

Overcoming the Imbalance Between Tag Recommendation Approaches and Real-World Folksonomy Structures with Cognitive-Inspired Algorithms

Dominik Kowald

Know-Center & Graz University of Technology, Austria
dkowald@know-center.at

Elisabeth Lex

Graz University of Technology, Austria
elisabeth.lex@tugraz.at

ABSTRACT

Social tagging systems enable users to collaboratively annotate Web resources with freely chosen keywords (i.e., *tags*). In order to assist users in this annotation process, tag recommendation algorithms have been proposed, which suggest a set of tags for a given user and a given resource [5]. Essentially, tag recommendation algorithms aim to help not only the individual to find appropriate tags [4] but also the collective to consolidate the shared tag vocabulary and thus, to reach semantic stability and implicit consensus [10].

Problem. While there is already a large body of research available with respect to tag recommendation algorithms, most existing approaches are designed in a purely data-driven way. Consequently, they rely on dense (i.e., broad) folksonomy structures, in which multiple posts per resource exist to have an appropriate amount of tag assignments available to train the algorithms. In order to ensure these dense data structures, many researchers use p -core pruning, which was, however, shown to unnaturally bias the tag recommendation evaluation process [3].

As indicated in Table 1, the folksonomy structures of real-world social tagging systems are yet typically quite sparse (i.e., narrow), which is reflected by the average number of posts per resource (i.e., $|P|/|R|$). This degree of narrowness (i.e., sparsity) is smaller than two in four of the six systems. Thus, in these settings, current state-of-the-art tag recommendation approaches are not performing very well due to the lack of available training data [7]. This leads to an imbalance between the way how current approaches are designed and how the folksonomies of real-world social tagging systems are structured.

Approach. To overcome this imbalance, we propose to build upon cognitive-inspired tag recommendation algorithms. We ground our approach on related research on the underlying cognitive mechanisms of social tagging, which has shown that the way users choose tags for annotating resources corresponds to processes and structures in human memory [2]. In this respect, a prominent example is the activation equation of the cognitive architecture ACT-R, which formalizes activation processes in human memory [1]. Specifically, the activation equation determines the probability that a specific

Dataset	$ U $	$ R $	$ T $	$ Y $	$ P $	$ P / R $
Flickr	9,590	856,755	125,119	3,328,590	856,755	1.000
CiteULike	18,474	811,175	273,883	3,446,650	900,794	1.110
BibSonomy	10,179	683,478	201,254	2,986,396	772,108	1.129
Delicious	15,980	963,741	184,012	4,266,206	1,447,267	1.501
LastFM	1,892	12,522	9,748	186,474	71,062	5.674
MovieLens	4,009	7,601	15,238	95,580	55,484	7.299

Table 1: Statistics of publicly available datasets for Flickr, CiteULike, BibSonomy, Delicious, LastFM and MovieLens. Here, $|U|$ is the number of users, $|R|$ is the number of resources, $|T|$ is the number of tags, $|Y|$ is the number of tag assignments, $|P|$ is the number of posts and $|P|/|R|$ accounts for the degree of narrowness (i.e., sparsity).

memory unit (e.g., a word or tag) will be needed in a specific context. Thus, such an approach should enable us to efficiently model the tag vocabulary of a user in a cognitive-plausible way without the need for dense data structures associated with the resource to be tagged.

Specifically, in [7, 8], we have implemented a novel tag recommendation algorithm based on the activation equation of ACT-R. This algorithm ranks a user's tags based on three factors: (i) past usage frequency, (ii) past usage recency, and (iii) similarity with current semantic context cues. We have evaluated our approach using six real-world folksonomy datasets (i.e., Flickr, CiteULike, BibSonomy, Delicious, MovieLens and LastFM) and compared it to various state-of-the-art approaches using our *TagRec* framework [6], which is freely-available on Github¹. This evaluation included popularity-based approaches (i.e., MP_r and $MP_{u,r}$), Collaborative Filtering (CF), Latent Dirichlet Allocation (LDA), FolkRank (FR), an alternative recency-based approach (GIRPTM) [11] as well as our cognitive-inspired algorithm (ACT-R).

Results. The results of this evaluation by means of the nDCG@10 metric are given in Table 2. We see that the performance of the classic, data-driven tag recommendation approaches (i.e., MP_r , $MP_{u,r}$, CF, LDA, PITF and FR) strongly depends on the sparsity of the datasets. While their performance is quite weak for sparse (i.e., narrow) ones such as Flickr, CiteULike, BibSonomy and Delicious, they perform very well in dense (i.e., broad) settings such as LastFM and MovieLens. We observe a different behavior with respect to the time-dependent approach GIRPTM. Compared to the data-driven algorithms, this approach achieves the best results in the narrow folksonomy Flickr and the worst in the broad folksonomies LastFM and MovieLens. The overall best results are reached by our cognitive-inspired ACT-R approach.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESCS'17, London, England

© 2017 ACM. 978-x-xxxx-xxxx-x/YY/MM...\$15.00

DOI: 10.1145/nnnnnnn.nnnnnnn

¹<https://github.com/learning-layers/TagRec>

Dataset	MP _r	MP _{u,r}	CF	LDA	PITF	FR	GIRPTM	ACT-R
Flickr	-	.569	.666	.280	.535	.561	.686	.711
CiteULike	.063	.392	.359	.138	.294	.392	.422	.438
BibSonomy	.091	.407	.369	.219	.327	.408	.409	.434
Delicious	.187	.358	.356	.271	.302	.292	.393	.431
LastFM	.283	.386	.317	.388	.414	.399	.397	.425
MovieLens	.271	.328	.254	.296	.324	.319	.326	.338

Table 2: Evaluation results for various tag recommendation algorithms in real-world folksonomy structures by means of the nDCG@10 metric. Our cognitive-inspired approach (i.e., ACT-R) outperforms data-driven methods in all datasets.

Discussion. Our results demonstrate that there indeed exists an imbalance between current state-of-the-art tag recommendation algorithms and the folksonomy structures of real-world social tagging systems. While these algorithms are designed for dense structures, most social tagging systems exhibit a sparse nature. To overcome this imbalance, we showed that cognitive-inspired algorithms, which model the tag vocabulary of a user in a cognitive-plausible way, can be helpful. Our present approach does this via implementing the activation equation of the cognitive architecture ACT-R, which determines the usefulness of units in human memory (e.g., tags).

Apart from that, in [9], we have shown that ACT-R can be generalized for related use cases such as hashtag recommendations in Twitter. For future work, we plan to build upon this results in order to propose cognitive-inspired recommender systems (e.g., for resource recommendation) as an alternative to data-driven ones. In this sense, our long-term goal is to design hybrid approaches, which combine the advantages of both worlds in order to adapt to the current setting (i.e., sparse vs. dense ones).

Acknowledgments. The authors would like to thank Stefanie Lindstaedt, Paul Seitlinger, Simone Kopeinik and Tobias Ley for supporting this research. This work is funded by the Know-Center Graz and the H2020 project AFEL (grant agreement: 687916).

KEYWORDS

Tag Recommendations; Data Sparsity; Folksonomy Structures; ACT-R; Recommender Evaluation

REFERENCES

- [1] John R. Anderson, Daniel Bothell, Michael D. Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. 2004. An integrated theory of the mind. *Psychological review* 111, 4 (2004), 1036.
- [2] Ulrike Cress, Christoph Held, and Joachim Kimmerle. 2013. The collective knowledge of social tags: Direct and indirect influences on navigation, learning, and information processing. *Computers & Education* 60, 1 (2013), 59–73.
- [3] Stephan Doerfel and Robert Jäschke. 2013. An analysis of tag-recommender evaluation procedures. In *Proceedings of the 7th ACM conference on Recommender systems*. ACM, 343–346.
- [4] Robert Jäschke, Leandro Marinho, Andreas Hotho, Lars Schmidt-Thieme, and Gerd Stumme. 2008. Tag recommendations in social bookmarking systems. *Ai Communications* 21, 4 (2008), 231–247.
- [5] Dominik Kowald. 2015. Modeling Cognitive Processes in Social Tagging to Improve Tag Recommendations. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 505–509. DOI: <http://dx.doi.org/10.1145/2740908.2741746>
- [6] Dominik Kowald, Simone Kopeinik, and Elisabeth Lex. 2017. The TagRec Framework As a Toolkit for the Development of Tag-Based Recommender Systems. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*. ACM, New York, NY, USA, 23–28. DOI: <http://dx.doi.org/10.1145/3099023.3099069>
- [7] Dominik Kowald and Elisabeth Lex. 2015. Evaluating Tag Recommender Algorithms in Real-World Folksonomies: A Comparative Study. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. ACM, New York, NY, USA, 265–268. DOI: <http://dx.doi.org/10.1145/2792838.2799664>
- [8] Dominik Kowald and Elisabeth Lex. 2016. The Influence of Frequency, Recency and Semantic Context on the Reuse of Tags in Social Tagging Systems. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media (HT '16)*. ACM, New York, NY, USA, 237–242. DOI: <http://dx.doi.org/10.1145/2914586.2914617>
- [9] Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. 2017. Temporal Effects on Hashtag Reuse in Twitter: A Cognitive-Inspired Hashtag Recommendation Approach. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1401–1410. DOI: <http://dx.doi.org/10.1145/3038912.3052605>
- [10] Claudia Wagner, Philipp Singer, Markus Strohmaier, and Bernardo A. Huberman. 2014. Semantic stability in social tagging streams. In *Proceedings of the 23rd international conference on World Wide Web*. ACM, 735–746.
- [11] Lei Zhang, Jian Tang, and Ming Zhang. 2012. Integrating temporal usage pattern into personalized tag prediction. In *Web Technologies and Applications*. Springer, 354–365.