that u has a lower status than v
- For a triad, status theory suggests that if we take each negative relation, reverse its direction, and flip its sign to positive, then the resulting triangle (with all positive edge signs) should be acyclic

Status theory

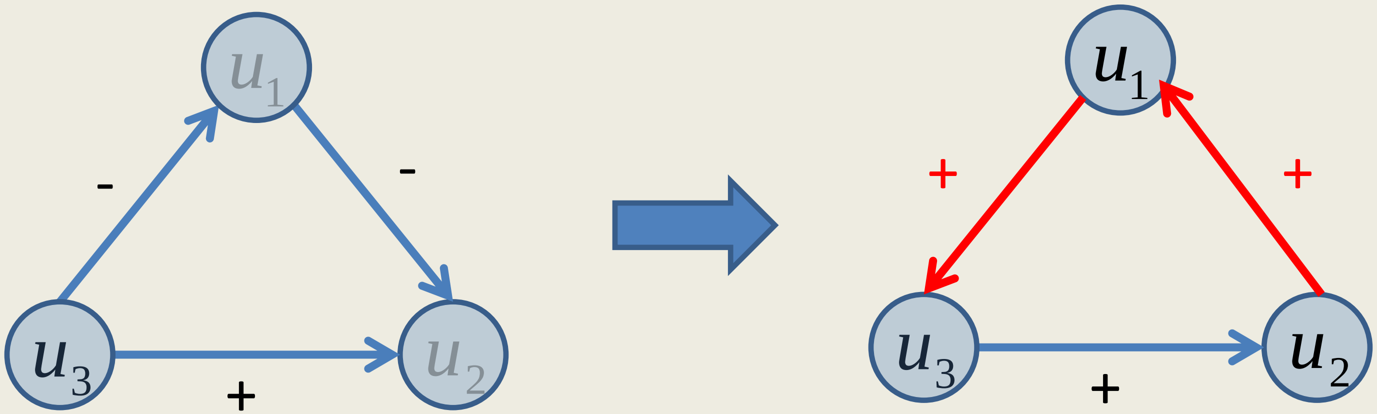

Cyclic,
invalid

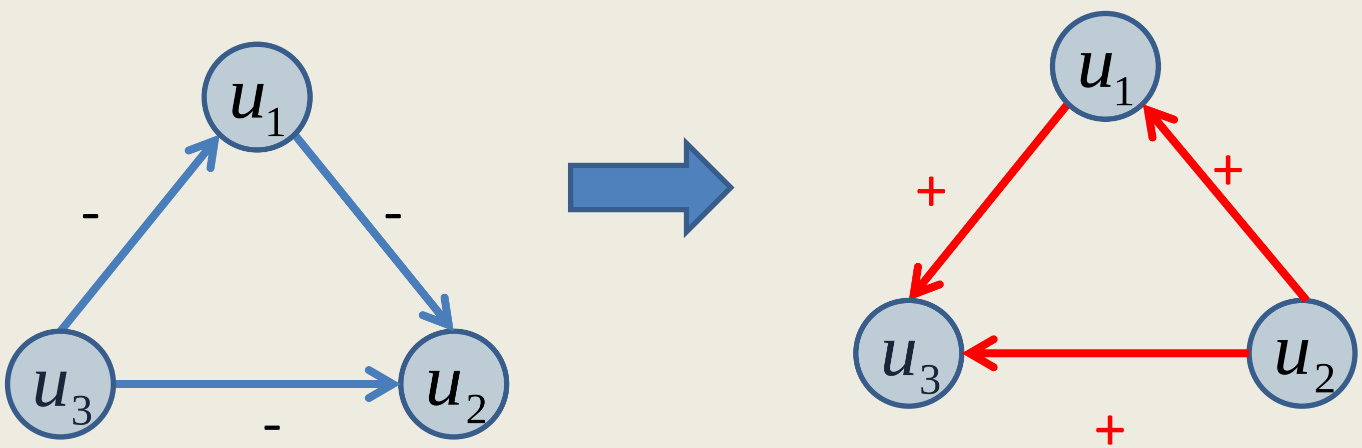

Acyclic,
valid

Trust and Distrust Propagation [Guha et al., 2004]

- A single step of distrust propagation in trust propagation

Trust propagation:

$$\mathbf{C} = \alpha_1 \mathbf{T} + \alpha_2 \mathbf{T}^\top \mathbf{T} + \alpha_3 \mathbf{T}^\top + \alpha_4 \mathbf{T}\mathbf{T}^\top$$

One step distrust propagation:

$$\tilde{\mathbf{G}} = \sum_{k=1}^{K} \gamma^k \mathbf{C}^k (\mathbf{T} - \mathbf{D})$$

- Multiple steps of distrust propagation in trust propagation

Trust and distrust propagation:

$$\mathbf{E} = \alpha_1 \mathbf{F} + \alpha_2 \mathbf{F}^\top \mathbf{F} + \alpha_3 \mathbf{F}^\top + \alpha_4 \mathbf{F}\mathbf{F}^\top$$

Propagation aggregation:

$$\tilde{\mathbf{G}} = \sum_{k=1}^{K} \gamma^k \mathbf{E}^k$$

Trust and Distrust Matrix Factorization [Tang and Liu, 2014]

- For each user u_i, we introduce one dimensional latent variable r_i, and then F_{ij} is modeled as $F_{ij} = r_i r_j$ to capture balance theory
  - Distrust relations are represented by negative values in **F = T-D**

- With modeling balance theory, trust and distrust matrix factorization disMF models F_{ij} as

$$F_{ij} = U_i V U_j^T + \lambda r_i r_j$$

$U_i V U_j^T$ models correlation between user preferences

$r_i r_j$ models balance theory

λ controls contributions from balance theory

Experimental Settings for Task 2

- Each time we choose x% of pairs of users with trust relations A_T^x as old trust relations and the remaining as new trust relations A_T^n

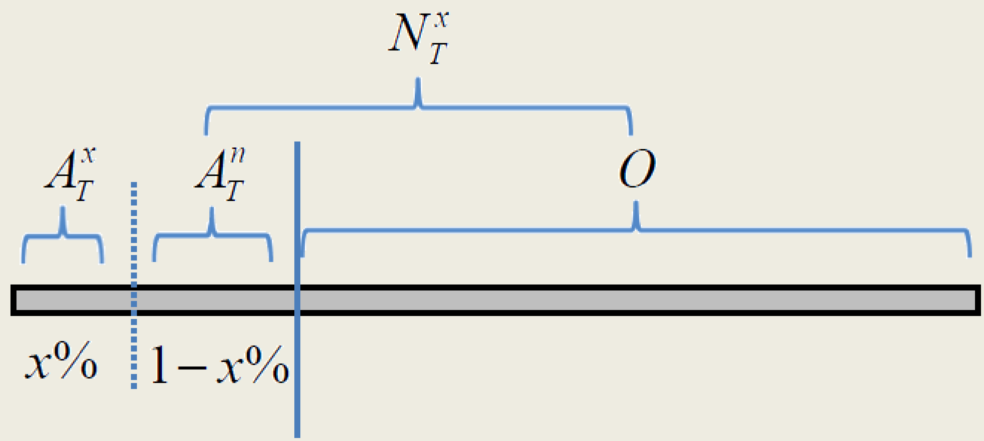

- Through Task 2 with A_T^x and D, we predict $|A_T^n|$ pairs of users P from N_T^x as trust

–D is the set of pairs with distrust as ground truth in the data set

- The performance is computed as

$$PA = \frac{|A_T^n \cap P|}{|A_T^n|}$$

Evaluation of Trust and Distrust Propagation

- Incorporating distrust propagation into trust propagation can improve the performance of trust measurement
- One step distrust propagation usually outperforms multiple step distrust propagation

	50%	55%	60%	65%	70%	80%	90%
TP	**0.1376**	0.1354	0.1293	0.1264	0.1201	0.1156	0.1098
disTP-s	**0.1435**	0.1418	0.1372	0.1359	0.1296	0.1207	0.1176
disTP-m	**0.1422**	0.1398	0.1359	0.1355	0.1279	0.1207	0.1173
Random	**0.0023**	0.0023	0.0020	0.0019	0.0018	0.0015	0.0013

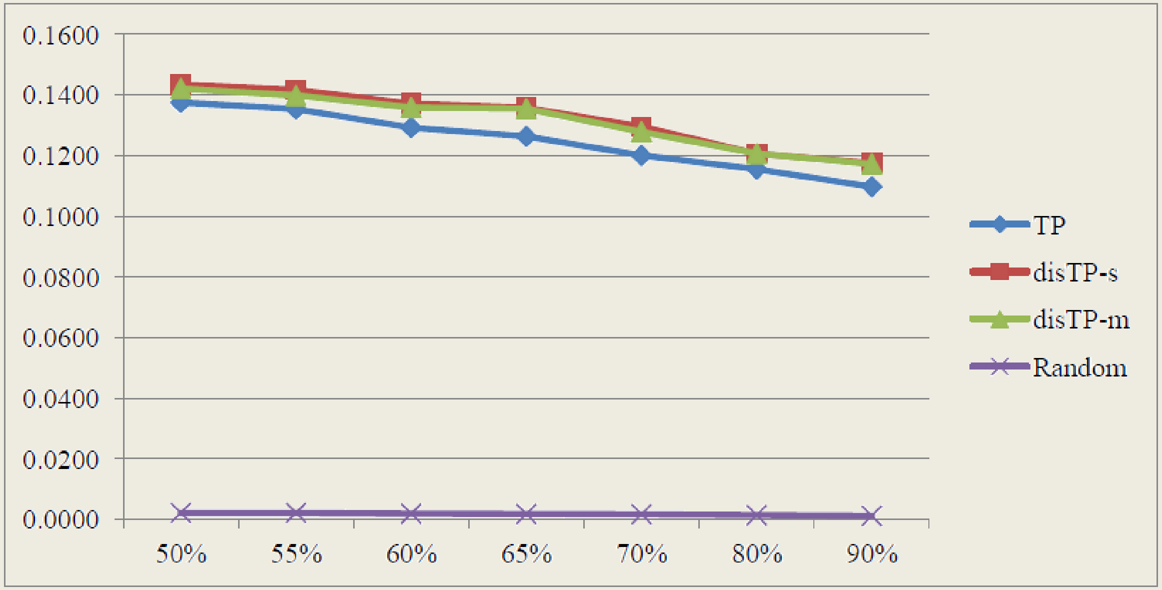

Evaluation of Trust and Distrust Matrix Factorization

- Incorporating distrust with balance theory can significantly improve the performance of trust prediction

	50%	55%	60%	65%	70%	80%	90%
MF	**0.1531**	0.1502	0.1489	0.1444	0.1391	0.1332	0.1277
disMF	**0.1665**	0.1654	0.1639	0.1601	0.1563	0.1498	0.1415
Random	**0.0023**	0.0023	0.0020	0.0019	0.0018	0.0015	0.0013

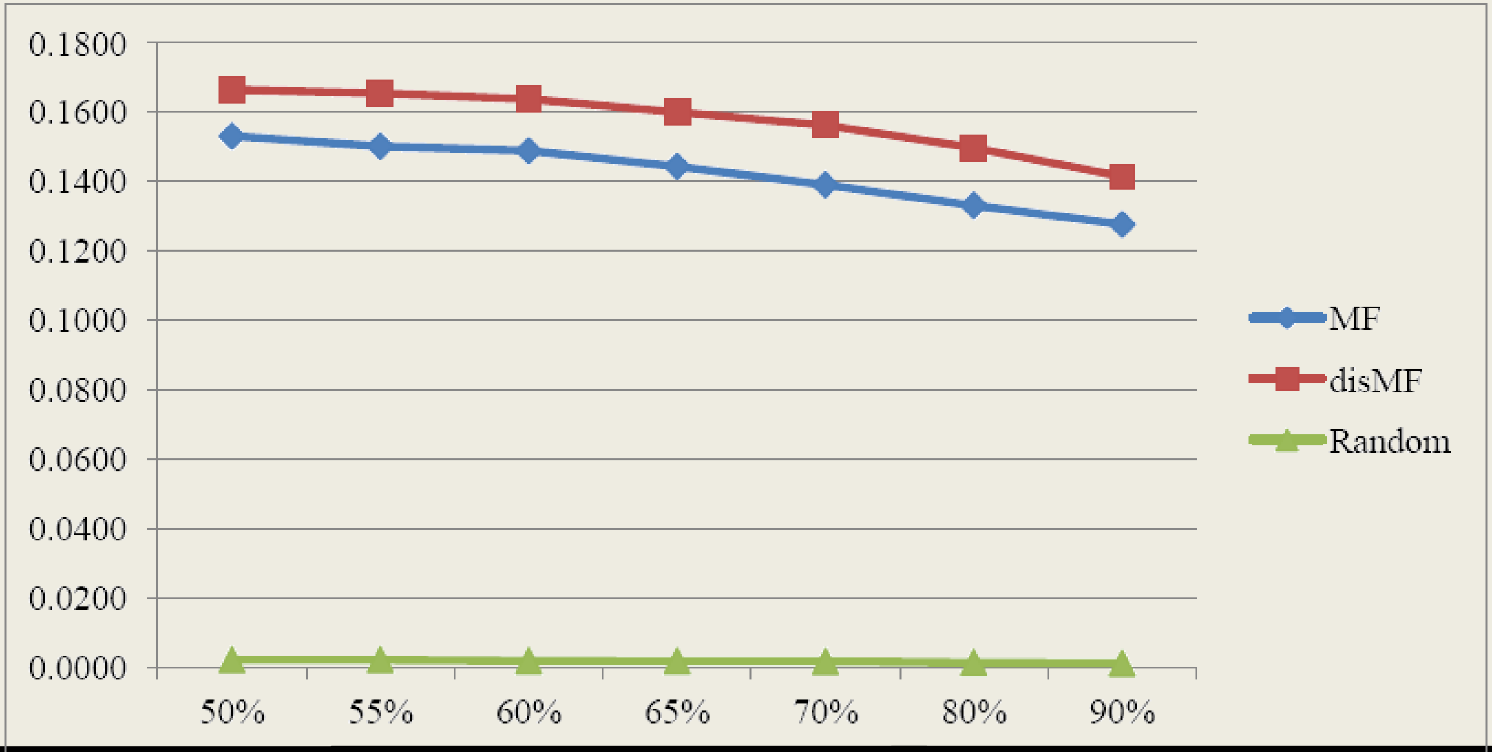

Findings from Task 1 and Task 2

- Task 1 shows that distrust is not the negation of trust
  – Low trust is not equivalent to distrust
  – Distrust is not the negation of trust

- Task 2 shows that the performance of trust prediction is improved by incorporating distrust
  – Distrust has added value in addition to trust
  – Incorporating distrust can improve trust computation

Distrust in Trust-aware Recommender Systems

- Genuine distrust information tends to be more noticeable and credible, and weighed more in decision making than trust information of a similar magnitude

- Users might or might not accept recommendations from their trusted users, but will certainly exclude recommendations from their distrusted users

Distrust in Memory-based Trust-aware Systems [Victor et al., 2009]

- Distrust as a filter

– Use distrust to filter out ``unwanted'' users in the recommendation processes

$$\hat{r}_{u,i} = \bar{r}_u + \frac{\boxed{\sum_{v \setminus \mathcal{D}}} (r_{v,i} - \bar{r}_v) \times t_{u,v}}{\sum t_{u,v}}$$

- Distrust as a dissimilarity measure

– Consider distrust scores as negative weights

$$\hat{r}_{u,i} = \bar{r}_u + \frac{\sum_{v \setminus \mathcal{D}} (r_{v,i} - \bar{r}_v) \times t_{u,v}}{\sum t_{u,v}} \boxed{- \frac{\sum_{v \in \mathcal{D}} (r_{v,i} - \bar{r}_v) \times d_{u,v}}{\sum d_{u,v}}}$$

Distrust in Model-based Trust-aware systems [Ma et al., 2009]

- Users with distrust relations should have very different user preferences

$$\max \sum_{i} \sum_{u_k \in D(u_i)} D_{ik} \| U_i - U_k \|$$

  – Distrust is considered a dissimilarity measure
  – Forcing users preference of two users with a distrust relation far away from each other

Mistrust, Untrust, and Distrust [Marsh and Dibben, 2005]

- Mistrust is misplaced trust
  – A trustee betrays the trust of the trustor

- Untrust is a measure of how little the trustee is actually trusted
  – The trustor has little confidence in the trustee

- Distrust is a measure of how much the trustor believes that the trustee will actively work against them in a given situation

Illustration of Mistrust, Untrust, and Distrust [Marsh and Dibben, 2005]

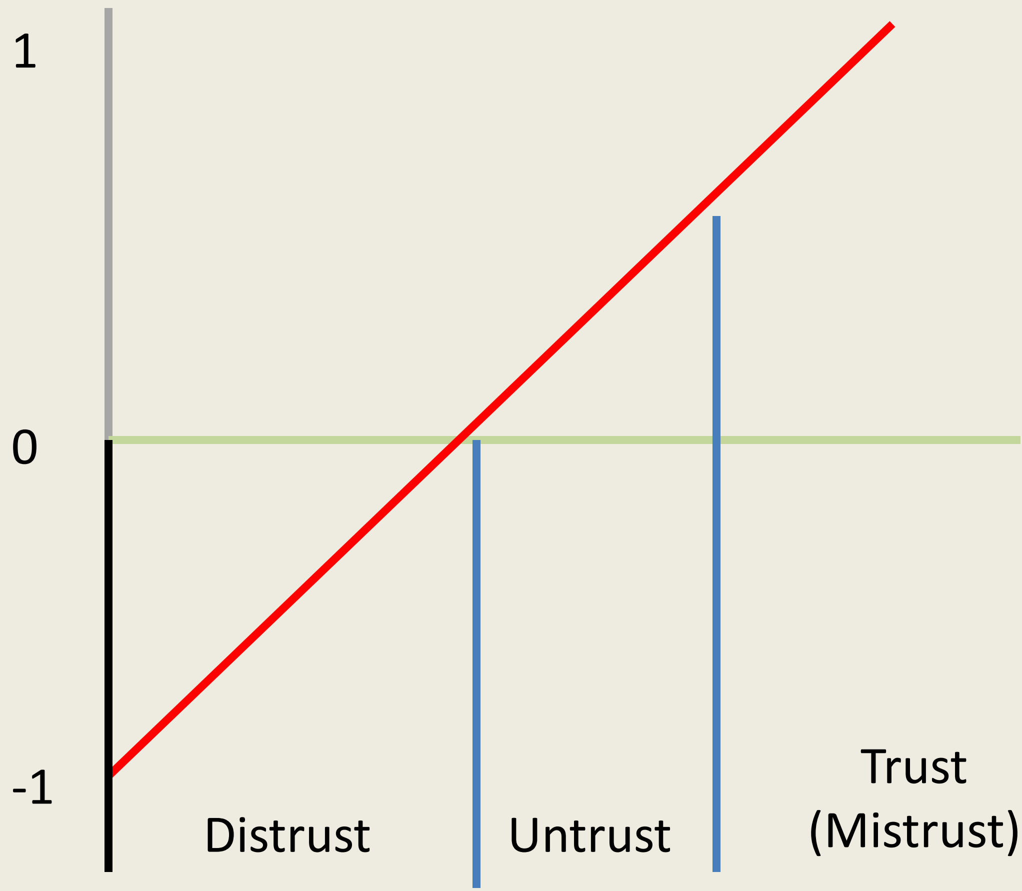

- Untrust is still a positive measurement
- While distrust is a negative measurement

Summary

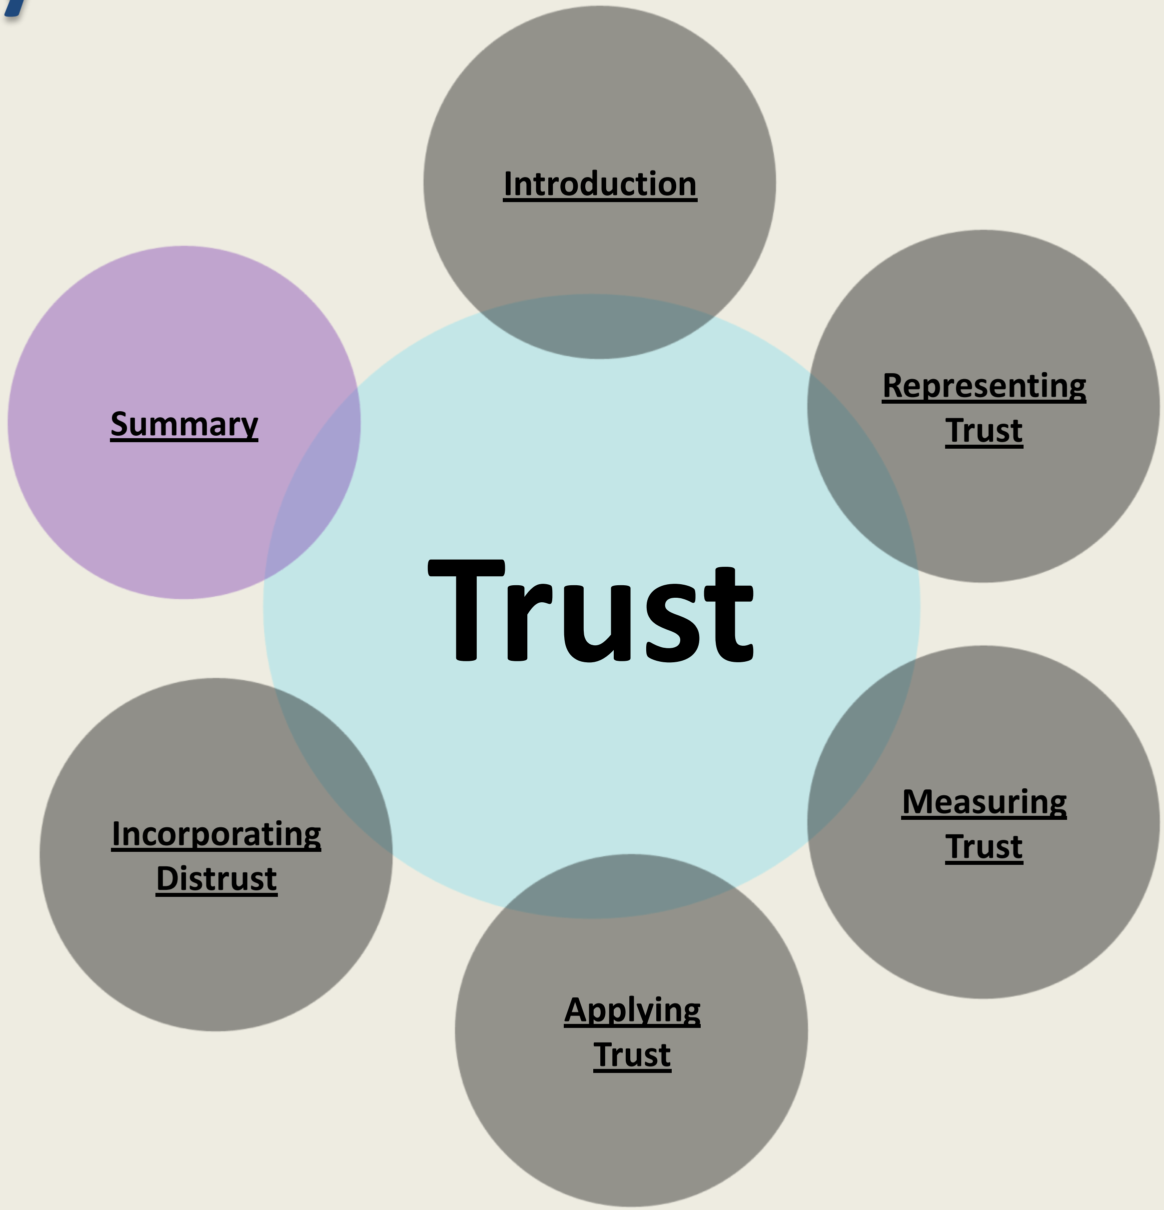
Introduction
Summary
Representing Trust
Trust
Incorporating Distrust
Measuring Trust
Applying Trust

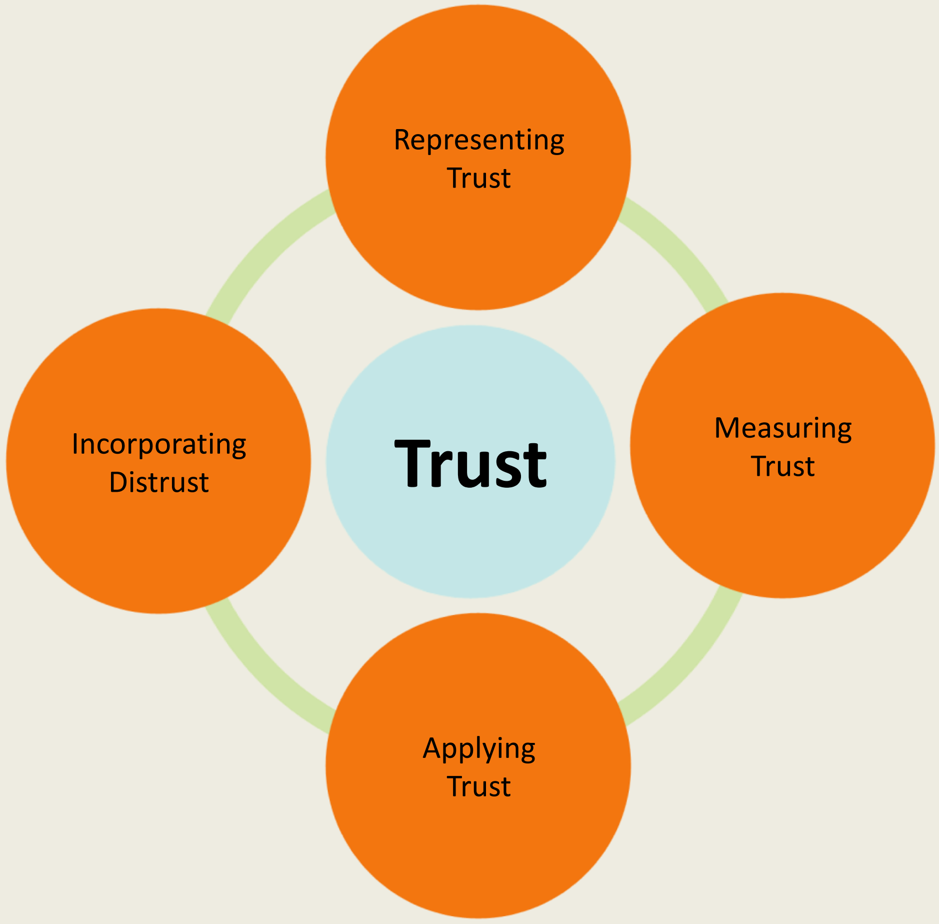
Representing Trust
Incorporating Distrust
Trust
Measuring Trust
Applying Trust

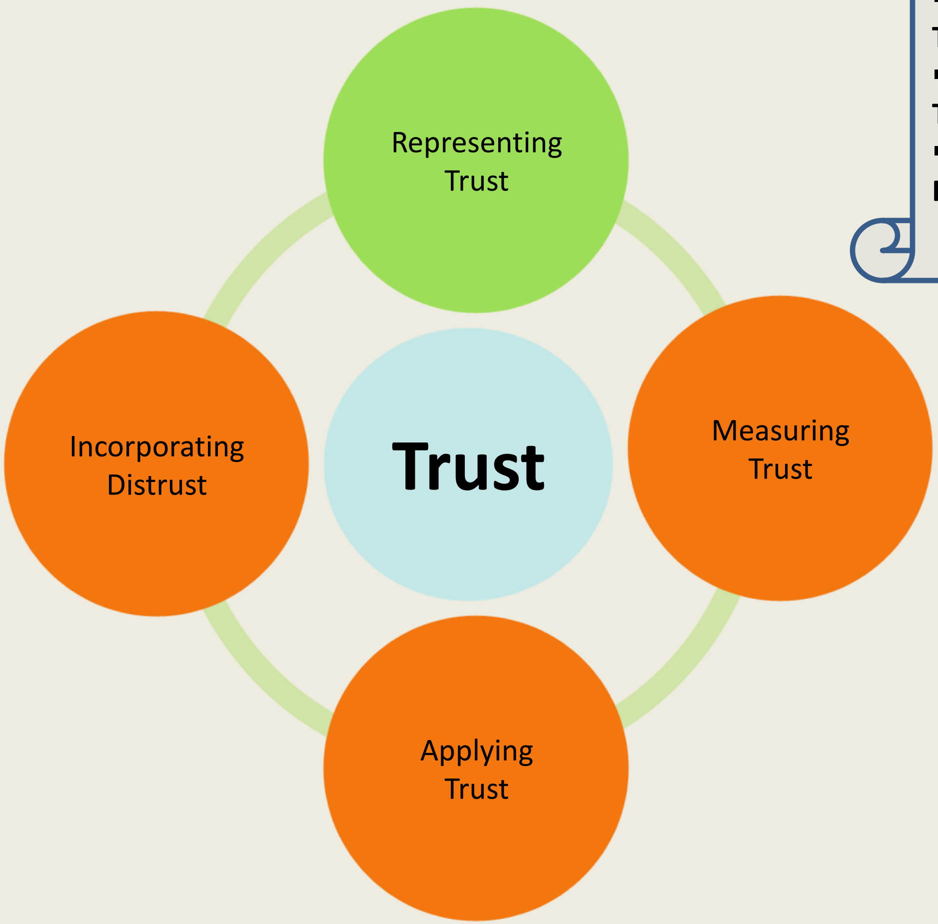

- Properties of Trust
- Classifications of Trust Representations
- Probabilistic vs Gradual Trust Representations
- Single vs Multi-dimensional Trust Representations
- Trust vs Trust and Distrust Representations

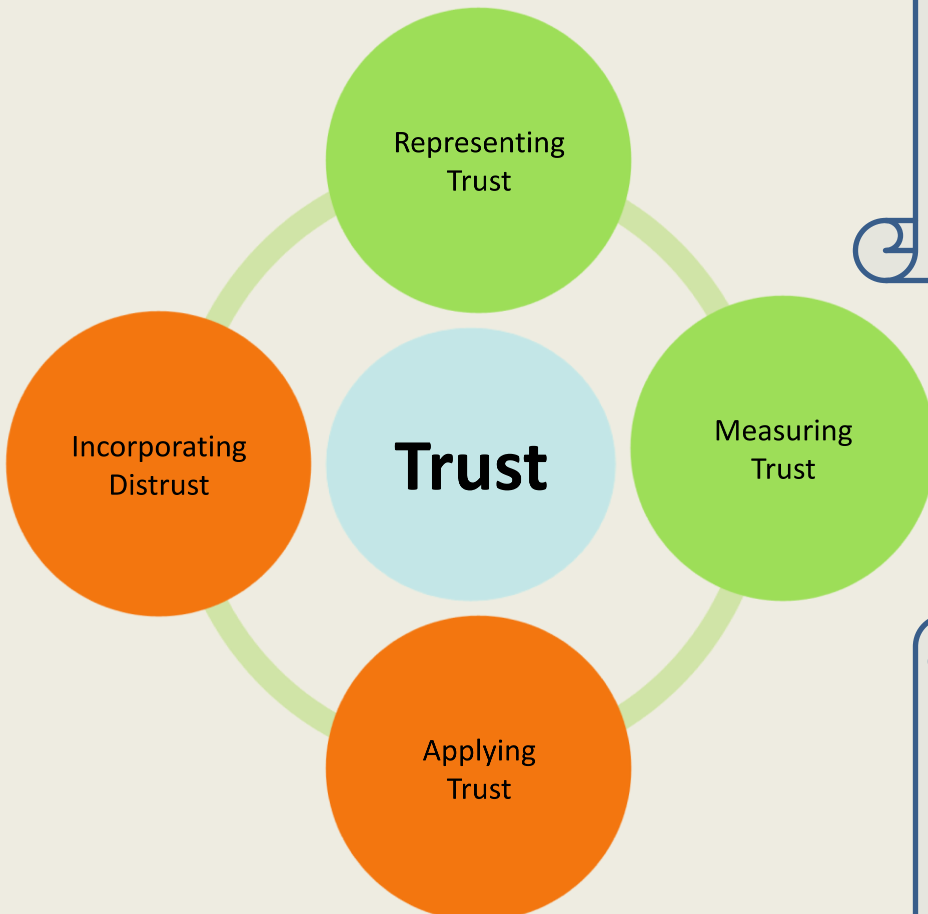

- Properties of Trust
- Classifications of Trust Representations
- Probabilistic vs Gradual Trust Representations
- Single vs Multi-dimensional Trust Representations
- Trust vs Trust and Distrust Representations

- Classifications of Trust Metrics
- Global and Local Trust Metrics
- Supervised and Unsupervised Trust Metrics
- Binary and Continuous Trust Metrics
- Evaluation

- Properties of Trust
- Classifications of Trust Representations
- Probabilistic vs Gradual Trust Representations
- Single vs Multi-dimensional Trust Representations
- Trust vs Trust and Distrust Representations

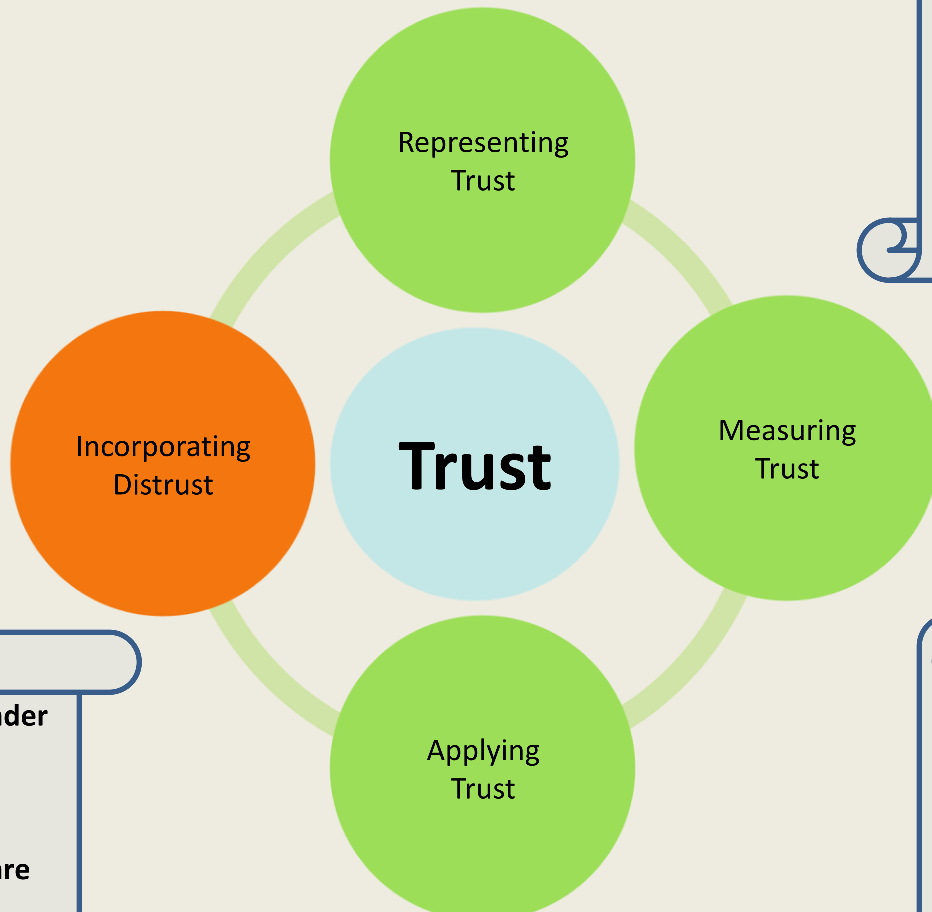

- Trust-aware Recommender Systems
- Opportunities and Challenges
- Model-based Trust-aware Recommender Systems
- Memory-based Trust-aware Recommender Systems
- Evaluation

- Classifications of Trust Metrics
- Global and Local Trust Metrics
- Supervised and Unsupervised Trust Metrics
- Binary and Continuous Trust Metrics
- Evaluation

- Trust in Social Sciences
- Computational Understanding of Distrust
- Distrust in Trust Representations
- Distrust in Trust Measurements
- Distrust in Trust-aware Recommender Systems

- Properties of Trust
- Classifications of Trust Representations
- Probabilistic vs Gradual Trust Representations
- Single vs Multi-dimensional Trust Representations
- Trust vs Trust and Distrust Representations

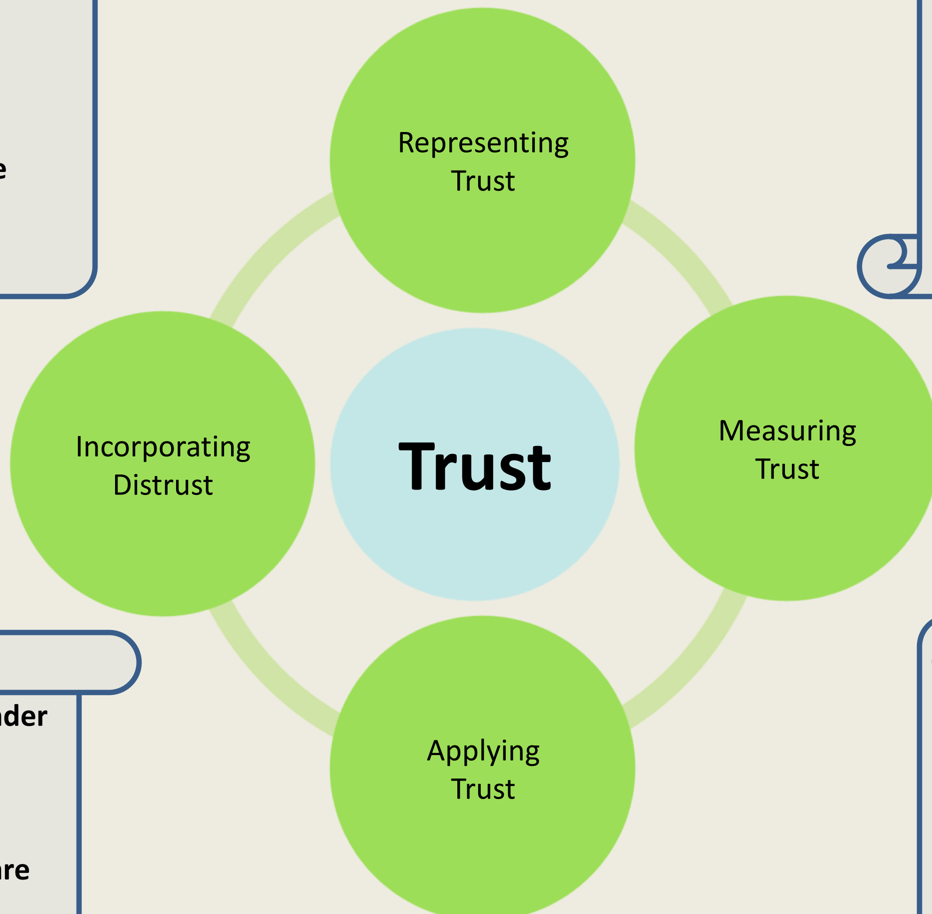

- Trust-aware Recommender Systems
- Opportunities and Challenges
- Model-based Trust-aware Recommender Systems
- Memory-based Trust-aware Recommender Systems
- Evaluation

- Classifications of Trust Metrics
- Global and Local Trust Metrics
- Supervised and Unsupervised Trust Metrics
- Binary and Continuous Trust Metrics
- Evaluation

Research Directions in Measuring Trust

- Solving data sparsity problem in trust measurements
  – Power-law-like distributions suggest that most of users only have few trust relations
  – Most of existing algorithms might fail for users with few trust relations
  – Integrating multiple sources such as item and helpfulness ratings in Epinions might help
  – Incorporating social theories such as homophily and influence

- Measuring trust when trust is not explicitly available
  – There are no explicit trust relations in some social media websites such as Twitter
  – Large amounts of user generated content and user interactions are available
  – Can we measure implicit trust with available data?

Research Directions in Trust-aware Recommendation

- Capturing multi-faceted trust and trust evolution

– Trust relations under different contexts may differ

– Trust relations may change over time

- Incorporating cross-media data

– Trust relations on one site are often sparse and incomplete

– Integrating data from multiple sites may provide a comprehensive and complete view about users

- Exploiting weak trust relations

– Weak tie theory suggests that it is more often that novel information flows through weak rather than strong ties

Research Directions in Incorporating Distrust

- Predicting distrust with publicly available data
  - Trust is a desired property while distrust is an unwanted one for most social networking sites or services
  - Various online services implement trust mechanisms
  - Few of them allow online users to specify distrust relations
  - Distrust relations are usually not available publicly
  - Interaction data is, however, pervasively available, which might indicate distrust relations

- Exploiting distrust in trust applications
  - Using distrust for recommendation is still an open problem
  - Effects of distrust in information propagation
  - Distrust in content filtering
  - Distrust in e-commerce

References

[**Massa, 2007**] P. Massa. A survey of trust use and modeling in real online systems. *Trust in Eservices: Technologies, Practices and Challenges*, page 51, 2007.

[**McKnight and Chervany, 2001**] D. H. McKnight and N. L. Chervany. Trust and distrust definitions: One bite at a time. In *Trust in Cyber-societies*, pages 27–54. Springer, 2001.

[**Golbeck, 2005**] J. A. Golbeck. Computing and applying trust in web-based social networks. 2005.

[**Tang et al., 2013a**] J. Tang, H. Gao, X. Hu, and H. Liu. Exploiting homophily effect for trust prediction. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 53–62. ACM, 2013.

[**Tang et al. 2012a**] J. Tang, H. Gao, and H. Liu. mtrust: discerning multi-faceted trust in a connected world. In *Proceedings of the fifth ACM international conference on Web search and data mining, pages 93–102. ACM, 2012.*

[**Ziegler and Lausen, 2005**] C.-N. Ziegler and G. Lausen. Propagation models for trust and distrust in social networks. *Information Systems Frontiers*, 7(4-5):337–358, 2005.

[**Page et al., 1999**] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. 1999.

References

[**Kamvar et al., 2003**] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th international conference on World Wide Web*, pages 640–651. ACM, 2003.

[**Golbeck, 2006**] J. Golbeck. *Generating predictive movie recommendations from trust in social networks*. Springer, 2006.

[**Massa and Avesani, 2005**] P. Massa and P. Avesani. Controversial users demand local trust metrics: An experimental study on epinions. com community. In *Proceedings of the National Conference on artificial Intelligence*, volume 20, page 121. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005.

[**Liu et al., 2008**] H. Liu, E.-P. Lim, H.W. Lauw, M.-T. Le, A. Sun, J. Srivastava, and Y. Kim. Predicting trusts among users of online communities: an epinions case study. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 310–319. ACM, 2008.

[**Nguyen et al., 2009**] V.-A. Nguyen, E.-P. Lim, J. Jiang, and A. Sun. To trust or not to trust? Predicting online trusts using trust antecedent framework. In *Ninth IEEE International Conference on Data Mining*, pages 896–901. IEEE, 2009.

[**Koren, 2008**] Y. Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 426–434. ACM, 2008.

References

[**Tang et al., 2013b**] J. Tang, X. Hu, and H. Liu. Social recommendation: a review. *Social Network Analysis and Mining*, 3(4):1113–1133, 2013.

[**Jamali and Ester, 2009**] M. Jamali and M. Ester. Trustwalker: a random walk model for combining trust-based and item-based recommendation. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 397–406. ACM, 2009.

[**Ma et al., 2008**] H. Ma, H. Yang, M. R. Lyu, and I. King. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings of the 17th ACM conference on Information and knowledge management*, pages 931–940. ACM, 2008.

[**Tang et al., 2013c**] J. Tang, X. Hu, H. Gao, and H. Liu. Exploiting local and global social context for recommendation. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 2712–2718. AAAI Press, 2013.

[**Ma et al., 2009a**] H. Ma, I. King, and M. R. Lyu. Learning to recommend with social trust ensemble. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 203–210. ACM, 2009.

[**Ma et al., 2011**] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King. Recommender systems with social regularization. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 287–296. ACM, 2011.

References

[**Marsh and Dibben, 2005**] S. Marsh and M. R. Dibben. Trust, untrust, distrust and mistrust–an exploration of the dark (er) side. In *Trust management*, pages 17–33. Springer, 2005.

[**Tang and Liu, 2014**] J. Tang and H. Liu. Is distrust the negation of trust? the value of distrust in social media. Technical Report, Arizona State University, 2014.

[**Guha et al., 2004**] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM, 2004.

[**Jøsang et al., 2003**] A. Jøsang, E. Gray, and M. Kinateder. Analysing topologies of transitive trust. In *Proceedings of the First International Workshop on Formal Aspects in Security and Trust*, pages 9–22, 2003.

[**Heider, 1946**] F. Heider. Attitudes and cognitive organization. *The Journal of psychology, 21(1):107*-112, 1946.

[**Leskovec et al., 2010a**] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Signed networks in social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pages* 1361–1370. ACM, 2010.

[**Leskovec et al., 2010b**] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web, 2010.*

References

[**Victor et al., 2009**] P. Victor, C. Cornelis, M. De Cock, and P. Pinheiro da Silva. Gradual trust and distrust in recommender systems. *Fuzzy Sets and Systems, 160(10):1367–1382, 2009.*

[**Ma et al., 2009b**] H. Ma, M. R. Lyu, and I. King. Learning to recommend with trust and distrust relationships. In *Proceedings of the third ACM conference on Recommender systems,* pages 189–196. ACM, 2009.

[**Jamali and Ester, 2010**] M. Jamali and M. Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 135–142. ACM, 2010.

A book based on this tutorial

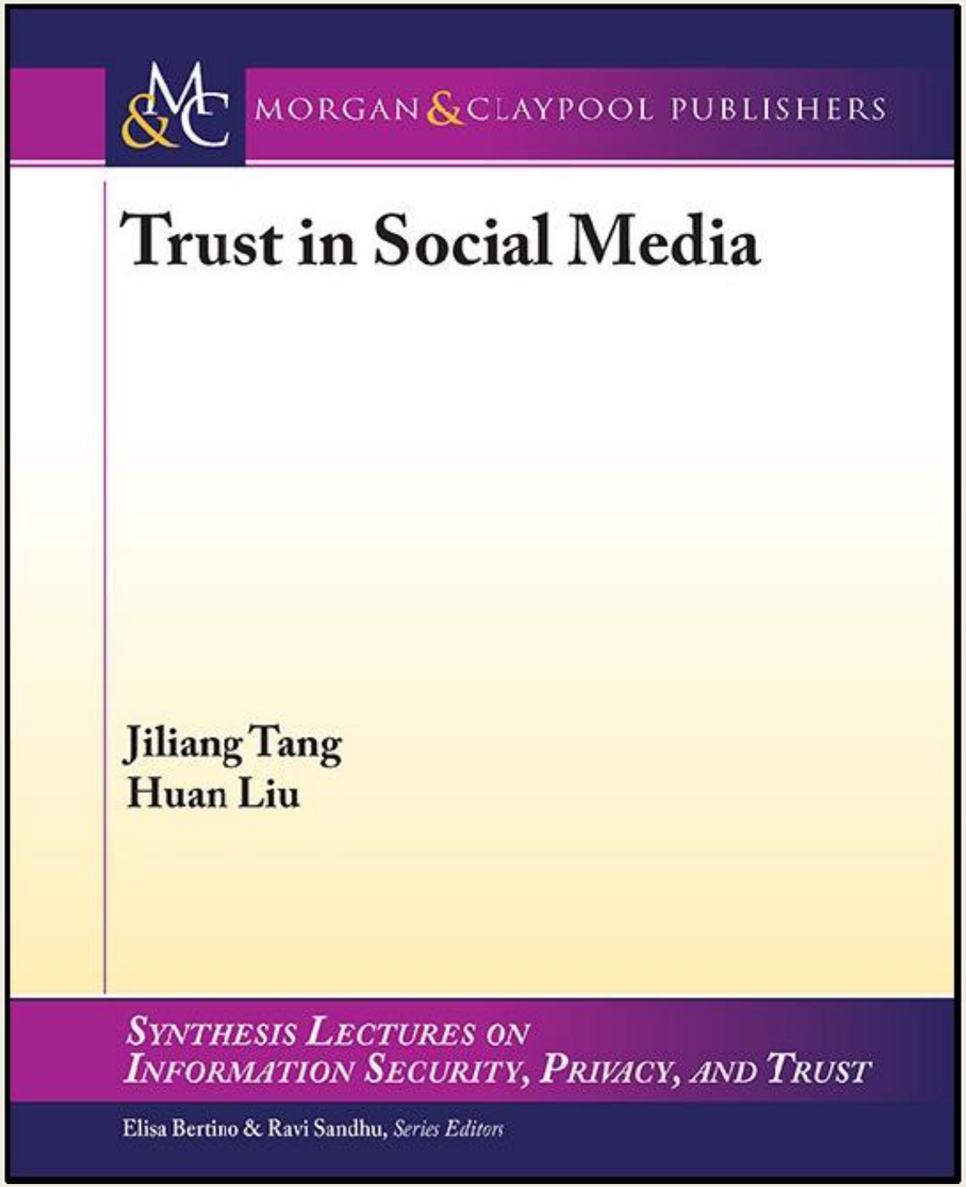

Acknowledgements

- Projects are partially supported by **National Science Foundation, Army Research Office** and **The Office of Naval Research**

Revolutionary Research . . . Relevant Results

- Members of Data Mining and Machine Learning Lab at ASU provided valuable feedback and suggestions