Stereotype and Most-Popular Recommendations in the Digital Library *Sowiport*

Joeran Beel
Trinity College Dublin
SCSS / KDEG
ADAPT Centre
joeran.beel@adaptcentre.ie

Siddharth Dinesh
Birla Institute for Technology and Science, India
dinesh@mr-dlib.org

Philipp Mayr
GESIS – Leibniz Institute for the Social Sciences, Germany
philipp.mayr@gesis.org

Zeljko Carevic
GESIS – Leibniz Institute for the Social Sciences, Germany
zeljko.carevic@gesis.org

Jain Raghvendra
National Institute of Informatics (NII), Japan
jain@nii.ac.jp

Abstract

Stereotype and *most-popular* recommendations are widely neglected in the research-paper recommender-system and digital-library community. In other domains such as movie recommendations and hotel search, however, these recommendation approaches have proven their effectiveness. We were interested to find out how stereotype and most-popular recommendations would perform in the scenario of a digital library. Therefore, we implemented the two approaches in the recommender system of *GESIS'* digital library *Sowiport*, in cooperation with the recommendations-as-a-service provider *Mr. DLib*. We measured the effectiveness of most-popular and stereotype recommendations with click-through rate (CTR) based on 28 million delivered recommendations. Most-popular recommendations achieved a CTR of 0.11%, and stereotype recommendations achieved a CTR of 0.124%. Compared to a "random recommendations" baseline (CTR 0.12%), and a content-based filtering baseline (CTR 0.145%), the results are discouraging. However, for reasons explained in the paper, we concluded that more research is necessary about the effectiveness of stereotype and most-popular recommendations in digital libraries.

Keywords: recommender systems, digital libraries, evaluation, stereotype recommendations, most-popular recommendations, content-based filtering

Preprint
Joeran Beel, Siddharth Dinesh, Philipp Mayr, Zeljko Carevic, and Jain Raghvendra. "Stereotype and Most-Popular Recommendations in the Digital Library Sowiport." In Proceedings of the 15th International Symposium of Information Science (ISI). To appear in March 2017.

1. Introduction

Recommender systems for research papers typically apply content-based filtering, item-based collaborative filtering, co-occurrence calculations or graph-based recommendations (Beel, Gipp, Langer, & Breitinger, 2015). Two less common recommendation classes are stereotyping and most-popular recommendations. Stereotyping is one of the earliest user-modeling and recommendation classes. In a stereotype recommender system, some generalizing assumptions are made about users (e.g. males like cars and females like perfume), and then items are recommended that presumably are interesting for those stereotype users. A most-popular recommender system adopts a one-fits-all approach and recommends items that have the highest popularity. For instance, a news website could recommend those news articles that were most often read or that had the highest average rating over all users (Lommatzsch, Johannes, Meiners, Helmers, & Domann, 2016). The basic assumption behind such a recommender system is that users will like what most other users read, download, like, etc.

Both stereotype and most-popular recommendations received little attention in the community of research-paper recommender systems, although the two recommendation classes proved effective in other domains (Kay, 2000; A. Kobsa, 1993; Alfred Kobsa, 2001; Lamche, Pollok, Wörndl, & Groh, 2014; Mattioli, 2012; Rich, 1979). Our research goal is to explore the effectiveness of stereotype and most-popular recommendations in digital libraries, more specifically in GESIS' digital library *Sowiport*[1]. The research question we attempt to answer is:

How effective are "Stereotype" and "Most-Popular" recommendations for recommending scholarly literature in digital libraries, Sowiport respectively?

2. Related Work[2]

Stereotype Recommendations

Stereotype recommendations were introduced by *Rich* in the book-recommender system *Grundy* (Rich, 1979). *Rich* was inspired by stereotypes from psychology where stereotypes allowed psychologists to quickly judge people based on a few characteristics. For instance, Rich assumed that male

[1] http://sowiport.gesis.org
[2] Some explanations of stereotype and most-popular recommendations are from Beel, Gipp, et al. (2015).

users have “a fairly high tolerance for violence and suffering, as well as a preference for thrill, suspense, fast plots, and a negative interest in romance”. Consequently, Grundy’s stereotype recommendation approach recommended action books and thrillers to male users.

One major problem with stereotypes is that they may pigeonhole users. While many men have a negative interest in romance, certainly not all do. In addition, building stereotypes is often labor intensive, as the items typically need to be manually classified for each stereotype. This limits the number of e.g. books that could be recommended (Barla, 2011).

Advocates of stereotypes argue that once the stereotypes are created, the recommender system needs little computing power and may perform quite well in practice. For instance, Weber & Castillo (2010) observed that female users were usually searching for the composer Richard Wagner when they entered the search query ‘Wagner’ on *Yahoo!*. In contrast, male users entering the same query usually were looking for the Wagner paint sprayer. Weber & Castillo modified *Yahoo!’*s search algorithm to show the Wikipedia page for Richard Wagner to female users, and the homepage of the Wagner paint sprayer company to male users searching for ‘Wagner.’ As a result, user satisfaction increased. Similarly, the travel agency *Orbitz* observed that Macintosh users were “40% more likely to book a four- or five-star hotel than PC users” and when booking the same hotel, Macintosh users booked the more expensive rooms (Mattioli, 2012). Consequently, Orbitz assigned its website visitors to either the “Mac User” or “PC user” stereotype, and Mac users received recommendations for pricier hotels than PC users. All parties benefited – users received more relevant search results, and Orbitz received higher commissions.

In the domain of research-paper recommender systems, stereotype recommendations have only been applied in the recommender system of the reference manager *Docear* (Beel, Langer, Gipp, & Nürnberger, 2014; Beel, Langer, Kapitsaki, Breitinger, & Gipp, 2015). The developers of the recommender system manually created a list of books and research articles relating to academic writing, and these documents were then recommended to the users of Docear. The authors report a mediocre effectiveness of the stereotype approach with an average click-through rate of 3.08%. In contrast, a standard content-based filtering approach achieved click-through rates slightly below 4%, and a novel content-based filtering approach, tailored to the users of Docear, achieved click-through rates around 7% (for more details about click-through rate as evaluation metric, please refer to Beel & Langer (2015) and the methodology section of the current paper).

We see a need for further research on stereotype recommendations in the domain of digital libraries. The Docear team recommended only documents about one topic, i.e. academic writing, and the recommendations were only tested in Docear. However, recommendation approaches may perform very differently in different scenarios (Beel, Breitinger, Langer, Lommatzsch, & Gipp, 2016; Beel, Langer, Nürnberger, & Genzmehr, 2013). Therefore, we see the need to conduct the research in a different scenario than Docear, and with additional topics than academic writing.

Most-Popular Recommendations

In the domain of research-paper recommender systems, several recommender systems use popularity as an *additional* ranking factor (Bethard & Jurafsky, 2010; He, Pei, Kifer, Mitra, & Giles, 2010; Ren, 2016; Totti, Mitra, Ouzzani, & Zaki, 2016; Zarrinkalam & Kahani, 2013). These systems first determine a list of recommendation candidates, for instance, with content-based filtering. Then, the recommendation candidates are re-ranked based on document popularity. For instance, out of the 20 recommendation candidates that are calculated with content-based filtering, the ten most cited papers might be recommended. Common metrics to calculate popularity are PageRank (Bethard & Jurafsky, 2010), HITS (He et al., 2010), Katz (He et al., 2010), citation counts (Bethard & Jurafsky, 2010; He et al., 2010; Rokach, Mitra, Kataria, Huang, & Giles, 2013), venues' citation counts (Bethard & Jurafsky, 2010; Rokach et al., 2013), citation counts of the authors' affiliations (Rokach et al., 2013), authors' citation count (Bethard & Jurafsky, 2010; Rokach et al., 2013), h-index (Bethard & Jurafsky, 2010), and recency of articles (Bethard & Jurafsky, 2010).

To the best of our knowledge, there is no research on how effective it is to recommend items in a digital library only based on the items' popularity (e.g. loans, views, downloads, citations).

3. Methodology

For our research we used the digital library *Sowiport* (Hienert, Sawitzki, & Mayr, 2015). Sowiport is operated by '*GESIS – Leibniz-Institute for the Social Sciences'*, which is the largest infrastructure institution for the Social Sciences in Germany. Sowiport contains about 9.6 million literature references and 50,000 research projects from 18 different databases, mostly relating to the social and political sciences. Literature references usually cover keywords, classifications, author(s) and journal or conference information and if available: citations, references and links to full texts. On a weekly base, Sowiport reaches around 22,000 unique users. These users spend on average 2 minutes in the system. Sowiport co-operates with Mr. DLib[3], an open Web Service to provide scholarly literature-recommendations-as-a-service (Figure 1). This means that all computations relating to the recommendations run on Mr. DLib's servers, while the presentation takes place on Sowiport's website.

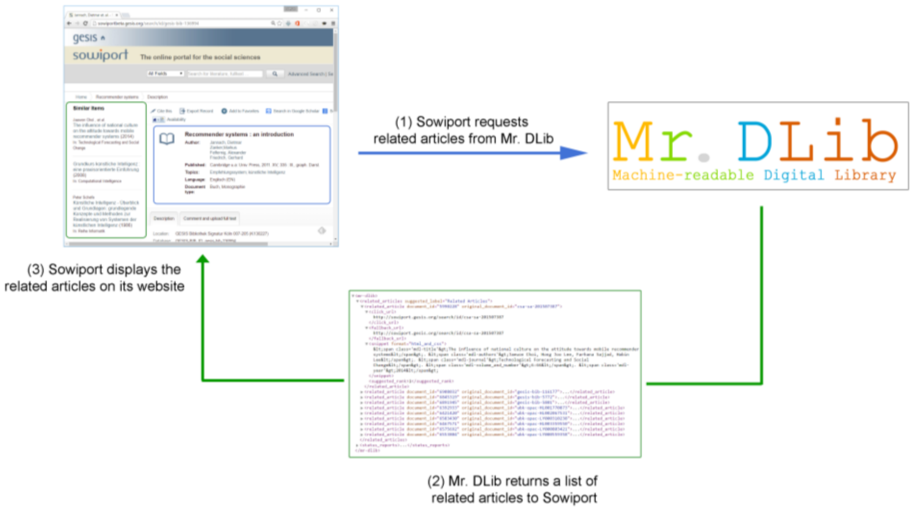

Figure 1: The recommendation process of Sowiport and Mr. DLib

Our recommender system shows related-article recommendations on each article's detail page in Sowiport (Figure 2). Whenever such a detail page is requested by a user, the recommender system randomly chooses one of four recommendation approaches to generate recommendations[4]: 1. stereotype recommendations, 2. most popular recommendations, 3. content-based filtering **(CBF)**, and 4. "random" recommendations, whereas CBF and random recommendations served as baselines. For content-based filtering recommendations, we used Lucene's "More Like This" function, a recommendation approach that is used by many research-paper recommender

[3] http://mr-dlib.org
[4] The approaches are chosen with different probabilities. For instance, random-recommendations were only chosen with a probability of 4% because we needed these kind of recommendations only as baseline.

systems (Beel, Gipp, et al., 2015). When the random approach is chosen, the recommender system randomly picks some documents out of the 9.6 million documents in the recommendation corpus.

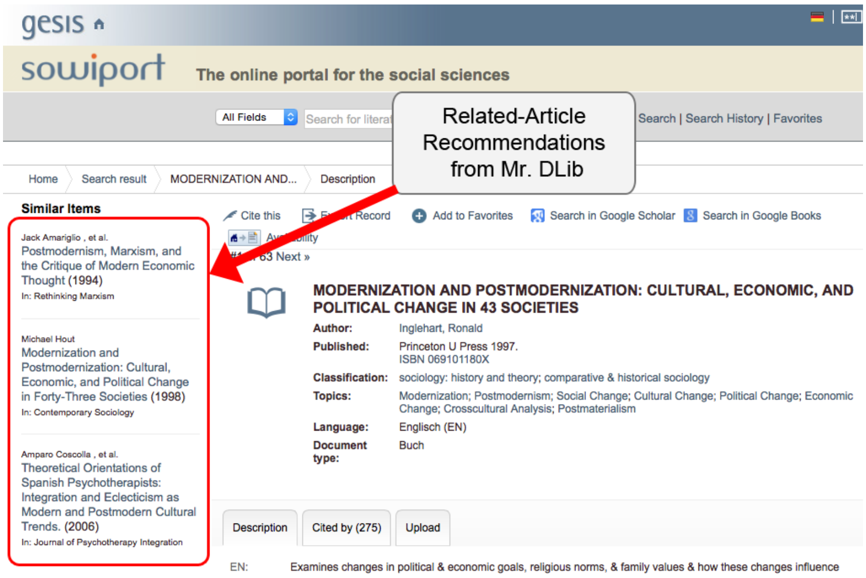

Figure 2: Screenshot of Sowiport's website with recommendations in the left part of the page

To create stereotype recommendations, we assumed that a major part of Sowiport users – who are mostly students and researchers – are interested in the topics "academic writing", "research methods", and "peer review & research evaluation". We used Sowiport's search function to find 16 documents that we considered to be relevant for the three research topics, and these documents were then recommended to the users of Sowiport. Figure 3 shows more details about the 16 documents.

For the most-popular recommendations we used two metrics to measure popularity. First, "views", which measure how often a document's detail-page was accessed by a visitor on Sowiport's website. Second, "exports", which measure how often documents' metadata was exported on Sowiport's website as e.g. BibTeX, EndNote, or email. For both metrics, we identified the 50 most popular documents for the month August, and recommended these documents to the users of Sowiport. Figure 4 shows some of the 2 x 50 documents, a complete list is available from us upon request.

	Sowiport ID	Title	Year	Language
Academic Writing	dzi-solit-000215431	Erfolgreiches wissenschaftliches Schreiben	2015	de
	dzi-solit-0129221	Kreatives wissenschaftliches Schreiben: Tipps und Tricks gegen Schreibblockaden	2001	de
	fis-bildung-1018973	Writing for peer reviewed journals	2013	en
	fis-bildung-1068313	Kreatives Schreiben von Diplom- und Doktorarbeiten	1998	de
	fis-bildung-1071788	Kreatives wissenschaftliches Schreiben	2001	de
	fis-bildung-621436	Geniale Notizen	2002	de
	gesis-bib-126169	Erfolgreiches wissenschaftliches Arbeiten: Seminararbeit, Bachelor-/Masterarbeit (Diplomarbeit), Doktorarbeit	2008	de
Peer Review	csa-sa-201609258	Wissenschaftliches Publizieren: Peer Review	2014	de
	gesis-ssoar-2362	Exzellenz und Evaluationsstandards im internationalen Vergleich	2007	de
	gesis-ssoar-2530	Einleitung: Wie viel (In-)Transparenz ist notwendig? Peer Review Revisited	2006	de
	gesis-ssoar-733	Peer Review in der DFG: die Fachkollegiaten	2007	de
Research Methods	fis-bildung-949616	Empirische Forschungsmethoden	2010	de
	gesis-solis-00569924	Einführung in die Wissenschaftstheorie	2014	de
	gesis-solis-00598617	Forschungsmethoden und Statistik: ein Lehrbuch für Psychologen und Sozialwissenschaftler	2013	de
	gesis-solis-00606948	Forschungsmethoden	2013	de
	iab-litdok-K110511315	Handbuch Qualitative Forschungsmethoden in der Erziehungswissenschaft	2010	de

Figure 3: Details on the 16 documents that we selected as stereotype recommendations

	Sowiport ID	Title	Year	Language
Top Views	fis-bildung-999945	Guter Chemieunterricht	2013	de
	gesis-solis-00560882	Die Gesellschaft und ihre Gesundheit: 20 Jahre Public Health in Deutschland ; Bilanz und Ausblick einer Wissenschaft	2011	de
	gesis-solis-00551750	Thrillslider: Rutschen, Rausch und Rituale auf Spielplätzen, Festplätzen und in Aqua-Parks	2010	de
	gesis-solis-00526599	Weiterbildungsbeteiligung von Menschen mit Migrationshintergrund in Deutschland	2009	de
	fis-bildung-840181	Kommt der Herbst mit bunter Pracht	2008	de
		…		
Top Exported	gesis-solis-00605639	Organisieren am Konflikt: Tarifauseinandersetzungen und Mitgliederentwicklung im Dienstleistungssektor	2013	de
	gesis-solis-00606019	Soziale Arbeit und Stadtentwicklung: Forschungsperspektiven, Handlungsfelder, Herausforderungen	2013	de
	gesis-solis-00580567	Fokusgruppen in der empirischen Sozialwissenschaft: von der Konzeption bis zur Auswertung	2012	de
	gesis-solis-00563254	Handbuch zur Verwaltungsreform	2011	de
	gesis-solis-00568965	Die Zukunft auf dem Tisch: Analysen, Trends und Perspektiven der Ernährung von morgen	2011	de
		…		

Figure 4: Details on the most viewed and exported documents (excerpt)

We measured the effectiveness of the recommendation approaches with click-through rate **(CTR)**. CTR describes the ratio of delivered to clicked recommendations. For instance, when 10,000 recommendations based on CBF were delivered, and 50 of these recommendations were clicked, the average CTR of CBF would be $\frac{50}{10{,}000} = 0.5\%$. The assumption is that the higher the CTR, the more effective is the recommendation approach. There is some discussion to what extend CTR is appropriate for measuring recommendation effectiveness, but overall it has been demonstrated to be a meaningful and well-suited metric (Beel & Langer, 2015; Joachims, Granka, Pan, Hembrooke, & Gay, 2005; Schwarzer et al., 2016).

Table 1: Number of displayed and clicked recommendations by recommendation approach

	Total	Content Based Filtering	Most Popular			Stereotype				Random
			Top Views	Top Exports	Overall	Academic Writing	Research Methods	Research Evaluation	Overall	
Displayed	28,214,883	24,335,531	1,187,845	1,060,647	2,248,492	149,235	147,034	84,938	381,207	1,249,653
Clicks	31,872	27,423	1,373	1,107	2,480	175	192	107	474	1,495

Between 17 October 2016 and 28 December 2016, Mr. DLib's recommender system delivered 28,214,883 recommendations to Sowiport[5]. Whenever comparing results of different algorithms, we report the significance level *p*, which is calculated with a two-tailed t-test. All data relating to this paper is available on *Harvard's Dataverse*[6], including a list of the delivered and clicked recommendations as CSV file, the *R* script to analyze the data, and the figures and tables presented in this paper as PNG and CSV files (Beel, Dinesh, Mayr, Carevic, & Raghvendra, 2017).

4. Results

Figure 5 shows the click-through rates for the four recommendation approaches. Content-based filtering performed best with an average CTR of 0.145%, compared to a CTR of 0.12% for random recommendations (p=0.03). Stereotype recommendations performed second best with a CTR of 0.124% on average, which is an improvement compared to random-recommendations, however, with low significance (p=0.47). Most-popular recommendations were even slightly less effective (CTR = 0.11%) than random recommendations, with high statistical significance (p=0.01).

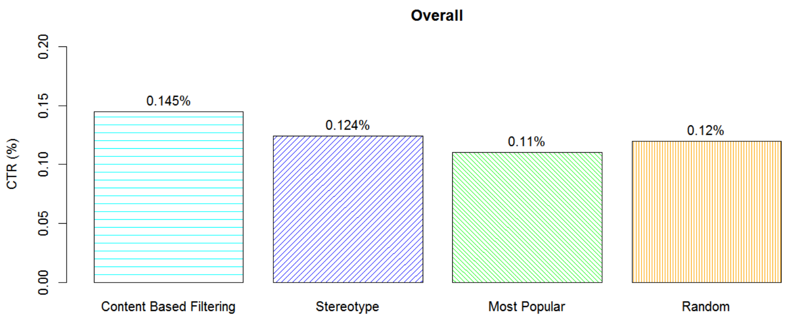

Figure 5: CTR for the different recommendation approaches

[5] Whenever an article's detail page was shown to a user, Mr. DLib returned between 1 and 15 related-article recommendations. Numbers include recommendations delivered to Bots that crawled the Sowiport website. Clicks were recorded via JavaScript. Hence, click-through rates overall are rather low. Numbers include only recommendations that required 3 or less seconds to calculate because in the other cases we could not be sure that the recommendations were actually displayed to a user.

[6] https://dataverse.harvard.edu/dataverse/Mr_DLib

For the most-popular recommendations it made no difference whether we used *exports* or *views* to determine the most popular recommendations (Figure 6). CTR was 0.104% and 0.116% respectively, i.e. both CTRs are below CTR of random recommendations.

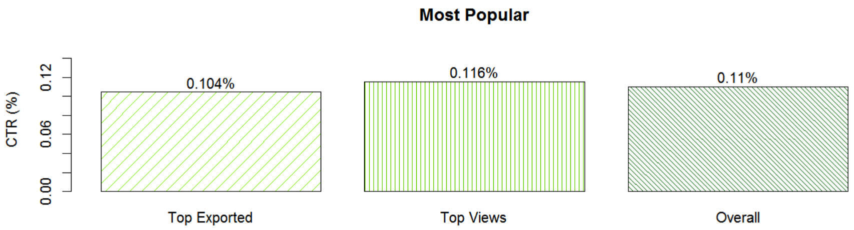

Figure 6: CTR for the most-popular recommendation categories

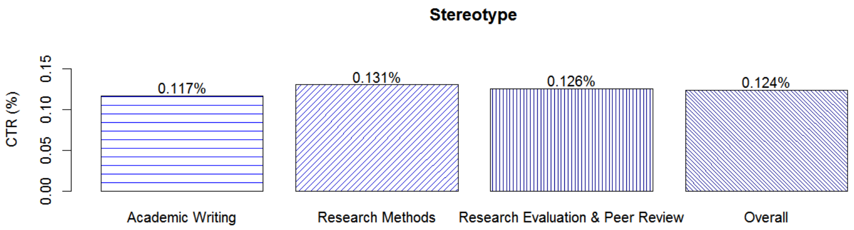

Figure 7: CTR for the different categories of stereotype recommendations

Looking at stereotype recommendations in detail reveals that CTR for the different categories varied (Figure 7). Recommendations for scholarly literature about academic writing achieved the lowest CTR (0.117%) among the stereotype recommendations. Recommendations about peer review and research evaluation achieved CTRs of 0.126%, and recommendations for literature about research methods achieved performed best with a CTR of 0.131%. However, the differences are statistically not significant.

5. Conclusion and Outlook

Overall, the results are somewhat disappointing. Stereotype recommendations were about as (in)effective as random recommendations with both having a CTR of 0.124% and 0.12% respectively. This result contradicts previous research about stereotype recommendations from the Docear researchers. Most-popular recommendations were even statistical significantly less effective (CTR = 0.11%) than random recommendations.

Based on the current results, it seems not sensible to apply stereotype and most-popular recommendations, at least not on Sowiport. However, to reach a final conclusion we consider more research to be necessary. Among others, additional evaluation metrics might be sensible. In addition, a better detection of web spiders crawling the Sowiport website (and hence requesting recommendations), would lead to more reliable data. It might also make sense to experiment with other popularity metrics than views and exports and longer or shorter periods of time to define a popular item. One interesting metric might be "libcitations" (White et al., 2009). Libcitations count a libraries' stock of a given book and give an indicator of its popularity in that library. In addition, the effectiveness of most-popular recommendations could be researched in other scenarios, for instance in smaller libraries with a more homogenous user base.

Further research about stereotype recommendations could focus on identifying, which type of items (e.g. research articles, reviews, blog posts, news, software tools, or research projects) and which kind of topics researchers are most interested in. It could also be interesting to build more tailored stereotypes. Currently, we only had one 'class' of stereotypes, i.e. we assumed that all Sowiport visitors had the same interests in academic writing etc. If the recommender system knew, for instance, a visitor's academic status (e.g. professor, post-doc, PhD student) or research discipline, the stereotype recommendations could be tailored better to the different user groups' needs.

Acknowledgements

This work was supported by a fellowship within the postdoc-program of the German Academic Exchange Service (DAAD). This publication also has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number 13/RC/2106. We are further grateful for the support provided by Sophie Siebert and Stefan Feyer.

References

Barla, M. (2011). Towards social-based user modeling and personalization. *Information Sciences and Technologies Bulletin of the ACM Slovakia*, *3*, 52–60.

Beel, J., Breitinger, C., Langer, S., Lommatzsch, A., & Gipp, B. (2016). Towards Reproducibility in Recommender-Systems Research. *User Modeling and User-Adapted Interaction (UMUAI)*, *26*(1), 69–101. doi:10.1007/s11257-016-9174-x

Beel, J., Dinesh, S., Mayr, P., Carevic, Z., & Raghvendra, J. (2017). Stereotype and Most-Popular Recommendations in the Digital Library Sowiport [Dataset]. Harvard Dataverse. doi:10.7910/DVN/HFIV1A

Beel, J., Gipp, B., Langer, S., & Breitinger, C. (2015). Research Paper Recommender Systems: A Literature Survey. *International Journal on Digital Libraries*, 1–34. doi:10.1007/s00799-015-0156-0

Beel, J., & Langer, S. (2015). A Comparison of Offline Evaluations, Online Evaluations, and User Studies in the Context of Research-Paper Recommender Systems. In *Proceedings of the 19th International Conference on Theory and Practice of Digital Libraries (TPDL)*, Lecture Notes in Computer Science (Vol. 9316, pp. 153–168). doi:10.1007/978-3-319-24592-8_12

Beel, J., Langer, S., Gipp, B., & Nürnberger, A. (2014). The Architecture and Datasets of Docear's Research Paper Recommender System. *D-Lib Magazine*, *20*(11/12). doi:10.1045/november14-beel

Beel, J., Langer, S., Kapitsaki, G. M., Breitinger, C., & Gipp, B. (2015). Exploring the Potential of User Modeling based on Mind Maps. In *Proceedings of the 23rd Conference on User Modelling, Adaptation and Personalization (UMAP)*, Lecture Notes of Computer Science (Vol. 9146, pp. 3–17). Springer. doi:10.1007/978-3-319-20267-9_1

Beel, J., Langer, S., Nürnberger, A., & Genzmehr, M. (2013). The Impact of Demographics (Age and Gender) and Other User Characteristics on Evaluating Recommender Systems. In *Proceedings of the 17th International Conference on Theory and Practice of Digital Libraries (TPDL 2013)* (pp. 400–404). Valletta, Malta: Springer.

Bethard, S., & Jurafsky, D. (2010). Who should I cite: learning literature search models from citation behavior. *Proceedings of the 19th ACM international conference on Information and knowledge management* (pp. 609–618). ACM.

He, Q., Pei, J., Kifer, D., Mitra, P., & Giles, L. (2010). Context-aware citation recommendation. *Proceedings of the 19th international conference on World wide web* (pp. 421–430). ACM.

Hienert, D., Sawitzki, F., & Mayr, P. (2015). Digital Library Research in Action – Supporting Information Retrieval in Sowiport. *D-Lib Magazine*, *21*(3/4). doi:10.1045/march2015-hienert

Joachims, T., Granka, L., Pan, B., Hembrooke, H., & Gay, G. (2005). Accurately interpreting clickthrough data as implicit feedback. *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 154–161). Acm.

Kay, J. (2000). Stereotypes, student models and scrutability. *International Conference on Intelligent Tutoring Systems* (pp. 19–30). Springer.

Kobsa, A. (1993). User modeling: Recent work, prospects and hazards. *Human Factors in Information Technology* (Vol. 10, pp. 111–111).

Kobsa, A. (2001). Generic user modeling systems. *User modeling and user-adapted interaction*, *11*(1-2), 49–63.

Lamche, B., Pollok, E., Wörndl, W., & Groh, G. (2014). Evaluating the Effectiveness of Stereotype User Models for Recommendations on Mobile Devices. *UMAP Workshops*.

Lommatzsch, A., Johannes, N., Meiners, J., Helmers, L., & Domann, J. (2016). Recommender Ensembles for News Articles based on Most-Popular Strategies. *Working Notes of the 7th International Conference of the CLEF Initiative, Evora, Portugal*.

Mattioli, D. (2012). On Orbitz, Mac users steered to pricier hotels. *Wall Street Journal, http://online.wsj.com/news/articles/SB10001424052702304458604577488822667325882*.

Ren, X. (2016). *Effective citation recommendation by information network-based clustering*.

Rich, E. (1979). User modeling via stereotypes. *Cognitive science*, *3*(4), 329–354.

Rokach, L., Mitra, P., Kataria, S., Huang, W., & Giles, L. (2013). A Supervised Learning Method for Context-Aware Citation Recommendation in a Large Corpus. *Proceedings of the Large-Scale and Distributed Systems for Information Retrieval Workshop (LSDS-IR)* (pp. 17–22).

Schwarzer, M., Schubotz, M., Meuschke, N., Breitinger, C., Markl, V., & Gipp, B. (2016). Evaluating Link-based Recommendations for Wikipedia. *Proceedings of the 16th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL)*, JCDL '16 (pp. 191–200). Newark, New Jersey, USA: ACM. doi:10.1145/2910896.2910908

Totti, L. C., Mitra, P., Ouzzani, M., & Zaki, M. J. (2016). A Query-oriented Approach for Relevance in Citation Networks. *Proceedings of the 25th International Conference Companion on World Wide Web* (pp. 401–406). International World Wide Web Conferences Steering Committee.

Weber, I., & Castillo, C. (2010). The demographics of web search. *Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval* (pp. 523–530). ACM.

White, H. D., Boell, S. K., Yu, H., Davis, M., Wilson, C. S., & Cole, F. T. H. (2009). Libcitations: A measure for comparative assessment of book publications in the humanities and social sciences. *Journal of the American Society for Information Science and Technology*, *60*(6), 1083–1096. doi:10.1002/asi.21045

Zarrinkalam, F., & Kahani, M. (2013). SemCiR - A citation recommendation system based on a novel semantic distance measure. *Program: electronic library and information systems*, *47*(1), 92–112.