Recommending people in social networks: algorithmic models and network diversity

Javier Sanz-Cruzado

IR Seminar @ University of Glasgow

21st February 2022

In collaboration with

Pablo Castells
Universidad Autónoma de Madrid

Sofía M. Pepa
GMV

Iadh Ounis
University of Glasgow

Craig Macdonald
University of Glasgow

Online social networks

Online social networks

- Establish new connections
- Communication
- Share and receive information
- Changes to our society
  - Politics
  - Privacy
  - Lifestyle
  - Communication

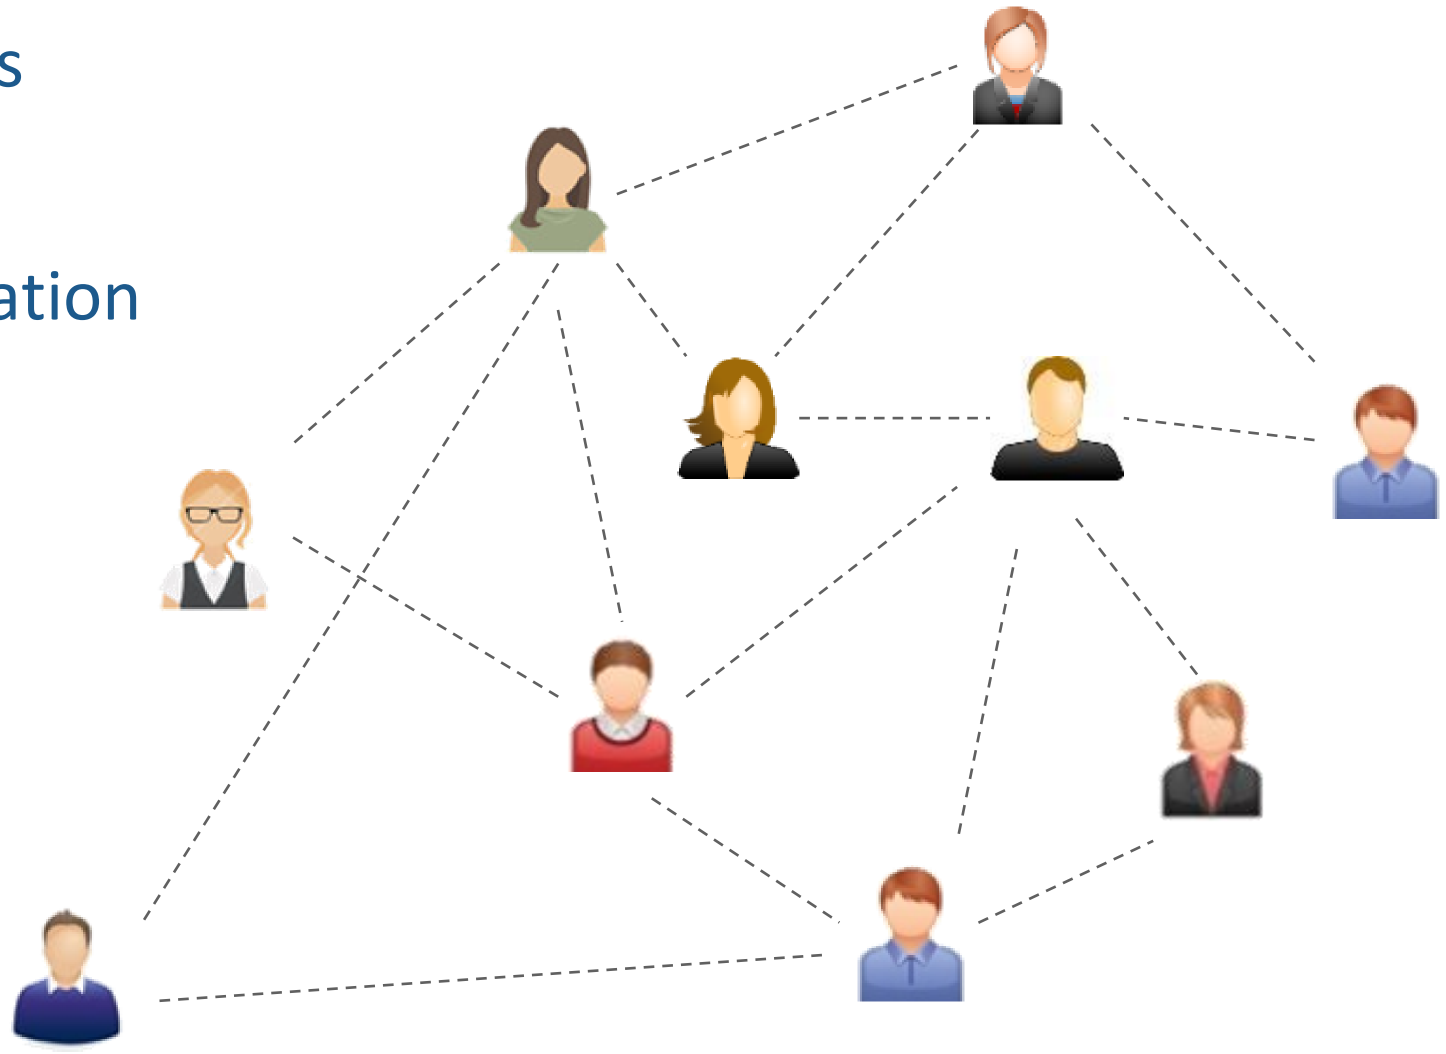

Recommender systems

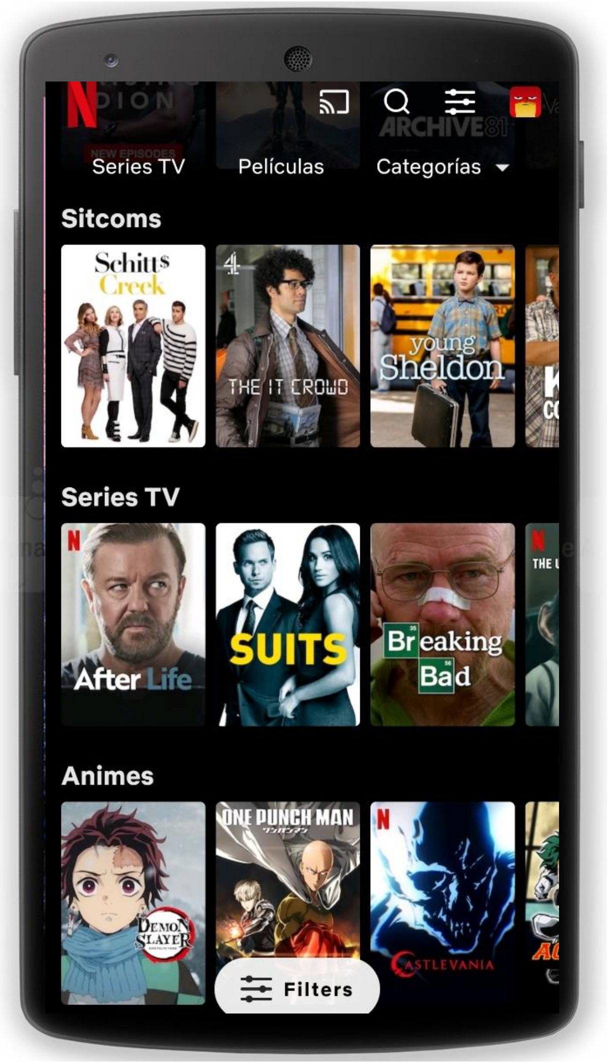

- **Goal:** From past user interactions, suggest items they might be interested in.

- **Multiple domains**
  - Audiovisual content: Netflix, Spotify
  - E-commerce: Amazon, eBay
  - Academic publications: Google Scholar, Mendeley
  - **Social networks: Twitter, Facebook**

The recommendation task

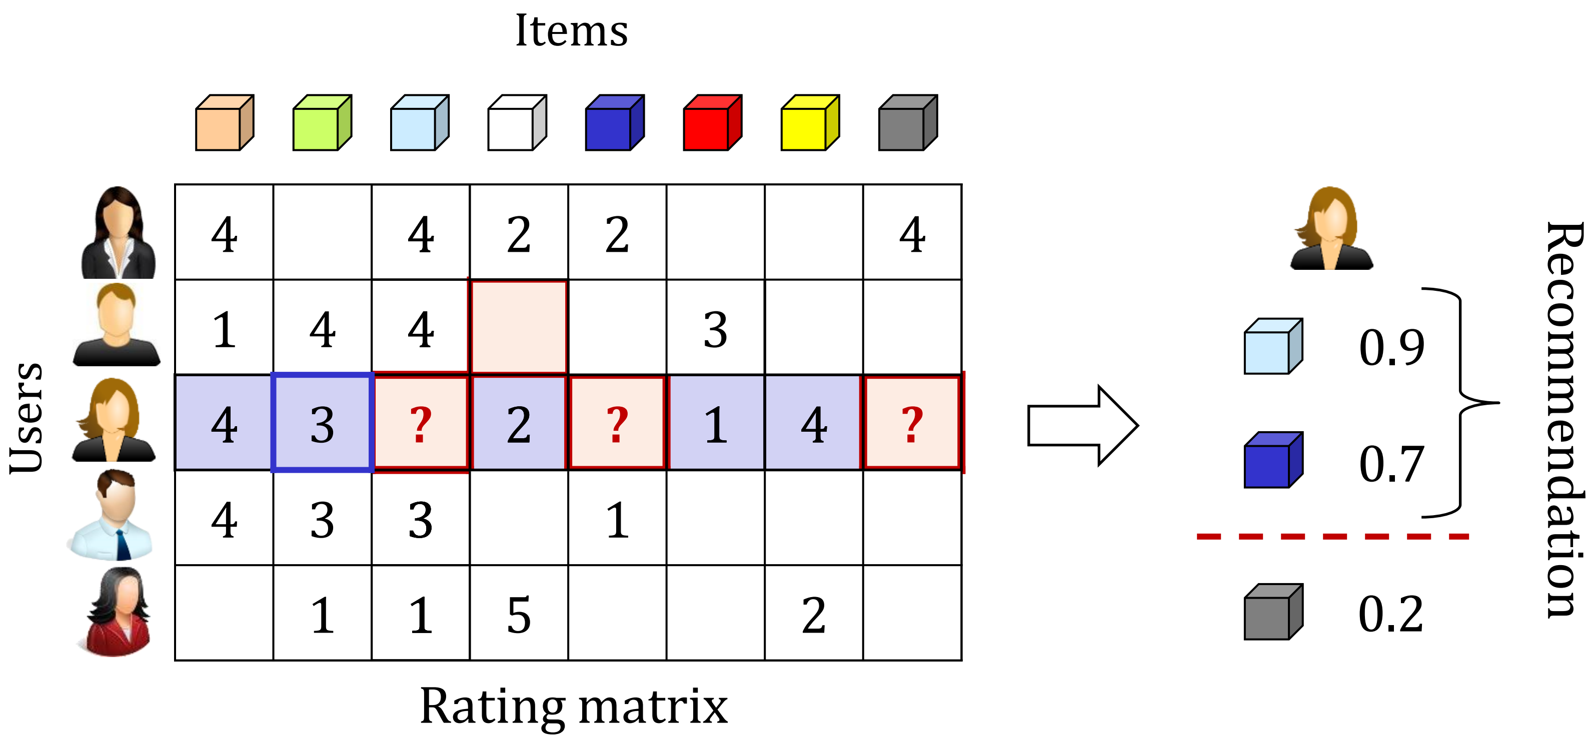

Contact recommendation

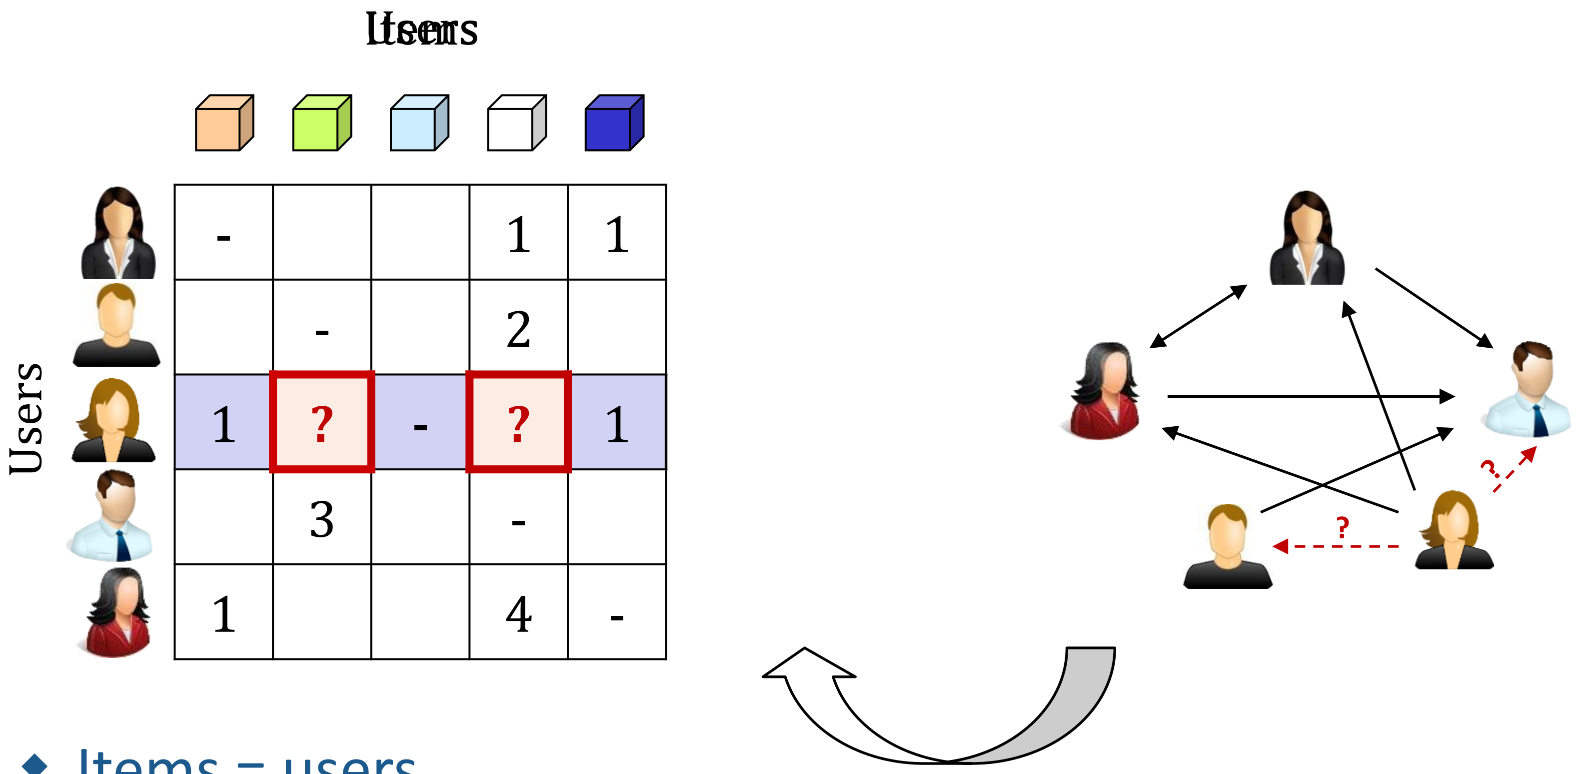

- Items = users
- Availability of social relationships
- Rating matrix = adjacency matrix

(Guy 2015, Sanz-Cruzado & Castells 2018)

Contact recommendation examples

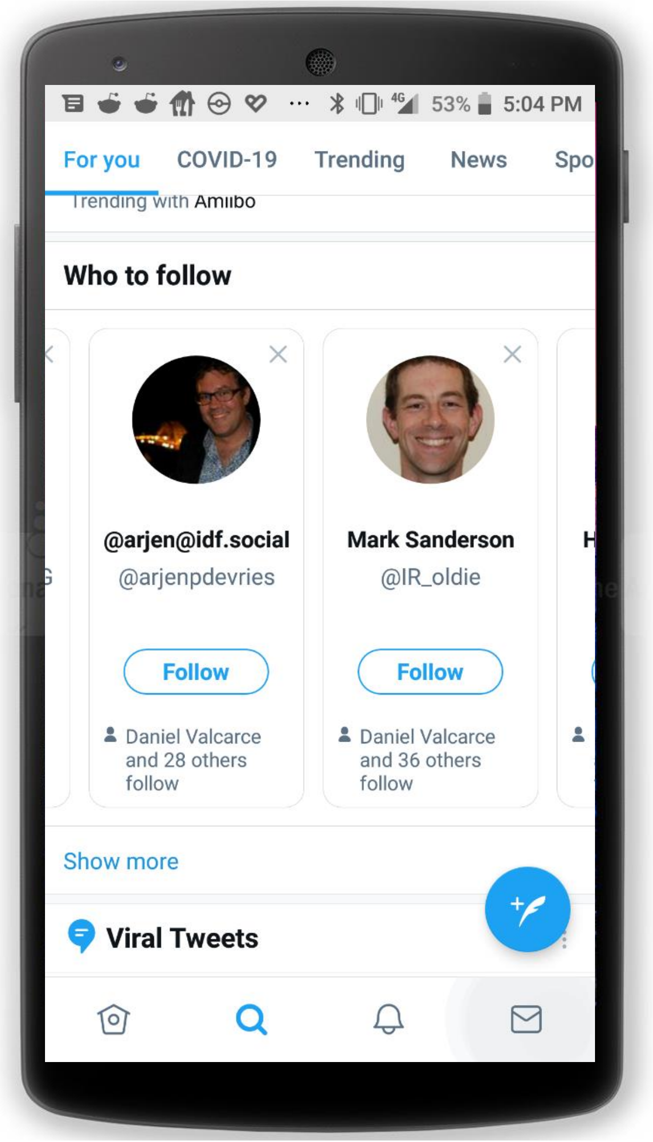

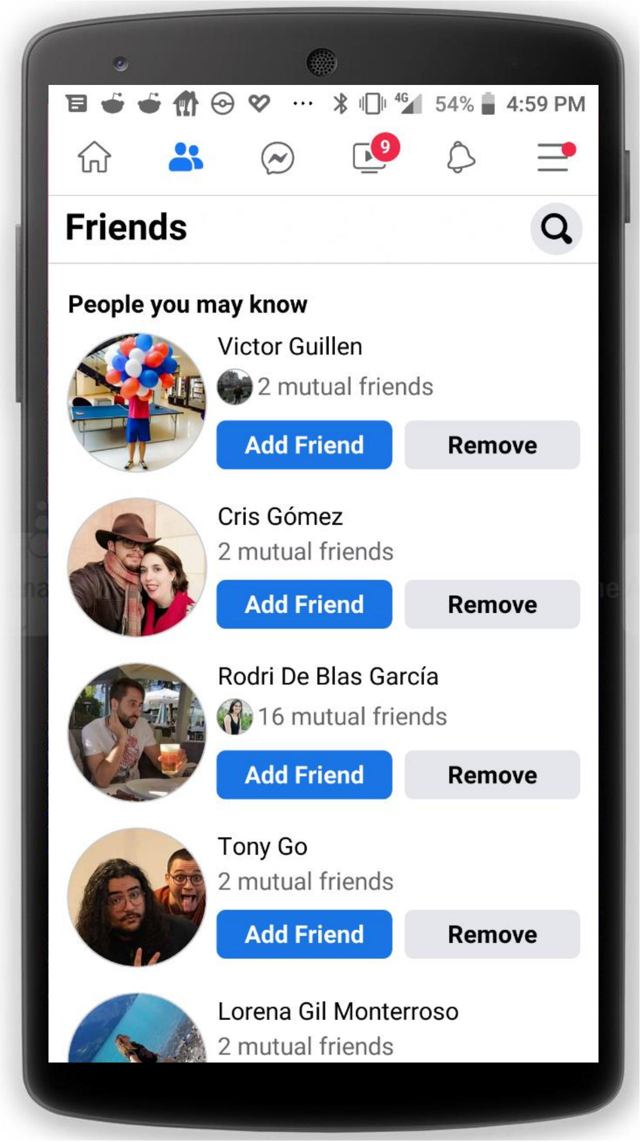

Why contact recommendation?

- Particular characteristics
  - Development of new methods
  - Use of social network analysis

- Creation of new links
  - Main asset of online social networks
  - Communication channels
  - Source of information
  - Increase engagement of users

This presentation

- Part I: Algorithmic models
  - Explore the adaptations of text information retrieval (IR) models to the contact recommendation task.
  - **Publications:** ECIR 2019, ECIR 2020, IP&M 2020
- Part II: Network diversity
  - Study the effect of contact recommendations on the properties of social networks.
  - **Publications:** MSM@WWW 2018, SoMePeaS@ECIR 2019, RecSys 2018

Part I

Algorithmic models

Motivations

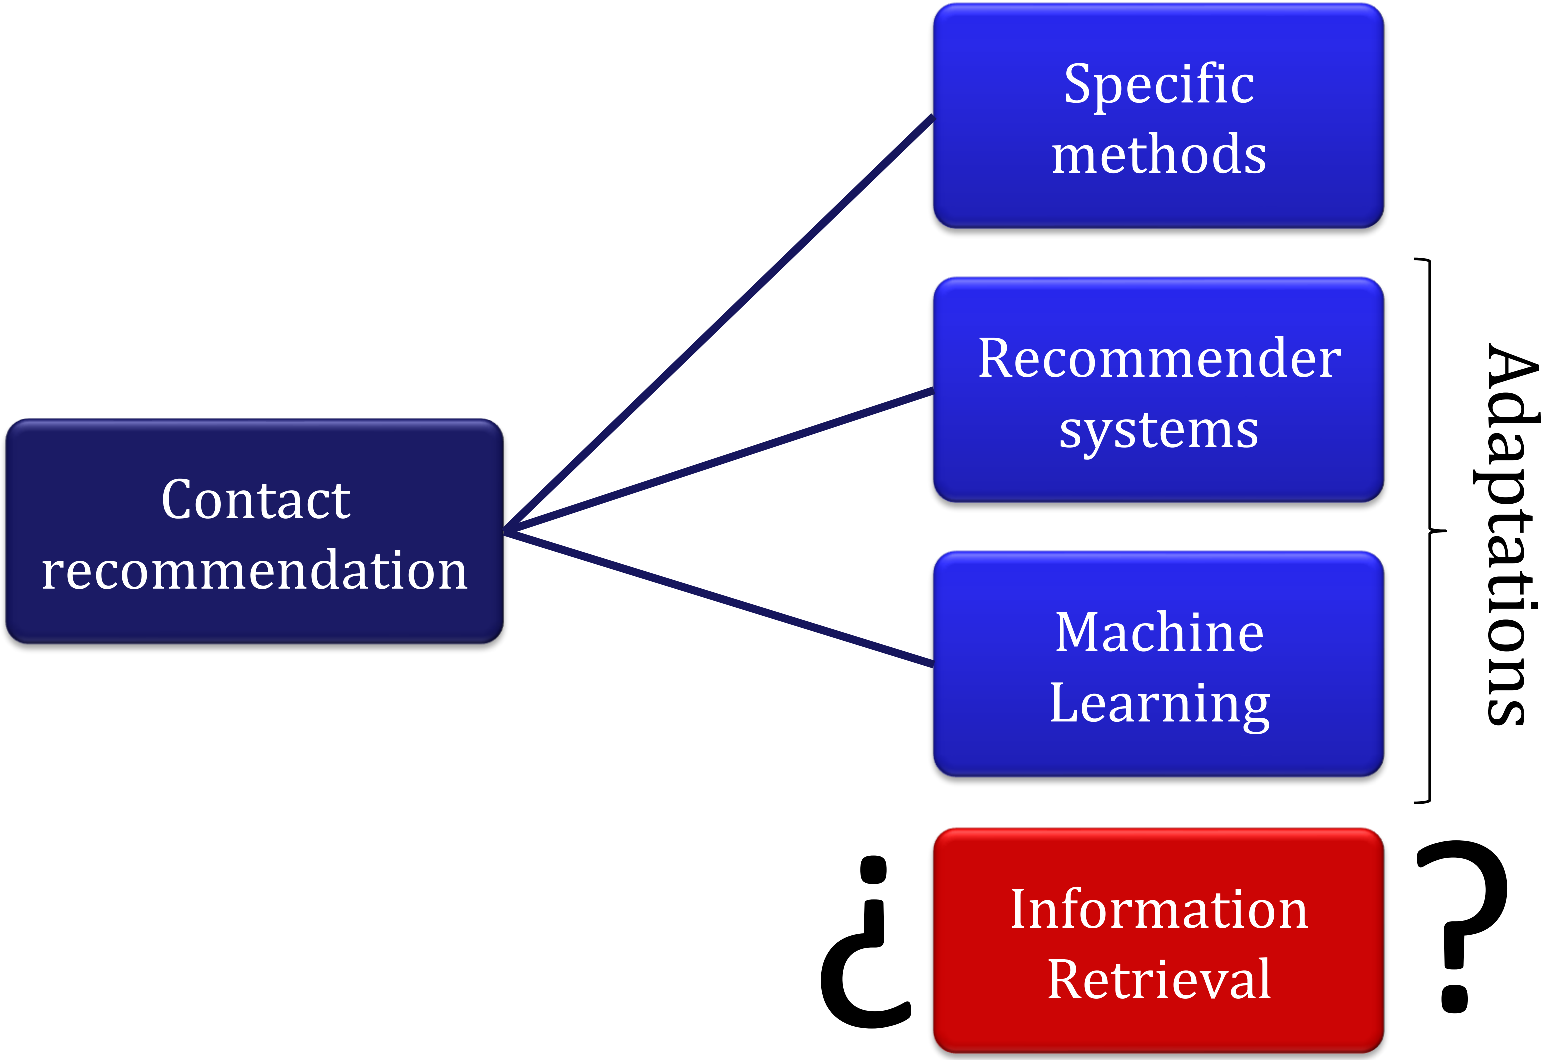

(Sanz-Cruzado & Castells 2018)

IR vs. Recommendation

Text IR

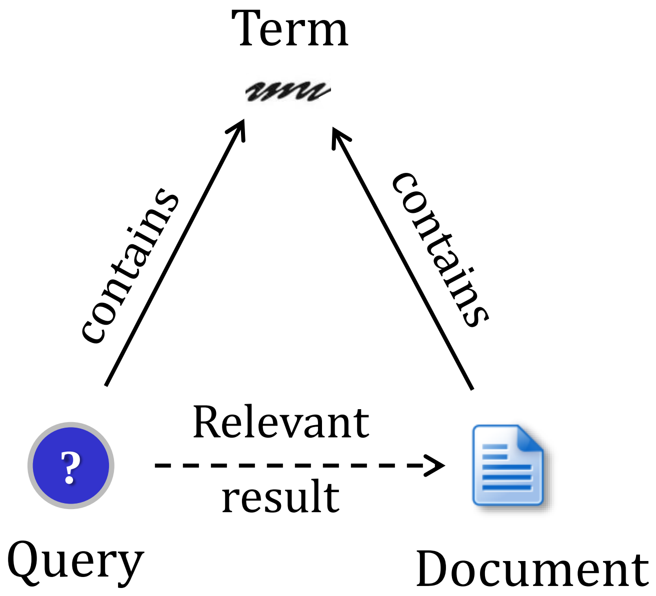

Content-based Recommendation

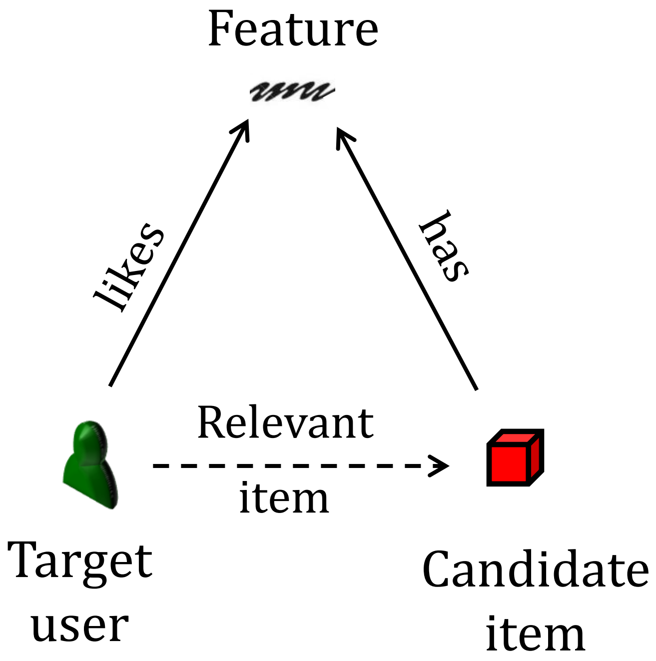

(Adomavicius & Tuzhilin 2005)

Collaborative filtering

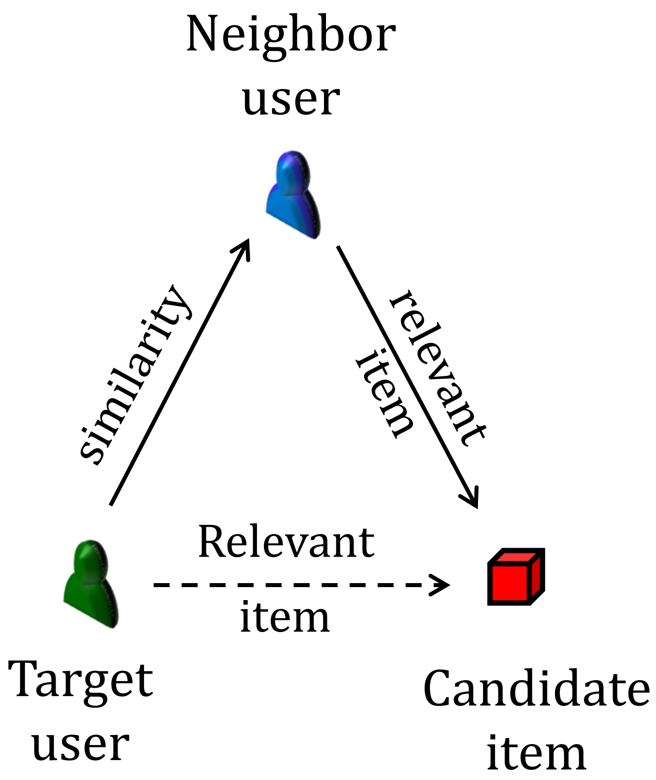

(Bellogín et al., Parapar et al. 2013, Wang et al. 2008, Valcarce et al. 2017)

IR vs. Contact recommendation

Text IR

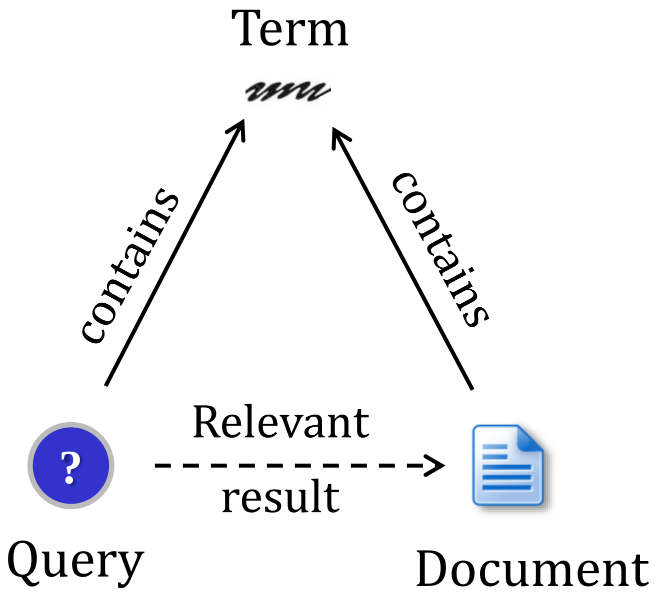

Collaborative filtering

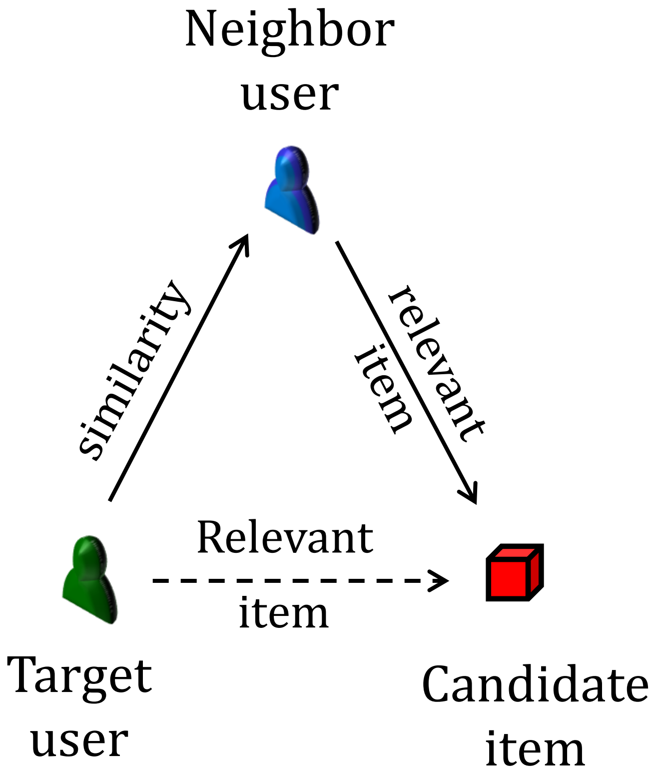

Contact recommendation

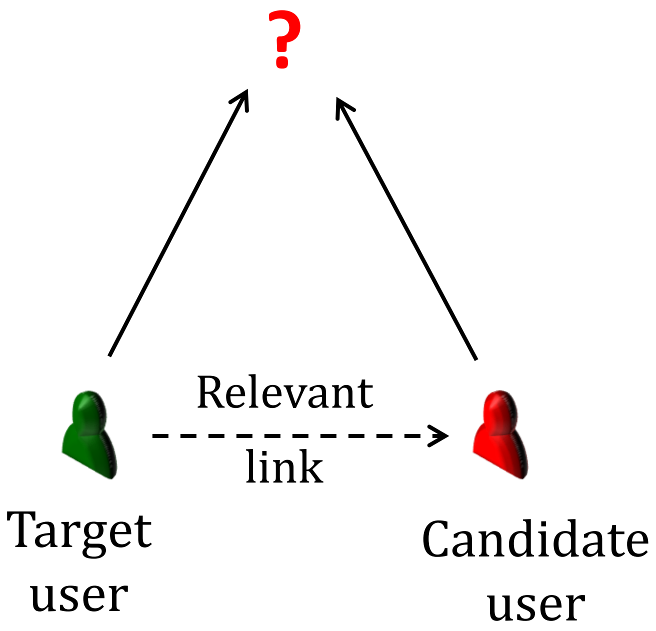

IR vs. Contact recommendation

Text IR

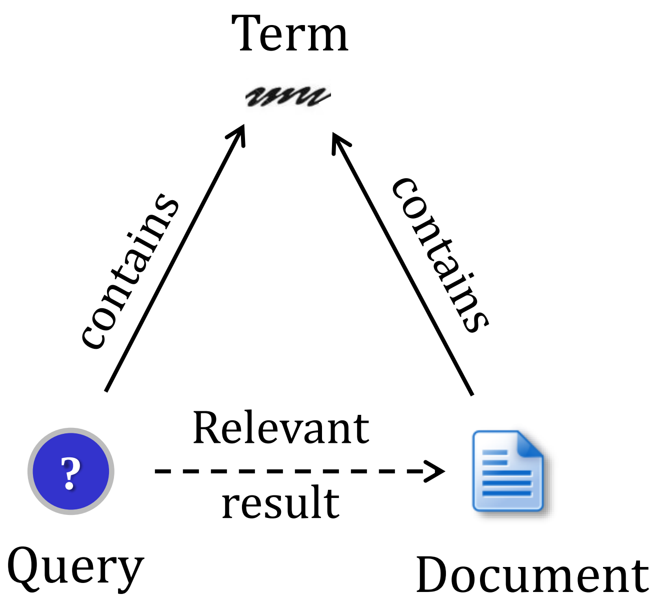

Collaborative filtering

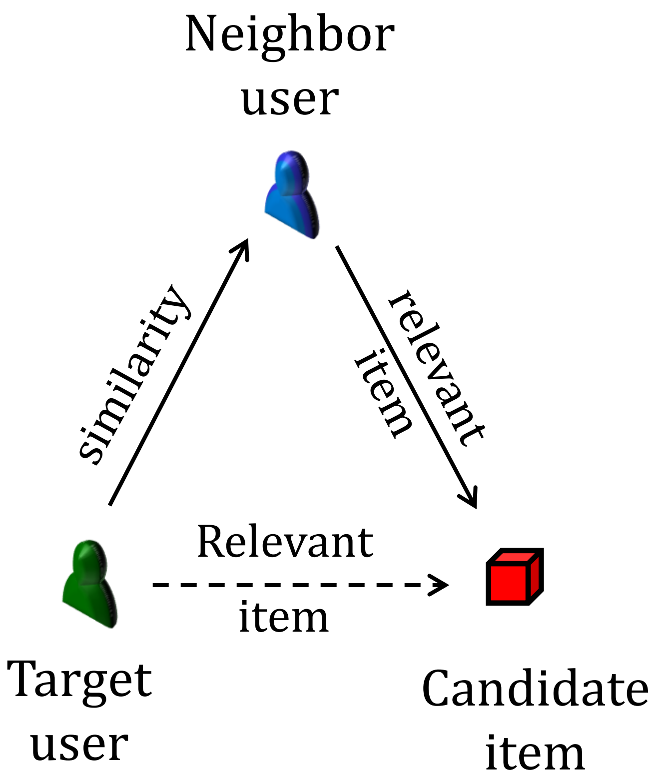

Contact recommendation

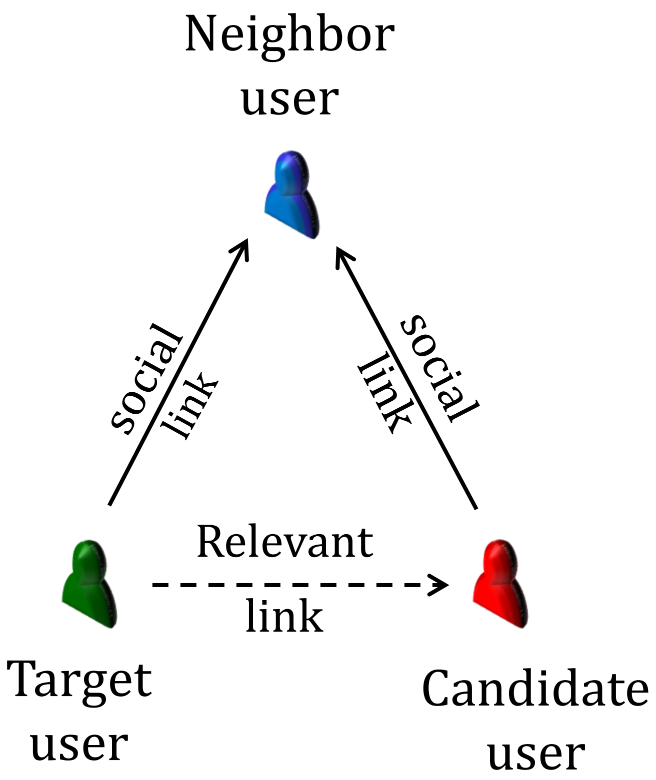

(Hannon et al. 2010)
(Sanz-Cruzado et al. 2020)

An example: BM25

Text IR:

$$f_q(d) = \sum_{t \in d \cap q} \frac{(k+1)\,\mathrm{freq}(t,d)}{k\left(1 - b + \frac{b|d|}{\mathrm{avg}_{d'}|d'|}\right) + \mathrm{freq}(t,d)} \mathrm{RSJ}(t)$$

$$\mathrm{RSJ}(w) = \log \frac{|D| - |D_t| - 0.5}{|D_t| - 0.5}$$

Where

- d: document → $\Gamma(v)$: candidate user
- q: query → $\Gamma(u)$: target user
- $t \in d \cap q$: term → $t \in \Gamma(u) \cap \Gamma(v)$: neighbor user
- D: set of all documents → $\mathcal{U}$: all users
- D_t: documents containing t → $\Gamma(t)$: v containing t in $\Gamma(v)$
- $\mathrm{freq}(t,d)$: frequency of $t \in d$ → $w(t,v)$: edge weight
- $|d|$: document d length → $\mathrm{len}(v) = \sum_{x \in \Gamma(v)} w(x,v)$

An example: BM25

Text IR:

$$f_q(d) = \sum_{t \in d \cap q} \frac{(k+1)\,\mathrm{freq}(t,d)}{k\left(1 - b + \frac{b|d|}{\mathrm{avg}_{d'}|d'|}\right) + \mathrm{freq}(t,d)} \mathrm{RSJ}(t)$$

$$\mathrm{RSJ}(w) = \log \frac{|D| - |D_t| - 0.5}{|D_t| - 0.5}$$

Contact recommendation:

$$f_u(v) = \sum_{t \in \Gamma(u) \cap \Gamma(v)} \frac{(k+1) w(t,v)}{k\left(1 - b + \frac{b \cdot \mathrm{len}(v)}{\mathrm{avg}_{v'}(\mathrm{len}(v'))}\right) + w(t,v)} \mathrm{RSJ}(t)$$

$$\mathrm{RSJ}(t) = \log \frac{|\mathcal{U}| - |\Gamma(t)| + 0.5}{|\Gamma(t)| + 0.5}$$

Experimental setup

- Offline evaluation
- Data from Twitter and Facebook
- **Twitter**
  - Snowball sampling
  - 2 samples
    - **1 month:** All tweets between 19th June and 19th July 2015
    - **200 tweets:** 200 last tweets by each user before 2nd August 2015
  - 2 graphs / dataset
    - Interaction networks: $(u, v) \in E$ if u mentions/retweets v
    - Follow networks
- **Facebook**
  - From Stanford Large Network Dataset Collection
  - Union of 10 ego-networks

Methodology

- Split:

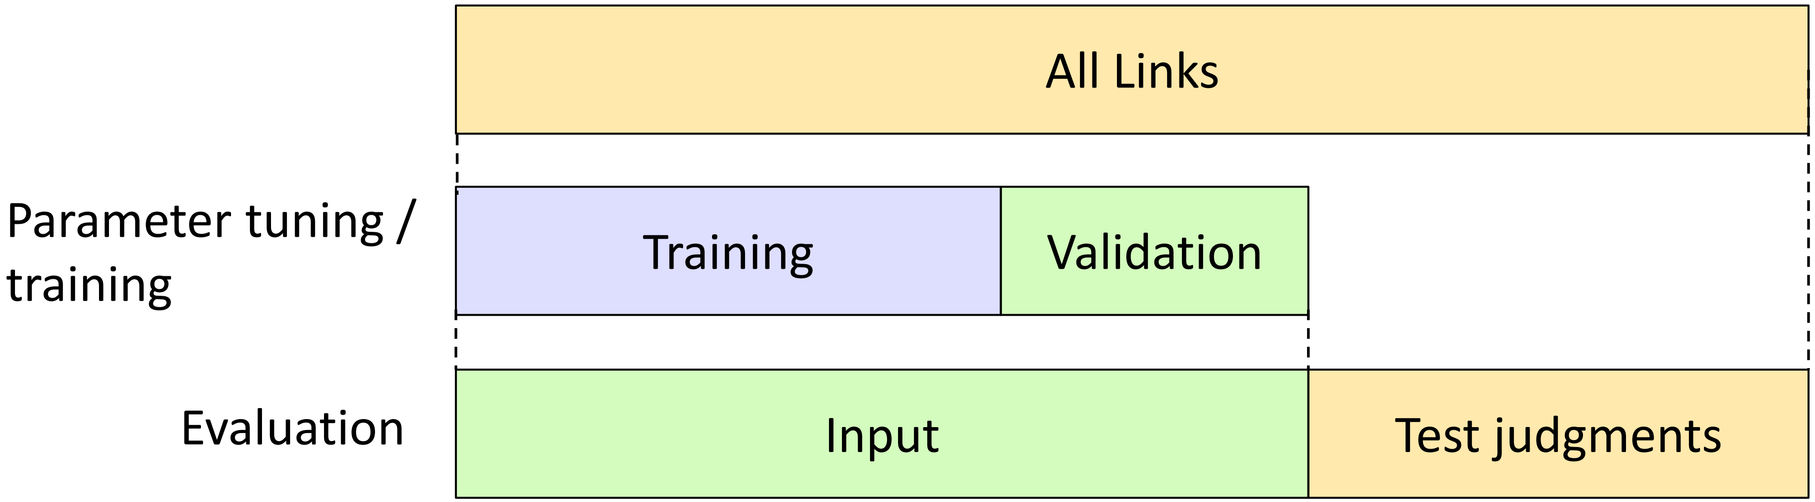

- Hyperparameter selection: grid search (nDCG@10)
- Evaluate using IR metrics on test: nDCG@10, MAP@10

Dataset statistics

	Twitter 1-month		Twitter 200-tweets		Facebook
	Interactions	Follows	Interactions	Follows	
Users	9,528	9,770	9,985	9,964	4,039
Input edges	170,425	645,022	104,866	427,568	56,466
Test edges	54,335	81,110	21,598	98,519	17,643
Directed	✔	✔	✔	✔	✗
Weighted	✔	✗	✔	✗	✗
Split type	Temporal	Temporal	Temporal	Temporal	Random
Density	0.0018	0.0067	0.0013	0.0048	0.0087

Algorithms

- IR models:
  - **Probability ranking principle:** BM25, BIR, ExtremeBM25
  - **Language models:** Query likelihood (QLJM, QLD, QLL)
  - **Divergence from randomness:** PL2, DFRee, DFReeKLIM, DLH, DPH
  - Vector space model (VSM)
- General collaborative filtering
  - User-based / Item-based kNN (cosine similarity)
  - Implicit matrix factorization (iMF)
- Specific approaches
  - **Friends of friends:** Adamic-Adar, MCN, Jaccard, cosine similarity
  - **Random walks:** Personalized PageRank, Money,…
  - **Path-based:** Local Path Index, Katz…
- Sanity check: Random and most popular

Results (nDCG@10)

	200-tweets		Facebook
Algorithm	Interaction	Follows	
BM25	**0.1097**	0.1159	0.5731
BIR	0.1004	0.114	0.572
PL2	0.0983	0.1166	0.5712
VSM	0.0425	0.0787	0.5237
iMF	0.1035	**0.1329**	0.521
User-based kNN	0.0954	0.1297	0.5457
Item-based kNN	0.0724	0.1205	0.4542
Adamic-Adar	0.0997	0.114	0.5746
MCN	0.0948	0.111	0.5585
Resource allocation	0.0913	0.1117	**0.5922**
Personalized PageRank	0.063	0.0843	0.5891
Cosine	0.048	0.0768	0.4943
Popularity	0.0422	0.0397	0.0523
Random	0.0003	0.0018	0.003

- IR models are effective
  - BM25 among top 5
  - **Best:** 200-tweets interactions
  - VSM lowest performing IR model
- Rest of algorithms
  - Implicit MF is best
  - Adamic-Adar and MCN are competitive
  - Jaccard/cosine are not very competitive
  - Rest seem very graph dependent

Can we do better?

	1-month		200-tweets		
Algorithm	Interaction	Follows	Interaction	Follows	Facebook
User-based kNN	0.1367	0.1413	0.0954	0.1297	0.5457
Item-based kNN	0.1174	0.1296	0.0724	0.1205	0.4542
Cosine	0.0393	0.0497	0.0480	0.0768	0.4943

User-based / Item-based kNN (cosine similarity) > Standalone cosine similarity

What if we try the same with IR models?

K Nearest Neighbors

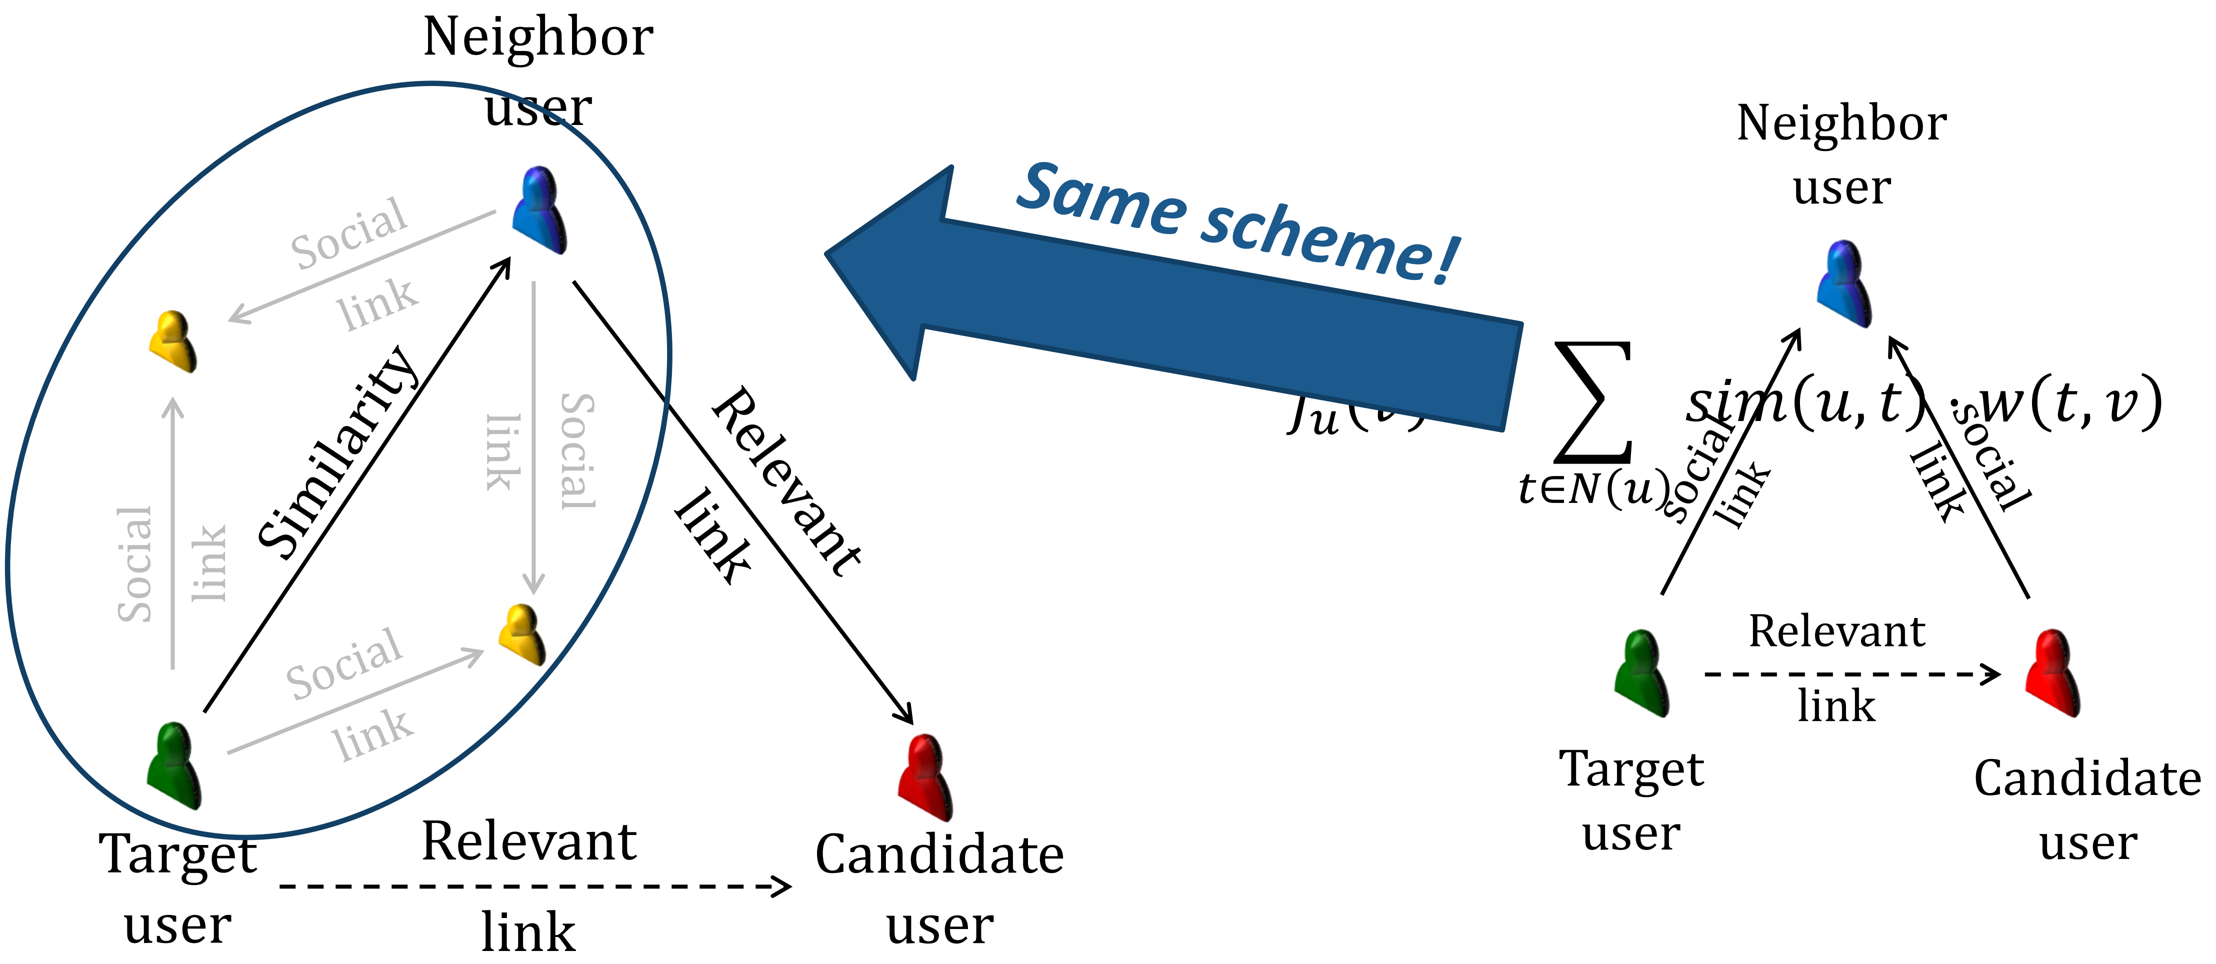

(Ning et al. 2015)

Results kNN + IR (nDCG@10)

Terrier Team | University of Glasgow

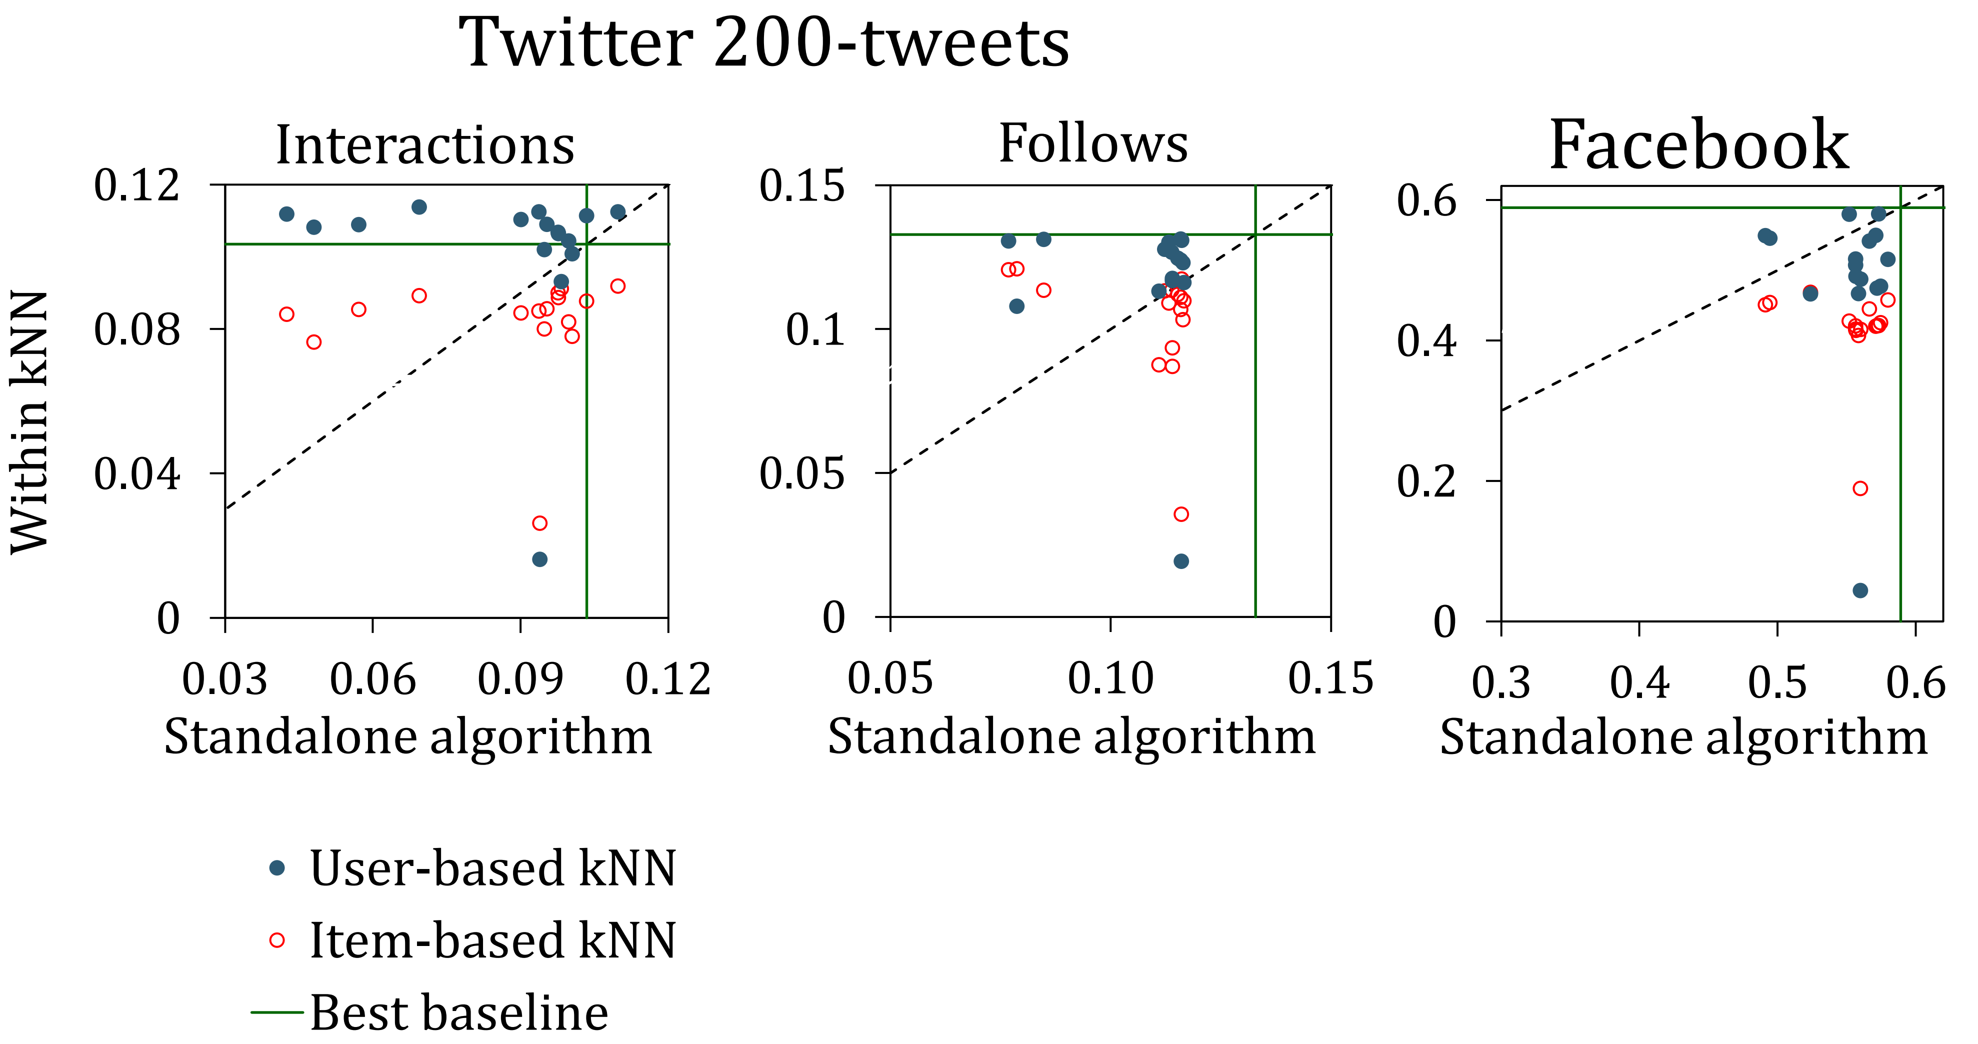

Can we do even better?

- **Idea:** Learning to rank (Liu 2007)
  - Supervised machine learning models
  - Very effective in IR
- How does it work?
  1. Sample candidates
  2. Generate features for each target-candidate user pair
  3. Generate recommendation ranking

Our experiments

- **LETOR algorithm:** LambdaMART (Burges 2010, Ganjissafar et al. 2011)

- **Features:** Scores of contact recommendation methods
  - IR models
  - Friends of friends (FOAF) approaches
  - User-based / Item-based kNN + IR / FOAF

- **Sample suitable candidates:** use IR models

Results (nDCG@10)

LambdaMART improves best recommendation baselines

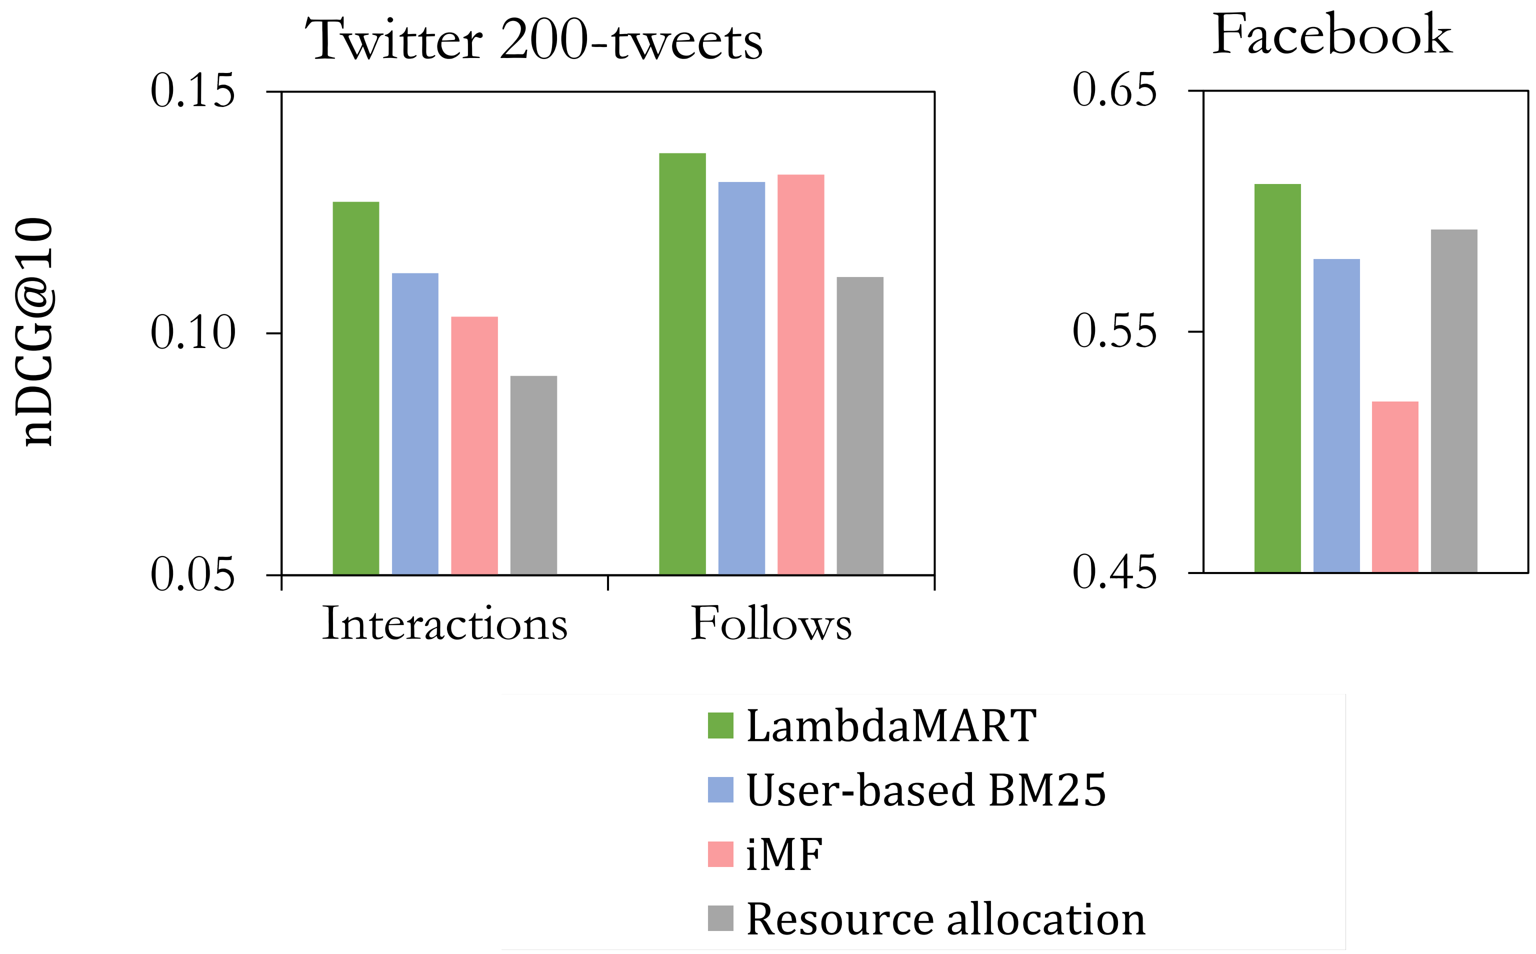

Conclusions

- We can use IR models as contact recommendation algorithms
- Direct IR models are both effective and efficient (BM25)
- IR-based models are better as neighborhood selectors for kNN
- Learning to rank techniques improve the accuracy of best state of the art algorithms
- IR models are effective in three different roles in contact recommendation
  - Direct recommenders
  - Neighborhood selectors in kNN
  - Samplers and features in learning to rank

Part II
Network diversity

Accuracy

- Fundamental goal of contact recommendation
- Increase network density
- Limitations:
  - **Local perspective:** average over isolated users
  - **Narrow perspective:** one-dimensional utility

Beyond accuracy

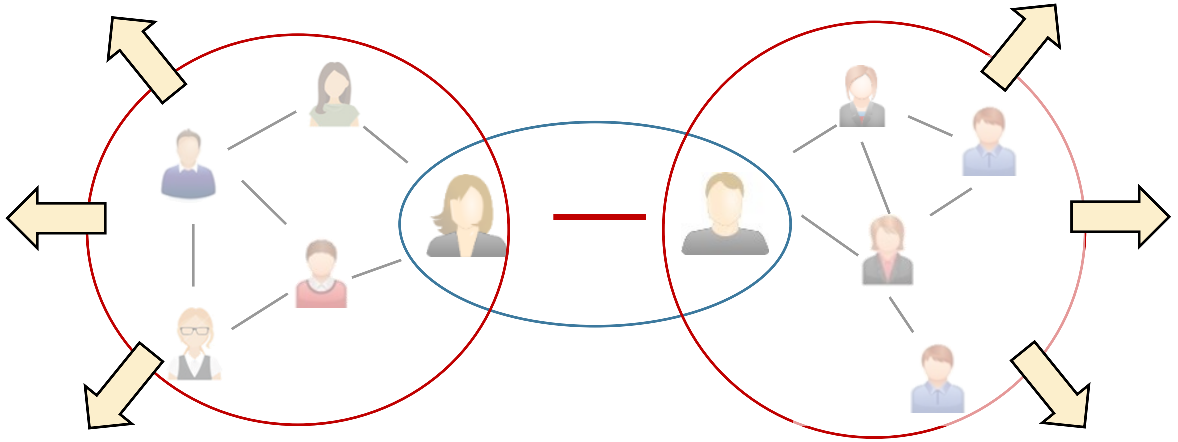

- Users in the network are not isolated
- A few links can cause global effects
- Different links – different effects
- Contact recommendation
  - 500 million new links/month on Twitter (as of 2015)
  - Potential to drive network evolution

Goals

1. Define suitable metrics to measure global benefits of recommendation

2. What do the metrics really mean? Do they capture relevant aspects of network functionality?

Effects on network structure

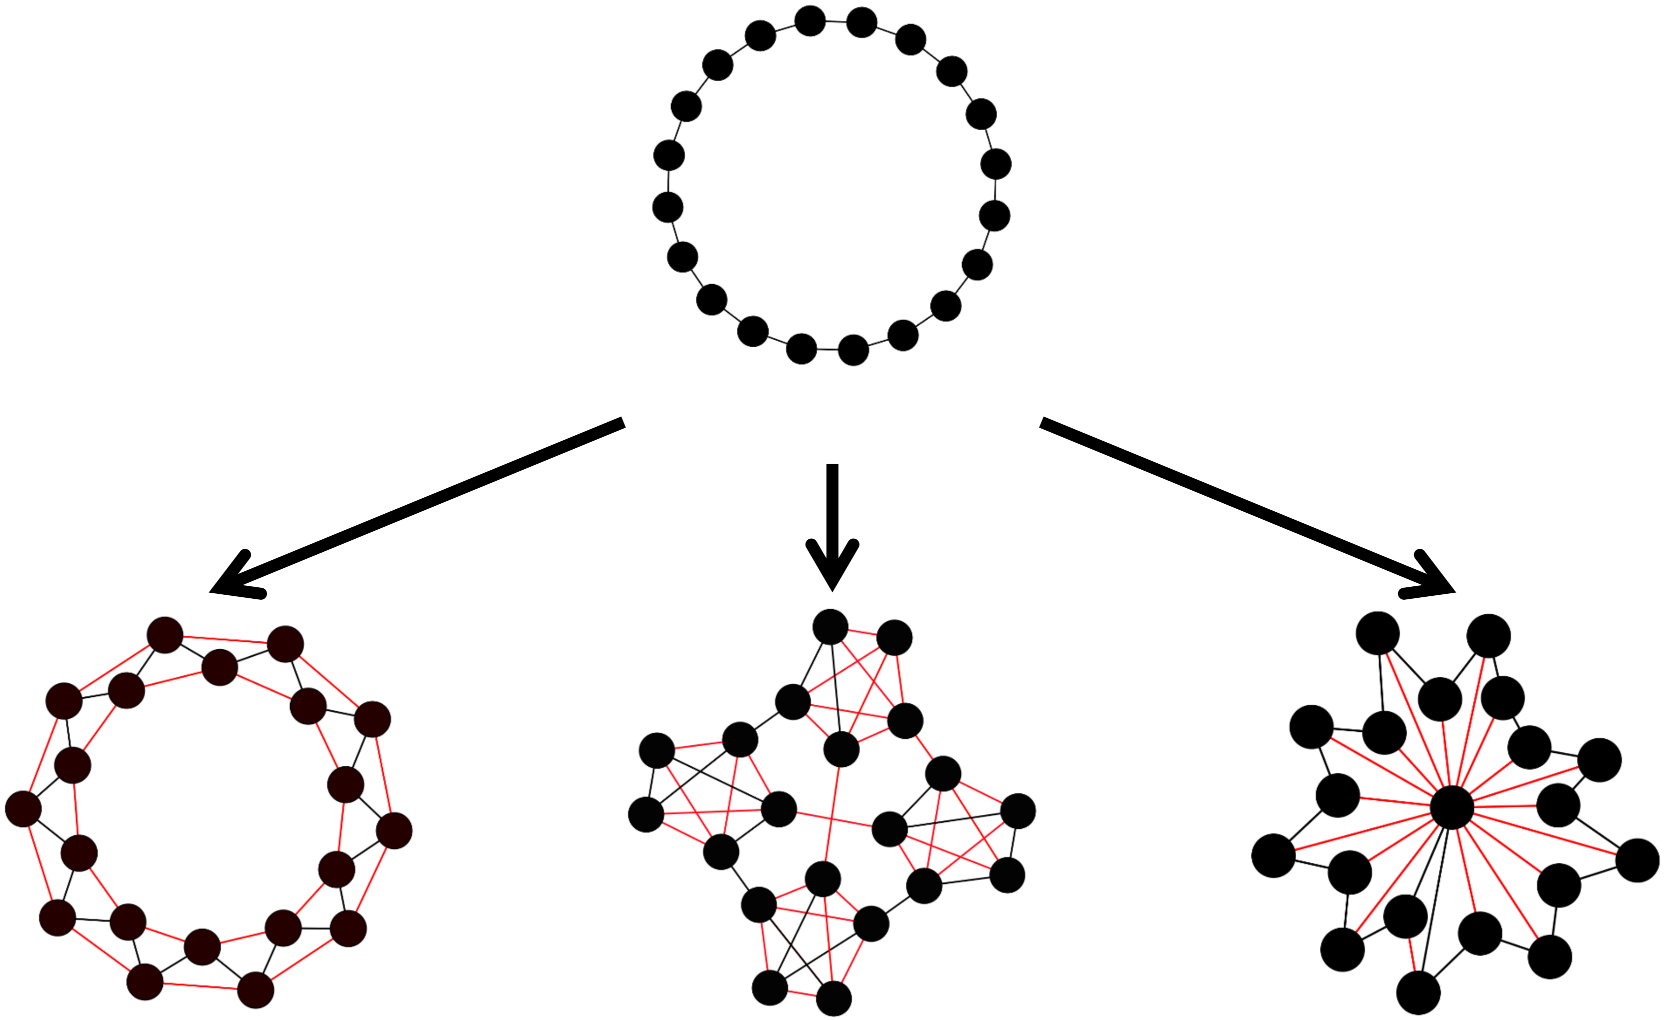

How to measure?

User	Score
u_2	0.9
u_3	0.8
u_4	0.1

Recommendation ranking

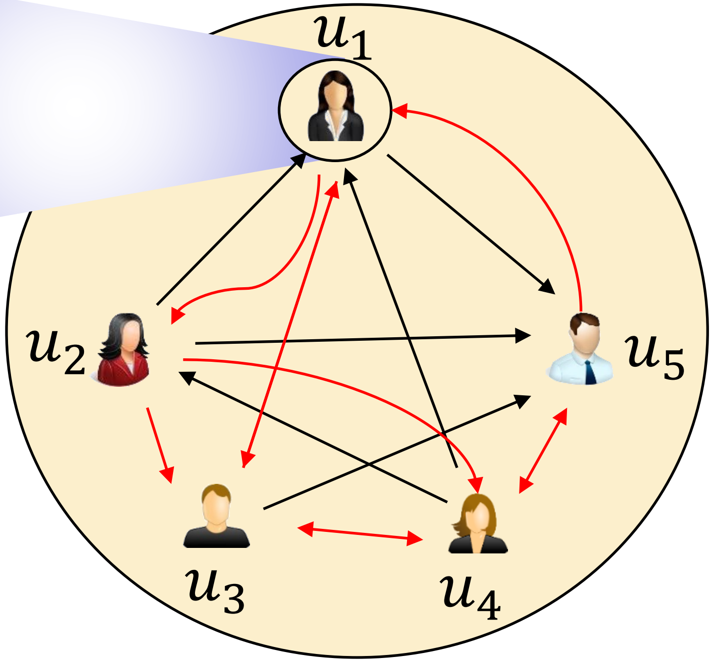

Structural metric

Potentially relevant structural features

- Structural diversity
  - Source of novel information.
  - Enrichment of the information flow.
  - Related to the notion of weak tie (Granovetter, 1978)
- Strength of a tie
  - Measures the involvement of users in the tie.
  - **Strong ties:** family, close friends.
  - **Weak ties:** people you meet in conferences, shopkeepers.
- In the network structure: **non-redundant links**

Weak links: local notions

- Consider the direct environment of the link.

- **Triadic closure:** minimum unit of structural redundancy.

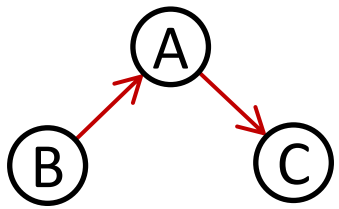

a) Non-redundant triad

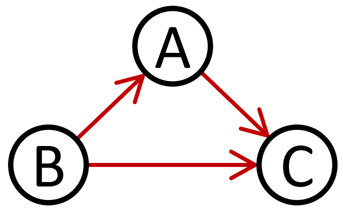

b) Redundant triad

- **Metric:** clustering coefficient complement.

Measures the proportion of non-redundant triads in the network.

Weak ties: global notions

- **Weak ties:** links between communities (De Meo et al. 2012)
  - Tightly connected groups of nodes
  - Few connections outside the group
- **Modularity complement (MC):** number of weak ties

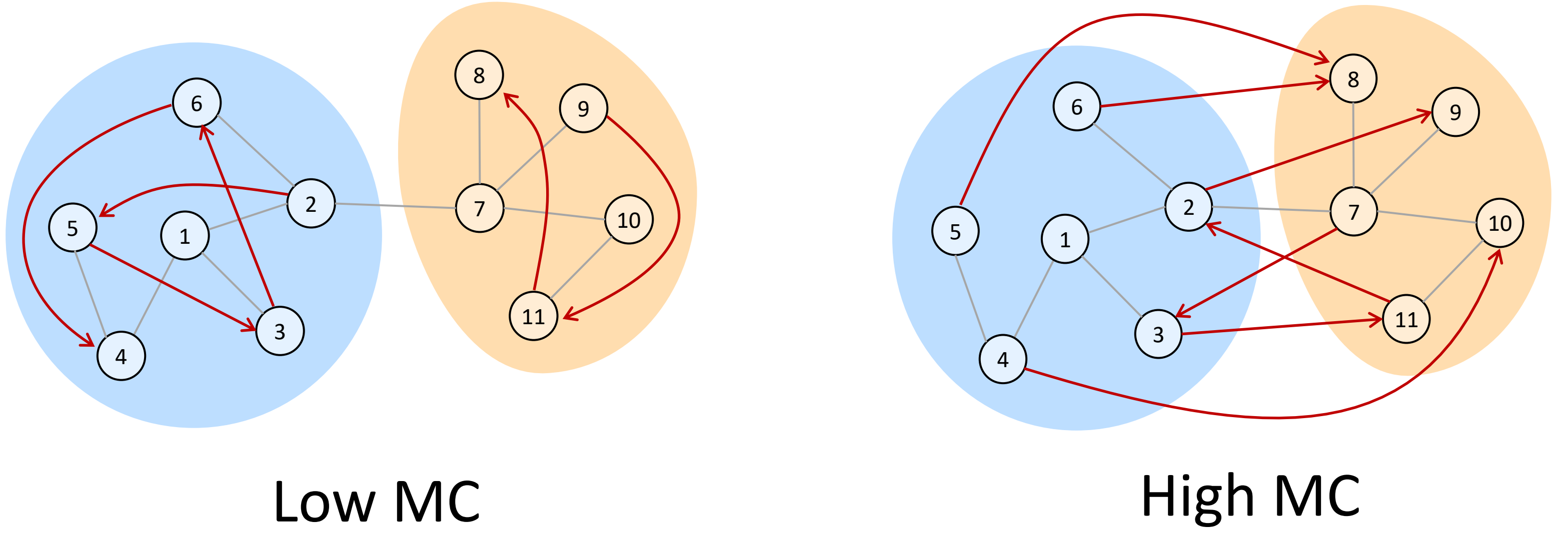

Weak ties: global notions (II)

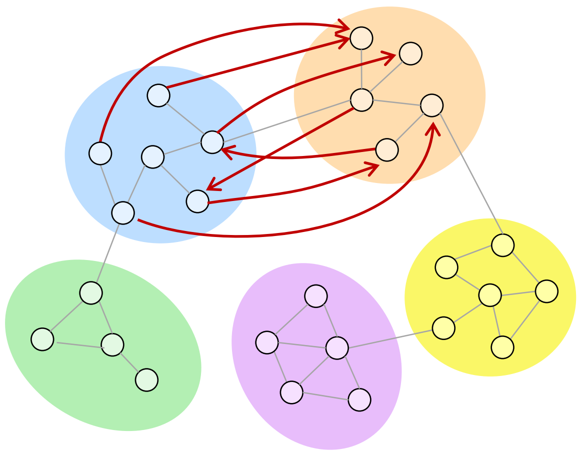

Weak-link redundancy

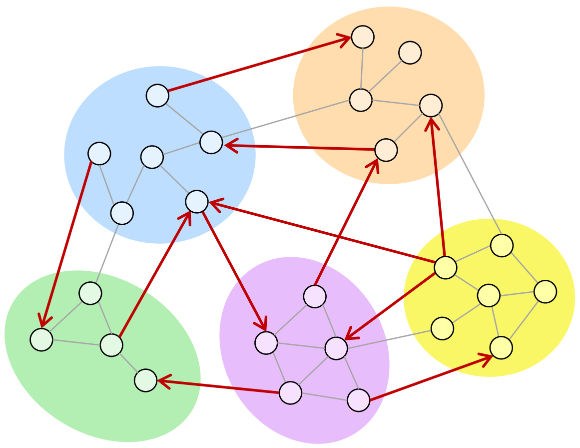

Weak-link diversity

- Community edge Gini complement (CEGC)
  - **New metric**
  - Distribution of weak links between pairs of communities
  - Based on the Gini index

Effect of different recommenders

Algorithm	nDCG@10	Clustering coefficient	Modularity	Community Gini
iMF	**<u>0.139</u>**	0.902	0.155	0.045
BM25	0.104	0.878	0.150	0.041
Adamic-Adar	0.098	0.882	0.149	0.041
MCN	0.092	0.879	0.145	0.040
Pers. PageRank	0.100	0.915	0.182	0.054
Popularity	0.057	0.924	**<u>0.295</u>**	0.061
Random	0.001	**<u>0.952</u>**	0.280	**<u>0.091</u>**
Original network	-	0.9437937	0.1463597	0.0390234

What do these numbers really mean for the network?

Filter bubbles

We analyze the potential of weak ties on reducing filter bubbles

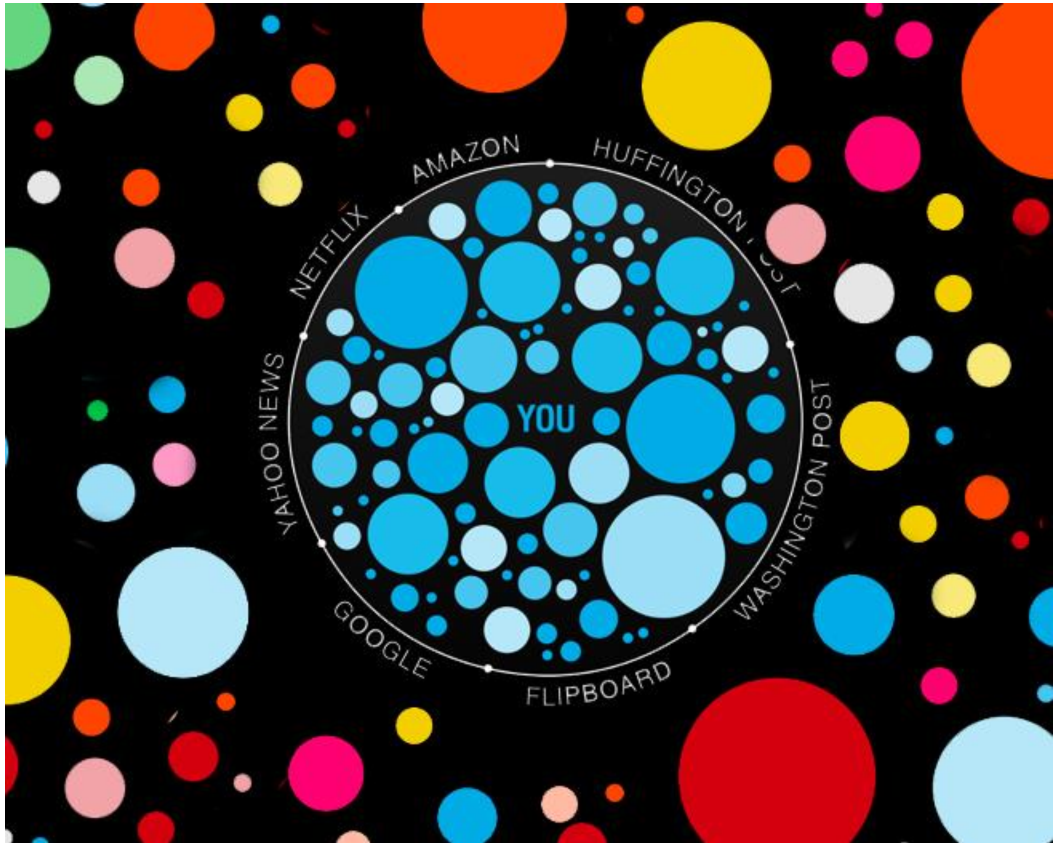

(Pariser 2011)

Diffusion experiment

Terrier Team | University of Glasgow

- Hypothesis

> The more structurally diverse the recommendation is, the more diverse and novel the information flow through the network will be.

- Experiment on interaction networks
  1. Start with a baseline: Implicit MF / BM25
  2. Apply gradual rerankers for optimizing a metric
  3. Extend the network with top k recommended links
  4. Run propagation of (real) tweets through the network
  5. Measure diffusion properties (novelty & diversity)

Simulation

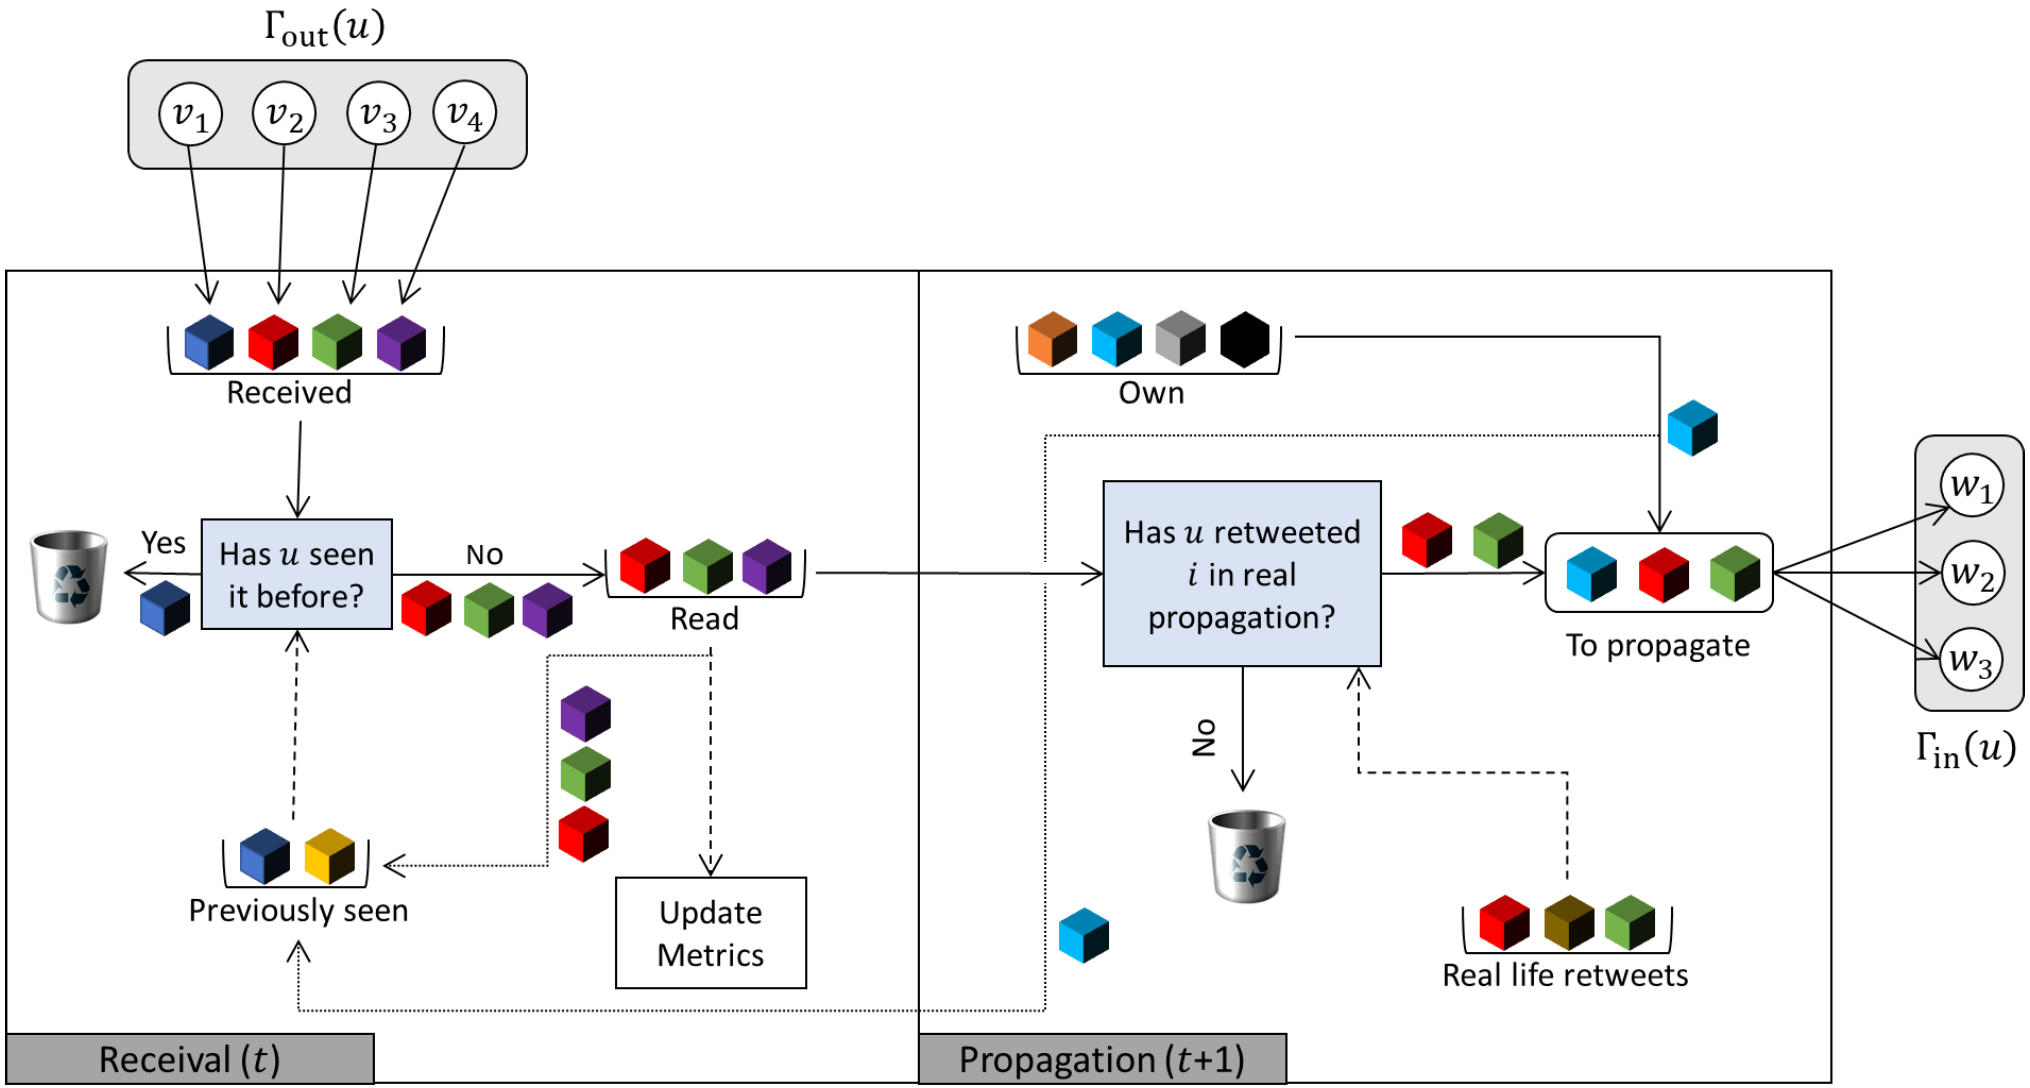

$\Gamma_{out}(u)$
v_1
v_2
v_3
v_4
Received
Yes
Has u seen it before?
No
Read
Previously seen
Update Metrics
Receival (t)
Own
Has u retweeted i in real propagation?
No
To propagate
Real life retweets
Propagation (t+1)
w_1
w_2
w_3
$\Gamma_{in}(u)$

Diffusion properties

- Measured in terms of tweet hashtags (as topics)
- Novelty
  - Proportion of the hashtags unknown to the users.
  - Known hashtags: hashtags in their original tweets.
- Diversity
  - How evenly are hashtags propagated over the population
  - Complement of the Gini index

Results

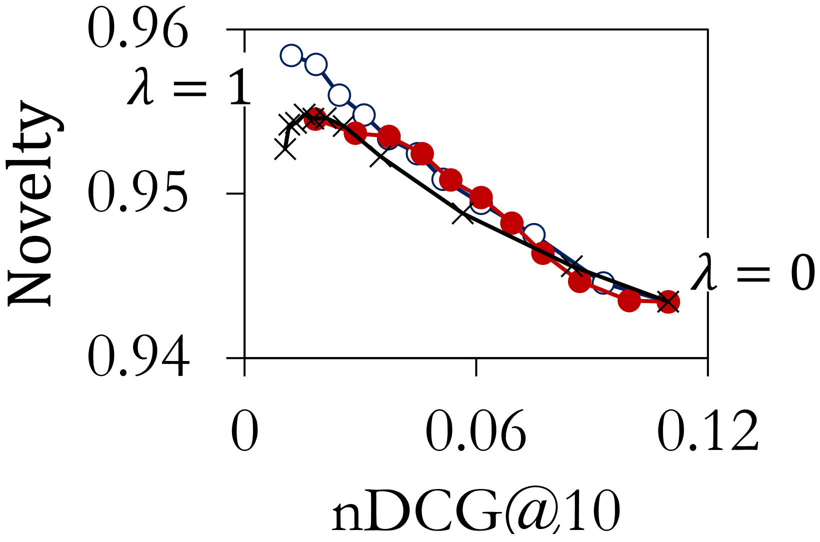

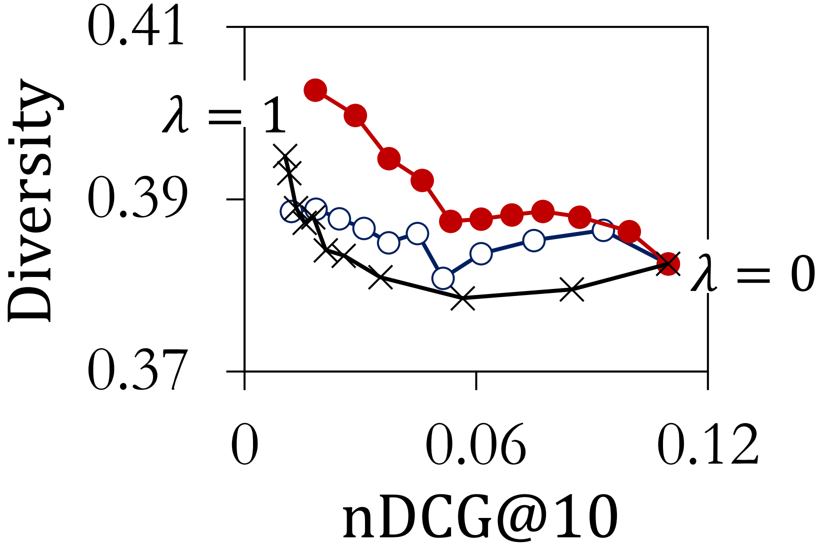

Graph: Twitter 200-tweets interactions

Baseline: BM25

- Community Gini
- Modularity
- Clustering coefficient

Enhancing weak ties has positive effects in the novelty and diversity of the information flow

Conclusions

- Accuracy is a partial perspective
- We propose evaluation perspectives beyond accuracy
  - Global network effects beyond (averaged) isolated user gains
  - New metrics elaborating on weak ties
- Enhancing the number of weak ties improves novelty & diversity of the information arriving to the users

Summary

- We can use IR models as contact recommendation algorithms
- IR models are both effective and efficient (BM25)
  - Direct recommenders (BM25)
  - Neighborhood selectors in kNN
  - Samplers and features in learning to rank
- Accuracy is a partial perspective
- We consider evaluation perspectives beyond accuracy
  - Global network effects beyond (averaged) isolated user gains.
  - New metrics elaborating on weak ties.
- Enhancing the number of weak ties improves novelty & diversity of the information arriving to the users

How to continue?

- Explore further relations with IR
  - Deep learning IR models
  - Other areas: query reformulation, relevance feedback

- Beyond accuracy
  - New dimensions: fairness
  - Find further benefits: reduce glass ceiling effect, radicalization

- Interactive recommendation
  - Analyze the evolution of the structural network properties

Want to know more?

J. Sanz-Cruzado. **Contact recommendation in social networks: algorithmic models, diversity and network evolution**. 2021. PhD thesis. Link

Algorithmic models:

- J. Sanz-Cruzado, P. Castells, C. Macdonald, I. Ounis. **Effective Contact Recommendation in Social Networks by Adaptation of Information Retrieval Models.** *Information Processing & Management* , 57 (5), 102285, September 2020.
- J. Sanz-Cruzado, C. Macdonald, I. Ounis, P. Castells. **Axiomatic Analysis of Contact Recommendation Methods in Social Networks: An IR Perspective.** 42nd European Conference on Information Retrieval (ECIR 2020). Online, April 2020, pp. 157-190.
- J. Sanz-Cruzado, P. Castells. **Information Retrieval Models for Contact Recommendation in Social Networks**. 41st European Conference on Information Retrieval (ECIR 2019). Cologne, Germany, April 2019, pp. 148-163.

Want to know more? (II)

Network diversity:

- J. Sanz-Cruzado, P. Castells. **Enhancing Structural Diversity in Social Networks by Recommending Weak Ties.** *12th ACM Conference on Recommender Systems (RecSys 2018)*, Vancouver, Canada, October 2018, pp. 233-241.
- J. Sanz-Cruzado, P. Castells. **Beyond Accuracy in Link Prediction.** *3rd Workshop on Social Media for Personalization and Search (SoMePeAS 2019)* co-located with *41st European Conference on Information Retrieval (ECIR 2019)*. Cologne, Germany, April 2019, pp. 79-94.
- J. Sanz-Cruzado, S.M. Pepa, P. Castells. **Structural Novelty and Diversity in Link Prediction.** *9th International Workshop on Modeling Social Media (MSM 2018)* co-located with *The Web Conference 2018 (WWW 2018)*. Companion of The Web Conference 2018 . Lyon, France, April 2018, pp. 1347-1351.

Thanks for your attention

E-mail: Javier.sanz-cruzadopuig@glasgow.ac.uk
Twitter: @JavierSanzCruza
Webpage: https://javiersanzcruza.github.io

Slides will be published in the webpage after the seminar

References

- Adomavicius, G., Tuzhilin, A. **Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions.** *IEEE TKDE* 17(6), 2005.
- Bellogín, A., Wang, J., Castells, P. **Bridging memory-based collaborative filtering and text retrieval.** *Information Retrieval* 16(6), pp. 697-724, 2013.
- Burges, C. **From RankNet to LambdaRank to LambdaMART: an overview**. *Microsoft Technical Report*, 2010.
- De Meo, P., Ferrara, E., Fiumara, G., Provetti, A. **On Facebook, Most Ties are Weak.** *Communications of the ACM* 57(11), pp. 78-84, 2012.
- Ganjisaffar, Y., Caruana, R., Lopes, C. **Bagging gradient-boosted trees for high precision, low variance ranking models.** *SIGIR* 2011, Beijing, China, pp. 85-94, 2011.
- Granovetter, M. **The Strength of Weak Ties**. *American Journal of Sociology* 78(6), pp. 1360-1380, 1973.
- Guy, I. **People recommendation on social media.** *Social Information Access*, pp. 570-623, 2018.
- Hannon, J., Bennet, M., Smyth, B. **Recommending Twitter users to follow using content and collaborative filtering approaches.** *ACM RecSys 2010*, Barcelona, Spain, pp. 199-206, 2010.

References (II)

- Liu, T. **Learning to rank for information retrieval.** *Foundations and Trends in Information Retrieval* 3(3), pp. 225-331, 2007.
- Ning, X., Desrosiers, C., Karypis, G. **A comprehensive survey of neighborhood-based recommendation methods.** *Recommender Systems Handbook*, 2nd edition, pp. 37-76, 2015.
- Parapar, J., Bellogín, A., Castells, P., Barreiro, A. **Relevance-based language modelling for recommender systems.** *Information Processing & Management* 49, pp. 966-980, 2013.
- Pariser, E. **The Filter Bubble: How the new personalized web is changing what we read and how we think.** 2011.
- Sanz-Cruzado, J., Castells, P. **Contact Recommendations in Social Networks.** In S. Berkovsky, I. Cantador, D. Tikk (Eds): *Collaborative Recommendations: Algorithms, Practical Challenges and Applications*, World Scientific Publishing, pp. 519-570, 2018.
- Valcarce, D., Parapar, J., Barreiro, A. **Axiomatic analysis of language modelling of recommender systems**. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 25, pp. 113-127, 2017.
- Wang, J., Robertson, S., de Vries, A.P., Reinders, M.J.T. **Probabilistic relevance ranking for collaborative filtering**. *ACM TOIS*, 26, article 16. 2008.