Improving novelty and diversity of nearest-neighbors recommendation by exploiting dissimilarities

Pablo Sánchez[1,3], Javier Sanz-Cruzado[2], **Alejandro Bellogín**[3]
[1]Universidad Pontificia Comillas, Spain
[2]University of Glasgow, United Kingdom
[3]Universidad Autónoma de Madrid, Spain

ECIR IR4Good, April 2025

Context – Nearest-neighbor recommendation

- Methods based on similarities between users or items
- Highly explainable
- Very flexible and customizable
- Strong baselines

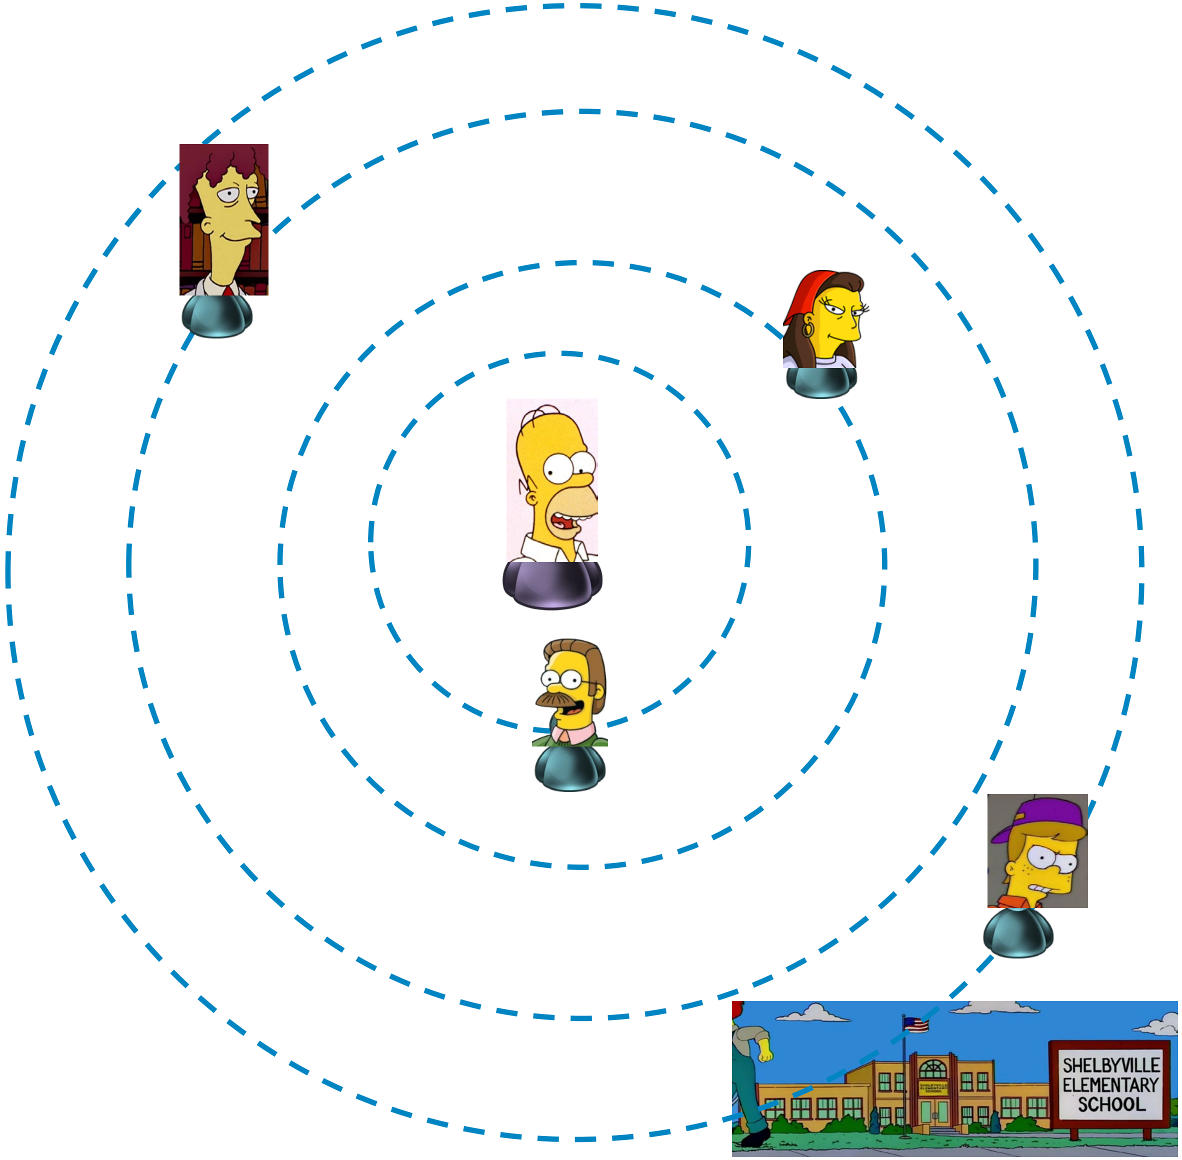

Challenge

- Focusing on similar neighbors or ratings is good for accuracy, but not so good for diversity and novelty

- Accuracy – beyond-accuracy tradeoff
  - Not only limited to neighbor-based recommendation!







Universidad Autónoma de Madrid

Escuela Politécnica Superior

Contribution

User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said, Technische Universität Berlin, alan@dai-lab.de
Ben Fields, Musicmetric (Semetric Ltd.), me@benfields.net
Brijnesh J. Jain, Technische Universität Berlin, jain@dai-lab.de
Sahin Albayrak, Technische Universität Berlin, sahin@dai-lab.de

Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, yangchaoedu@hotmail.com
Cong Cong Ai, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, aicongcong_h@hnu.edu.cn
Renfa Li, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, lirenfa@vip.sina.com

- Inspired by related works, we exploit dissimilarities

- In neighbor models
- In metrics

Contribution

- Inspired by related works, we exploit dissimilarities

User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said Technische Universität Berlin alan@dai-lab.de
Ben Fields Musicmetric (Semetric Ltd.) me@benfields.net
Brijnesh J. Jain Technische Universität Berlin jain@dai-lab.de
Sahin Albayrak Technische Universität Berlin sahin@dai-lab.de

Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China yangchaoedu@hotmail.com
Cong Cong Ai Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China aicongcong_h@hnu.edu.cn
Renfa Li Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China lirenfa@vip.sina.com

- In neighbor models
  - *nndiv*: only using nearest neighbors but considering similarity and dissimilarity scores
- In metrics

Universidad Autónoma de Madrid

Escuela Politécnica Superior

Contribution

- Inspired by related works, we exploit dissimilarities

User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said, Technische Universität Berlin, alan@dai-lab.de
Ben Fields, Musicmetric (Semetric Ltd.), me@benfields.net
Brijnesh J. Jain, Technische Universität Berlin, jain@dai-lab.de
Sahin Albayrak, Technische Universität Berlin, sahin@dai-lab.de

Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, yangchaoedu@hotmail.com
Cong Cong Ai, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, aicongcong_h@hnu.edu.cn
Renfa Li, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, lirenfa@vip.sina.com

- In neighbor models
    - *inds*: compute near and far neighbors and combine scores
- In metrics

Universidad Autónoma de Madrid

Escuela Politécnica Superior

Contribution

User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said
Technische Universität Berlin
alan@dai-lab.de

Ben Fields
Musicmetric (Semetric Ltd.)
me@benfields.net

Brijnesh J. Jain
Technische Universität Berlin
jain@dai-lab.de

Sahin Albayrak
Technische Universität Berlin
sahin@dai-lab.de

Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
yangchaoedu@hotmail.com

Cong Cong Ai
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
aicongcong_h@hnu.edu.cn

Renfa Li
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
lirenfa@vip.sina.com

- Inspired by related works, we exploit dissimilarities

- In neighbor models
  - *indr*: compute predictions based on near and far neigbors and combine final rankings

- In metrics

Contribution

- Inspired by related works, we exploit dissimilarities


User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said, Technische Universität Berlin, alan@dai-lab.de
Ben Fields, Musicmetric (Semetric Ltd.), me@benfields.net
Brijnesh J. Jain, Technische Universität Berlin, jain@dai-lab.de
Sahin Albayrak, Technische Universität Berlin, sahin@dai-lab.de


Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, yangchaoedu@hotmail.com
Cong Cong Ai, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, aicongcong_h@hnu.edu.cn
Renfa Li, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, lirenfa@vip.sina.com

- In neighbor models

- In metrics
  - *rat-diff*: consider difference in ratings consumed by both users

Contribution

- Inspired by related works, we exploit dissimilarities

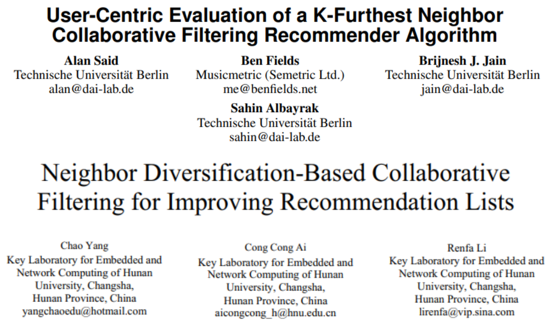
User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said
Technische Universität Berlin
alan@dai-lab.de

Ben Fields
Musicmetric (Semetric Ltd.)
me@benfields.net

Brijnesh J. Jain
Technische Universität Berlin
jain@dai-lab.de

Sahin Albayrak
Technische Universität Berlin
sahin@dai-lab.de

Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
yangchaoedu@hotmail.com

Cong Cong Ai
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
aicongcong_h@hnu.edu.cn

Renfa Li
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
lirenfa@vip.sina.com

- In neighbor models
- In metrics
  - *rdsupp*: consider difference and support of the differences

Universidad Autónoma de Madrid

Escuela Politécnica Superior

Contribution

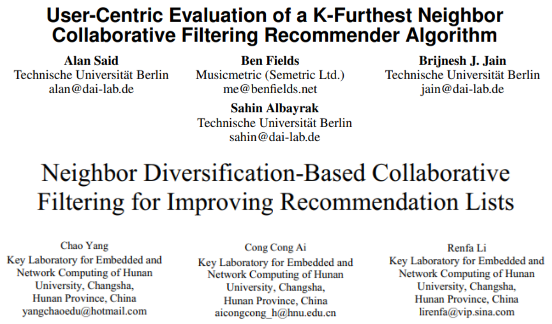

User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said, Technische Universität Berlin, alan@dai-lab.de
Ben Fields, Musicmetric (Semetric Ltd.), me@benfields.net
Brijnesh J. Jain, Technische Universität Berlin, jain@dai-lab.de
Sahin Albayrak, Technische Universität Berlin, sahin@dai-lab.de

Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, yangchaoedu@hotmail.com
Cong Cong Ai, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, aicongcong_h@hnu.edu.cn
Renfa Li, Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China, lirenfa@vip.sina.com

- Inspired by related works, we exploit dissimilarities

- In neighbor models

- In metrics
  - *bin-sets*: preferences are binarized and the sets are compared

Contribution

- Inspired by related works, we exploit dissimilarities

User-Centric Evaluation of a K-Furthest Neighbor Collaborative Filtering Recommender Algorithm

Alan Said
Technische Universität Berlin
alan@dai-lab.de

Ben Fields
Musicmetric (Semetric Ltd.)
me@benfields.net

Brijnesh J. Jain
Technische Universität Berlin
jain@dai-lab.de

Sahin Albayrak
Technische Universität Berlin
sahin@dai-lab.de

Neighbor Diversification-Based Collaborative Filtering for Improving Recommendation Lists

Chao Yang
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
yangchaoedu@hotmail.com

Cong Cong Ai
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
aicongcong_h@hnu.edu.cn

Renfa Li
Key Laboratory for Embedded and Network Computing of Hunan University, Changsha, Hunan Province, China
lirenfa@vip.sina.com

- In neighbor models
  - *nndiv*: only using nearest neighbors but considering similarity and dissimilarity scores
  - *inds*: compute near and far neighbors and combine scores
  - *indr*: compute predictions based on near and far neigbors and combine final rankings
- In metrics
  - *rat-diff*: consider difference in ratings consumed by both users
  - *rdsupp*: consider difference and support of the differences
  - *bin-sets*: preferences are binarized and the sets are compared

Universidad Autónoma de Madrid

Escuela Politécnica Superior

Results – Baseline approaches

- Performance in terms of accuracy (nDCG@5):
  - Letters denote statistical improvements (one-tailed t-test p<0.05)

Datasets	**Pop**$^{(a)}$	**UB**$^{(b)}$	**IB**$^{(c)}$	**HKV**$^{(d)}$	**BPRMF**$^{(e)}$	**EASEr**$^{(f)}$	**RP**$^{3}\beta^{(g)}$
Mov20M	0.126	$\underline{0.300}^{aceg}$	0.205^{ae}	0.312^{abceg}	0.182^{a}	$\mathbf{0.319}^{abcdeg}$	0.260^{ace}
GReads	0.024	0.099^{aef}	$\underline{\mathbf{0.114}}^{abdefg}$	0.104^{abefg}	0.050^{a}	0.081^{ae}	0.097^{aef}
Vinyls	0.003	$\underline{\mathbf{0.068}}^{acdefg}$	0.056^{adefg}	0.054^{aefg}	0.029^{a}	0.052^{aeg}	0.037^{ae}
Lastfm	0.096	$\underline{0.278}^{aceg}$	0.222^{a}	0.306^{aceg}	0.224^{a}	$\mathbf{0.308}^{aceg}$	0.201^{a}

- Nearest-neighbor recommenders are competitive

Results – Dissimilarity-based neighbor models

- Accuracy and beyond-accuracy (novelty: EPC, diversity: Gini and IC):

Dataset	Rec	nDCG	EPC	Gini	IC	UC
Mov20M	kNN/UB	0.300	0.772	0.004	4.8%	100%
	kNN$_{\text{nndiv}}$	0.301	0.772	0.004	4.7%	100%
	kNN$_{\text{inds}}$	**0.304***	**0.783***	**0.005**	**5.4%**	100%
	kNN$_{\text{indr}}$	0.282	0.763	0.004	4.3%	100%
GReads	kNN/IB	0.114	0.974	0.081	36.2%	99.38%
	kNN$_{\text{nndiv}}$	**0.119***	**0.976***	**0.086**	37.1%	99.38%
	kNN$_{\text{inds}}$	0.115	0.974	0.080	36.0%	99.38%
	kNN$_{\text{indr}}$	0.102	0.970	0.081	**37.5%**	99.38%

Dataset	Rec	nDCG	EPC	Gini	IC	UC
Vinyls	kNN/UB	0.068	0.998	0.078	37.9%	99.99%
	kNN$_{\text{nndiv}}$	0.068	0.998	0.076	37.1%	99.99%
	kNN$_{\text{inds}}$	**0.069**	0.998	0.075	37.1%	99.99%
	kNN$_{\text{indr}}$	0.062	0.998	**0.085**	**41.6%**	99.99%
Lastfm	kNN/UB	0.278	0.890	**0.007**	**4.3%**	99.89%
	kNN$_{\text{nndiv}}$	0.279	**0.891**	**0.007**	4.2%	99.89%
	kNN$_{\text{inds}}$	**0.281**	**0.891**	**0.007**	**4.3%**	99.89%
	kNN$_{\text{indr}}$	0.268	0.887	0.006	4.0%	99.89%

- Novelty and diversity improved while keeping (or increasing) same levels of accuracy

Universidad Autónoma de Madrid

Escuela Politécnica Superior

Results – Dissimilarity metrics comparison

- Only for *inds* model (best results overall):

	Mov20M			GReads			Vinyls			Lastfm		
Rec	**nDCG**	**EPC**	**Gini**	**nDCG**	**EPC**	**Gini**	**nDCG**	**EPC**	**Gini**	**nDCG**	**EPC**	**Gini**
kNN	0.3003	0.7721	0.0045	0.1140	0.9741	0.0810	0.0683	0.9982	0.0780	0.2778	0.8901	0.0070
$k\text{NN}_{\text{inds}}$ (rat-diff)	0.3019	0.7818	**0.0053**	0.1134	0.9741	0.0806	0.0665	**0.9986**	**0.1163**	0.2799	**0.9088**	**0.0117**
$k\text{NN}_{\text{inds}}$ (rdsupp)	0.3025	0.7792	0.0051	0.1145	**0.9745**	**0.0821**	**0.0686**	0.9982	0.0750	**0.2807**	0.8912	0.0072
$k\text{NN}_{\text{inds}}$ (bin-sets)	**0.3039**	**0.7832**	**0.0053**	**0.1148**	0.9743	0.0799	0.0684	0.9982	0.0767	0.2802	0.8907	0.0071

- In Mov20M, GReads, and Lastfm: all dimensions improved
- In Vinyls: not able to improve all with the same metric
- Not a clear winner, although *rat-diff* tends to increase novelty and diversity

Conclusions

- Various definitions of neighbor models and dissimilarity metrics based on different hypotheses
- We always obtain improvements in terms of novelty and diversity
  - Sometimes, also in terms of accuracy
- For the future: adapt to implicit information and contextual scenarios, measure content-based novelty or diversity metrics

Thank you

Improving novelty and diversity of nearest-neighbors recommendation by exploiting dissimilarities

Code and other info: https://github.com/pablosanchezp/kNNDissimilarities

Pablo Sánchez[1,3], Javier Sanz-Cruzado[2], Alejandro Bellogín[3]

[1]Universidad Pontificia Comillas, Spain

[2]University of Glasgow, United Kingdom

[3]Universidad Autónoma de Madrid, Spain

ECIR IR4Good, April 2025

Experimental settings

- Four datasets from different domains

Dataset	Users	Items	Ratings	Sparsity (%)	Rating scale
Mov20M	138,493	26,744	20,000,263	0.53998	[0.5-5]
GReads	18,892	25,475	1,378,033	0.28633	[0-5]
Vinyls (5-core)	75,258	64,443	1,097,592	0.02263	[1-5]
Lastfm	1,892	17,632	92,834	0.27828	[1-5]

- Vinyls: 5-core; Lastfm: transformation from implicit to explicit
- Random split 80/20
- TrainItems methodology
- Items in test with rating >= 4, relevant
- All metrics cutoffs, at 5

Neighbor models – Formulation

- *nndiv*: dissimilarity of close neighbors is included so differing stances with respect to the target user/item are taken into account

$$\sum_{v \in N_u^+} \left(w_{uv}^+ + w_{uv}^-\right) \cdot r_{vi}$$

- *inds*: two neigborhoods are computed independently and combined the predictions at the score level

$$\sum_{v \in N_u^+} w_{uv}^+ \cdot r_{vi} + \lambda \sum_{v \in N_u^-} w_{uv}^- \cdot r_{vi}$$

- *indr*: extension of previous model but preferences are integrated at the last step (ranking level)

$$c\left(\bar{n}\left(\mathrm{UB}^+(u,i)\right), \bar{n}\left(\mathrm{UB}^-(u,i)\right)\right)$$

(Dis)similarity metrics – Formulation

- *rat-diff*: two users have differing views if their ratings/scores are as different as possible

$$\text{sim}_{\text{rat-diff}}(u,v) = \left(1 - \frac{1}{|I_{uv}|} \cdot \sum_{i \in I_{uv}} \frac{1}{|r_{ui} - r_{vi}| + 1}\right)$$

- *rdsupp*: two users have differing views if their ratings/scores are different and the support of these differences is large

$$\text{sim}_{\text{rdsupp}}(u,v) = \text{sim}_{\text{rat-diff}}(u,v) \cdot \frac{|I_{uv}|}{|I_u \cup I_v|}$$

- *bin-sets*: two users have differing views if one likes what the other dislikes, and viceversa

$$\text{sim}_{\text{bin-sets}}(u,v) = \frac{|I_u^+ \cap I_v^-| + |I_u^- \cap I_v^+|}{|I_u^+ \cup I_v^-| + |I_u^- \cup I_v^+|}$$