Research-Paper Recommender Systems:
A Literature Survey

Joeran Beel[1], Bela Gipp[2], Stefan Langer[1], and Corinna Breitinger[1, 3]

[1]Docear, Magdeburg, Germany {beel | langer | breitinger}@docear.org
[2]University of Konstanz, Department of Computer and Information Science, Konstanz, Germany
bela.gipp@uni-konstanz.de
[3]Linnaeus University, School of Computer Science, Physics and Mathematics, Kalmar, Sweden

Abstract. In the last sixteen years, more than 200 research articles were published about *research-paper recommender systems*. We reviewed these articles and present some descriptive statistics in this paper, as well as a discussion about the major advancements and shortcomings and an overview of the most common recommendation concepts and approaches. We found that more than half of the recommendation approaches applied content-based filtering (55%). Collaborative filtering was applied by only 18% of the reviewed approaches, and graph-based recommendations by 16%. Other recommendation concepts included stereotyping, item-centric recommendations, and hybrid recommendations. The content-based filtering approaches mainly utilized papers that the users had authored, tagged, browsed, or downloaded. TF-IDF was the most frequently applied weighting scheme. In addition to simple terms, n-grams, topics, and citations were utilized to model users' information needs. Our review revealed some shortcomings of the current research. First, it remains unclear which recommendation concepts and approaches are the most promising. For instance, researchers reported different results on the performance of content-based and collaborative filtering. Sometimes content-based filtering performed better than collaborative filtering and sometimes it performed worse. We identified three potential reasons for the ambiguity of the results. A) Several evaluations had limitations. They were based on strongly pruned datasets, few participants in user studies, or did not use appropriate baselines. B) Some authors provided little information about their algorithms, which makes it difficult to re-implement the approaches. Consequently, researchers use different implementations of the same recommendations approaches, which might lead to variations in the results. C) We speculated that minor variations in datasets, algorithms, or user populations inevitably lead to strong variations in the performance of the approaches. Hence, finding the most promising approaches is a challenge. As a second limitation, we noted that many authors neglected to take into account factors other than accuracy, for example overall user satisfaction. In addition, most approaches (81%) neglected the user-modeling process and did not infer information automatically but let users provide keywords, text snippets, or a single paper as input. Information on runtime was provided for 10% of the approaches. Finally, few research papers had an impact on research-paper recommender systems in practice. We also identified a lack of authority and long-term research interest in the field: 73% of the authors published no more than one paper on research-paper recommender systems, and there was little cooperation among different co-author groups. We concluded that several actions could improve the research landscape: developing a common evaluation framework, agreement on the information to include in research papers, a stronger focus on non-accuracy aspects and user modeling, a platform for researchers to exchange information, and an open-source framework that bundles the available recommendation approaches.

Table of Contents

1 Introduction

In 1998, *Giles et al.* introduced the first research-paper recommender system as part of the *CiteSeer* project [14]. Since then, at least 216 articles relating to 120 research-paper recommendation approaches were published [1–13], [15–188], [190–218]. The amount of literature and approaches represents a problem for new researchers: they do not know which of the articles are most relevant, and which recommendation approaches are most promising. Even researchers familiar with research-paper recommender systems would find it difficult to keep track of the current developments, since the yearly number of articles steadily increases: 66 of the 217 articles (30%) were published just in 2012 and 2013 alone (Figure 1 & Table 1). The few existing literature surveys in the field [186–188] cover just a fraction of the articles, or focus on selected aspects, such as recommender-system evaluation [191]. Thus, they do not provide an overview of the research field, or identify the most promising approaches.


	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013
Cumulated	2	4	8	15	26	39	51	57	70	85	100	121	137	151	177	217
New (per year)	2	2	4	7	11	13	12	6	13	15	15	21	16	14	26	40

Figure 1: Annual publications in the field of research paper recommender systems[1]

We survey the field of research-paper recommender systems with the goal of enabling researchers and developers to (a) learn about the status-quo of research-paper recommender systems, (b) identify promising fields of research, and (c) motivate the community to solve the most urgent problems that currently hinder the effective use of research-paper recommender systems in practice. For clarity, we use the term "article" to refer to the reviewed journal articles, patents, websites, etc., and the term "paper" to refer

---

[1] Numbers are based on our literature search. Although, we believe our survey to be the most comprehensive survey about research-paper recommender systems, we may have missed a few articles. In addition, most likely, more than 40 papers were published in 2013 since we conducted the literature search in January 2014. Articles presented at conferences in late 2013 most likely had not been published in conferences proceedings by January 2014, and hence were not found through our search. Hence, the total number of papers published is probably higher than 217.

to documents being recommended by research-paper recommender systems[2]. When referring to a large number of recommender systems with certain properties, we cite three exemplary articles. For instance, when we report how many recommender systems apply content-based filtering, we report the number and provide three references [6], [58], [80].

Table 1: List of reviewed articles by year

Year	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013
Reference	[14], [43]	[80], [155]	[106], [123], [131], [139]	[34], [42], [82], [91], [141], [154], [174]	[31–33], [57], [88], [90], [93], [102], [103], [150], [173]	[8], [30], [35–39], [48], [140], [157], [160], [164], [165]	[27], [55], [81], [94], [105], [107], [108], [117], [162], [163], [169], [172]	[64], [122], [130], [133], [134], [188]	[2], [15], [23], [46], [53], [54], [83], [87], [104], [129], [135], [144], [196]	[28], [63], [85], [109], [110], [114], [120], [143], [147], [148], [153], [170], [180], [182], [189]	[13], [21], [47], [52], [78], [92], [95], [99], [100], [127], [138], [152], [166], [176], [184]	[3], [11], [18], [19], [24], [44], [45], [66], [89], [96–98], [124], [136], [137], [146], [156], [161], [167], [168], [175]	[12], [22], [26], [40], [51], [62], [65], [79], [115], [119], [121], [145], [151], [158], [181], [185]	[9], [10], [17], [29], [41], [50], [77], [84], [86], [101], [111], [116], [118], [186]	[1], [16], [49], [56], [58–61], [67], [68], [71–76], [126], [132], [142], [149], [171], [177–179], [183], [187]	[4–7], [20], [25], [69], [70], [112], [113], [125], [159], [190–195], [197–218]

To identify relevant literature for our survey, we conducted a literature search on *Google Scholar*, *ACM Digital Library, Springer Link,* and *ScienceDirect*. We searched for *[paper | article | citation] [recommender | recommendation] [system | systems]* and downloaded all articles that had relevance for research-paper recommender systems. Our relevance judgment made use of the title and the abstract if the title alone did not indicate a recognizable relevance to research-paper recommender systems. We examined the bibliography of each article. If an entry in the bibliography pointed to a relevant article not yet downloaded, we downloaded that article. In addition, we checked on Google Scholar which articles cited the relevant article. If one of the citing articles seemed relevant, we also downloaded it. We expanded our search to websites, blogs, patents, and presentations on major academic recommender systems. These major academic services include the academic search engines *CiteSeer*(x)[3], *Google Scholar* (*Scholar Update*)[4], and

---

[2] Some recommender systems also recommended "citations" but in our opinion, differences between recommending papers and citations are marginal, which is why we do not distinguish between these two terms in this paper.
[3] http://citeseerx.ist.psu.edu
[4] http://scholar.google.com/scholar?sciupd=1&hl=en&as_sdt=0,5

PubMed[5]; the social network *ResearchGate*[6]; and the reference managers *CiteULike*[7], *Docear*[8]*,* and *Mendeley*[9]. While these systems offer recommender systems along with their main services, there are also a few stand-alone recommender systems, namely *BibTip*[10], *bX*[11], *RefSeer*[12], *TheAdvisor*[13] and an experimental system called *Sarkanto*[14].

The first literature search was conducted in June 2013 and found 188 relevant articles [1–188]. Three of the 188 articles were literature surveys [186–188], which were ignored in our survey. The remaining 185 articles consist of peer-reviewed conference articles (59%), journal articles (16%), pre-prints (5%), and other formats such as PhD theses, patents, presentations, and web pages (Table 2). Overall, the reviewed articles were comprehensive, with a median page count of eight. More than one-third (36%) had 10 or more pages (Figure 2). Another 23% had eight or nine pages, while only 26% of the articles had four or less pages.

Table 2: Percentage of different article types reviewed

| Journal articles	Conference papers	PhD Theses	Master's Theses	Patents	Pre-prints/unpublished	Other
16%	59%	3%	2%	1%	5%	15%

Figure 2: Page count of reviewed articles

Citation counts follow a typical power-law distribution: a few articles gained many citations (the maximum was 528 citations for [43]) and many articles had few citations, see Figure 3. The mean citation count was 30, and median was seven. From the reviewed articles, 31% had no citations. Citation counts were retrieved from Google Scholar in early 2014. Some researchers have reservations about using Google Scholar as source for citation counts [219–221], but the numbers can give a rough idea of a paper's popularity.

---

[5] http://www.ncbi.nlm.nih.gov/pubmed
[6] http://www.researchgate.net/
[7] http://www.citeulike.org/
[8] http://www.docear.org
[9] http://www.mendeley.com/
[10] http://www.bibtip.com/
[11] http://www.exlibrisgroup.com/category/bXUsageBasedServices
[12] http://refseer.ist.psu.edu/
[13] http://theadvisor.osu.edu/
[14] http://lab.cisti-icist.nrc-cnrc.gc.ca/Sarkanto/

We reviewed the 185 articles, which contained information on 96 research-paper recommendation approaches [1–185]. In an initial review, we focused on the evaluation of the approaches. The review included an analysis of which evaluation methods were applied (e.g. user-studies or offline evaluations), which evaluation metrics were used (e.g. precision or recall), how many participants the user studies had, and how strongly datasets were pruned.

Of the 96 research-paper recommendation approaches presented in 185 articles, 62 approaches were chosen for an in-depth analysis, presented in 127 articles [1–127]. We chose these 62 approaches, because we classified the remaining 34 approaches, i.e. 58 articles, as not sufficiently related to 'research-paper recommender systems' [129–185], [189]. We classified articles as insufficiently related if they provided no evaluation, or if their approach did not differ significantly from that of previous authors. We also excluded articles that could not be clearly interpreted due to grammar and language use; or when they were outside of the scope (even if the article's title suggested relevance to research-paper recommender systems). One example of an article outside of the research scope was 'Research Paper Recommender Systems - A Subspace Clustering Approach' [130]. The title appears relevant for this survey, but the article presents a collaborative filtering approach that is not intended for recommender systems for research-papers. Instead, the paper used the *Movielens* dataset, which contains ratings of movies. Our analysis aims to obtain a thorough overview of how current research-paper recommendation approaches work. This included an analysis of the distribution of recommendation classes such as collaborative filtering, the use of document fields such as title, abstract, or citation context, and the use of weighting schemes such as TF-IDF.
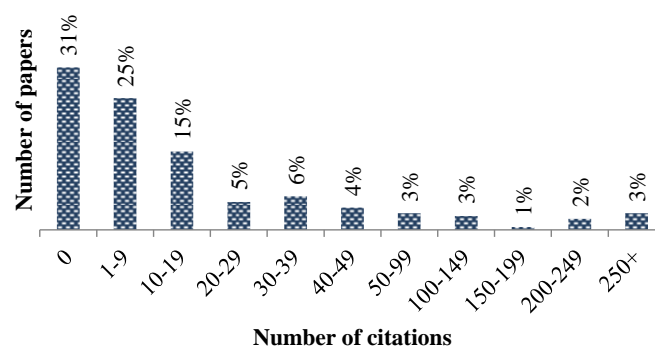

Figure 3: Citation Counts Overview for reviewed articles

In January 2014, we conducted a second literature search and found 29 additional articles of relevance [190–218]. The goal of this search was to identify the overall number of articles published in 2013, see Figure 1. However, time limitations prevented us from broadening the scope of the initially planned survey. Therefore, our review concentrates on the two subsets of the 217 articles: the 185 articles identified during the first round of the literature search; and the 127 articles that we chose for an in-depth review.

In the remaining paper, we present definitions (Section 2), followed by an introduction to related research fields (Section 3). We then present the survey of the 96 approaches' evaluations (Section 4), followed by an analysis of the 62 approaches that we chose for the in-depth review (Section 5). Finally, we examine the field of research-paper recommender systems in general, and point out some shortcomings, such as the neglect of user modeling, a strong focus on accuracy alone at the expense of other aspects being ignored and scarce information on the details of the algorithms used (Section 6).

2 Definitions

We use the term "idea" to refer to a hypothesis about how recommendations could be effectively generated. To differentiate how specific the idea is, we distinguish between recommendation classes, approaches, algorithms, and implementations (Figure 4).

We define a "recommendation class" as the least specific idea, namely a broad concept that broadly describes how recommendations might be given. For instance, the recommendation concepts *collaborative filtering* (CF) and *content-based filtering* (CBF) fundamentally differ in their underlying ideas: the underlying idea of CBF is that users are interested in items that are similar to items the users previously liked. In contrast, the idea of CF is that users like items that the users' peers liked. However, these ideas are rather vague and leave room for different approaches.

A "recommendation approach" is a model of how to bring a recommendation class into practice. For instance, the idea behind CF can be realized with user-based CF [222], content-boosted CF [223], and various other approaches [224]. These approaches are quite different but are each consistent with the central idea of CF. Nevertheless, these approaches to represent a concept are still vague and leave room for speculation on how recommendations are calculated.

A "recommendation algorithm" precisely specifies a recommendation approach. For instance, an algorithm of a CBF approach would specify whether terms were extracted from the title of a document or from the body of the text, and how terms are processed (e.g. stop-word removal or stemming) and weighted (e.g. TF-IDF). Algorithms are not necessarily complete. For instance, pseudo-code might contain only the most important information and ignore basics, such as weighting schemes. This means that for a particular recommendation approach there might be several algorithms.

Finally, the "implementation" is the actual source code of an algorithm that can be compiled and applied in a recommender system. It fully details how recommendations are generated and leaves no room for speculation. It is therefore the most specific idea about how recommendations might be generated.
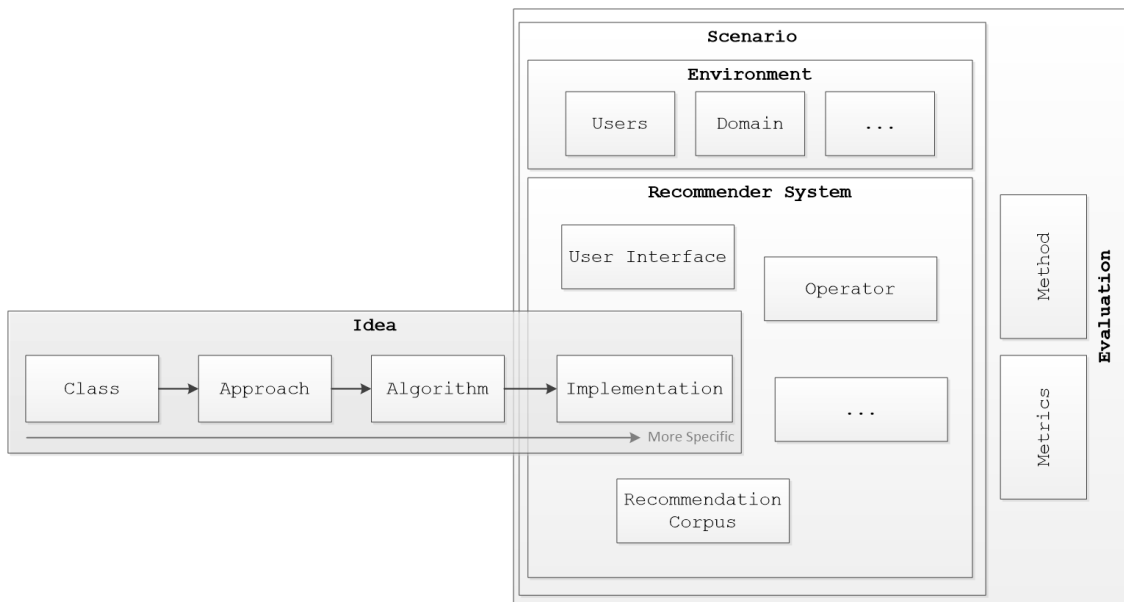
Figure 4: Illustration of recommendation system terminology and concepts

A "recommender system" is a fully functional software system that applies at least one implementation to make recommendations. In addition, recommender systems feature several other components, such as a user interface, a corpus of recommendation candidates, and an operator that owns/runs the system. Some recommender systems also use two or more recommendation approaches: CiteULike, a service for discovering and managing scholarly references, lets their users choose between two approaches [13], [17].

The "recommendation scenario" describes the entire setting of a recommender system, including the recommender system and the recommendation environment, i.e. the domain and user characteristics.

By "effectiveness," we refer to the degree to which a recommender system achieves its objective. The objective of a recommender system from a broad perspective is to provide "good" [225] and "useful" [226] recommendations that make users "happy" [227] by satisfying user needs [87]. The needs of users vary. Consequently, some users might be interested in *novel* research-paper recommendations, while others might be interested in *authoritative* research-paper recommendations. Of course, users require recommendations specific to their fields of research [117]. When we use the term "effectiveness," we refer to the specific objective the evaluator wanted to measure. We use the terms "performance" and "effectiveness" interchangeably.

"Evaluation" describes any kind of assessment that measures the effectiveness or merit of a concrete idea or approach. More details about research paper recommender system evaluation methods follow in Section 4.

3 Related Research Fields

Several research fields are related to user modeling and (research-paper) recommender systems. Although we do not survey these fields, we introduce them so interested readers can broaden their research.

Research on *academic search engines* deals with calculating relevance between research papers and search queries [228–230]. The techniques are often similar to those used by research-paper recommender systems. In some cases, recommender systems and academic search engines are even identical. As described later, some recommender systems require their users to provide keywords that represent their interests. In these cases, research-paper recommender systems do not differ from academic search engines where users provide keywords to retrieve relevant papers. Consequently, these fields are highly related and most approaches for academic search engines are relevant for research-paper recommender systems.

The *reviewer assignment problem* targets using information-retrieval and information-filtering techniques to automate the assignment of conference papers to reviewers [231]. The differences from research-paper recommendations are minimal: in the reviewer assignment problem a relatively small number of paper submissions *must* be assigned to a small number of users, i.e. reviewers; research-paper recommender systems recommend a few papers out of a large corpus to a relatively large number of users. However, the techniques are usually identical. The reviewer assignment problem was first addressed by *Dumais and Nielson* in 1992 [231], six years before *Giles et al.* introduced the first research-paper recommender system. A good survey on the reviewer assignment problem was published by *Wang et al.* [232].

Scientometrics deals with analyzing the impact of researchers, research articles and the links between them. Scientometrics researchers use several techniques to calculate document relatedness or to rank a collection of articles. Some of the measures – h-index [233], co-citation strength [234] and bibliographic coupling strength [235] – have also been applied by research-paper recommender systems [12], [123], [126]. However, there are many more metrics in scientometrics that might be relevant for research-paper recommender systems [236].

User modeling evolved from the field of Human Computer Interaction. One thing user modeling focuses on is reducing users' information overload making use of users' current tasks and backgrounds [237]. User modeling shares this goal with recommender systems, and papers published at the major conferences in both fields (UMAP[15] and RecSys[16]) often

---

[15] http://www.um.org/conferences

[16] http://recsys.acm.org/

overlap. User modeling is a central component of recommender systems because modeling the users' information needs is crucial for providing useful recommendations. For some comprehensive surveys about user modeling in the context of web personalization, refer to [238], [239].

Other related research fields include book recommender systems [240], educational recommender systems [241], academic alerting services [242], expert search [243], automatic summarization of academic articles [244–246], academic news feed recommenders [247], [248], academic event recommenders [249], venue recommendations [250], citation recommenders for patents [251], recommenders for academic datasets [252], and plagiarism detection. Plagiarism detection, like many research-paper recommenders, uses text and citation analysis to identify similar documents [253–255]. Additionally, research relating to crawling the web and analyzing academic articles can be useful for building research-paper recommender systems, for instance, author name extraction and disambiguation [256], title extraction [257–261], or citation extraction and matching [262]. Finally, most of the research on content-based [263] or collaborative filtering [264], [265] from other domains, such as movies or news can also be relevant for research-paper recommender systems.

4   Survey of the Evaluations

Recommender-systems research heavily relies on evaluations to assess the effectiveness of recommendation approaches. Researchers must find out which approaches to apply, either in practice or as a baseline for other evaluations. Among the key prerequisites for thorough evaluations are appropriate evaluation methods, a sufficient number of study participants, and a comparison of the novel approach against one or more state-of-the-art approaches [266]. The novel approach and its evaluation must be clearly described. The soundness of the evaluation, the re-implementation of the approach, and the reproducibility and replicability of the results is guaranteed only if a clear description is given.

We reviewed the evaluation methods, metrics, and datasets used; the number of participants in the user studies; the baselines used for comparing the novel approaches; and several other factors to judge the appropriateness of the evaluations of the 96 approaches. Originally, our goal was to identify the approaches for which evaluations were thoroughly conducted. The further review would have then concentrated on these thoroughly evaluated approaches, to identify the most promising approaches. However, as we will show in the following sections, the majority of evaluations contained limitations, which made it impossible to determine a number of promising approaches.

4.1   Evaluation Methods and their Adequacy

Of the 96 reviewed recommendation approaches, 21 (22%) were not evaluated by their authors [135], [152], [180]. In other cases, an evaluation was attempted, but the methods were questionable and were insufficiently described to be understandable or reproducible [137], [176], [181]. Of the remaining 75 evaluated approaches, 53 (71%) were evaluated using offline evaluations, 25 (33%) using quantitative user studies, two (3%) using qualitative user studies, and five (7%) using an online evaluation (Table 3)[17]. The different evaluation methods and their application in the field of research-paper recommender systems are introduced in the next sections.

Table 3: Evaluation methods used by reviewed recommendation approaches[17]

	Offline	User Study (Quant.)	User Study (Qual.)	Online
Absolute	53	25	2	5
Relative	71%	33%	3%	7%

---

[17] Some approaches were evaluated using several methods. As a result, percentages do not add up to 100%.

4.1.1 User Studies

User studies typically measure user satisfaction through explicit ratings. Users receive recommendations generated by different recommendation approaches, users rate the recommendations, and the approach with the highest average rating is considered most effective [264]. Study participants are typically asked to quantify their overall satisfaction with the recommendations. However, they might also be asked to rate individual aspects of a recommender system, for instance, how novel or authoritative recommendations are [117], or how suitable they are for non-experts [72]. A user study can also collect qualitative feedback, but qualitative feedback is rarely used in the field of (research-paper) recommender systems [156], [159].

We distinguish between "lab" and "real-world" user studies. In lab studies, participants are aware that they are part of a user study, which together with other factors might affect user behavior and thereby the evaluation's results [267], [268]. In real-world studies, participants are not aware of the study and rate recommendations for their own benefit, for instance because the recommender system improves recommendations based on user ratings (i.e. relevance feedback [263]), or user ratings are required to generate recommendations (i.e. collaborative filtering [222]). All reviewed user studies were lab-based.

Often, user studies are considered the optimal evaluation method [269]. However, the outcome of user studies often depends on the questions asked. *Cremonesi et al.* found that it makes a difference if users are asked for the "perceived relevance" or the "global satisfaction" of recommendations [270]. Similarly, it made a difference whether users were asked to rate the *novelty* or the *relevance* of recommendations [271]. A large number of participants are also crucial to user study validity, which makes user studies relatively expensive to conduct. The number of required participants, to receive statistically significant results, depends on the number of approaches being evaluated, the number of recommendations being displayed, and the variations in the results [272], [273]. However, as a rough estimate, at least a few dozen participants are required, often more.

Most participants rated only a few recommendations and four studies (15%) were conducted with fewer than five participants [62], [123], [171]; five studies (19%) had five to ten participants [66], [84], [101]; three studies (12%) had 11-15 participants [15], [146], [185]; and five studies (19%) had 16-50 participants [44], [118], [121]. Six studies (23%) were conducted with more than 50 participants [93], [98], [117]. Three studies (12%) failed to mention the number of participants [55], [61], [149] (Table 4). Given these findings, we conclude that most user studies were not large enough to arrive at meaningful conclusions.

Table 4: Number of participants in user studies of reviewed recommendation approaches

	Number of Participants					
	n/a	<5	5-10	11-15	16-50	>50
Absolute	3	4	5	3	5	6
Relative	12%	15%	19%	12%	19%	23%

4.1.2   Online Evaluations

Online evaluations were first used by the online advertising and e-commerce fields. They measure the acceptance rates of recommendations in real-world recommender systems. Acceptance rates are typically measured by click-through rates (CTR), i.e. the ratio of clicked recommendations to displayed recommendations. For instance, if a recommender system displays 10,000 recommendations and 120 are clicked, the CTR is 1.2%. Other metrics include the ratio of downloaded or bought items to the number of items displayed. Acceptance rate is typically interpreted as an implicit measure for user satisfaction. The assumption is that when a user clicks, downloads, or buys a recommended item, the user liked the recommendation. Of course, this assumption is not always reliable because users might buy a book but rate it negatively after reading it. However, metrics such as CTR can be an *explicit* measures of effectiveness, namely when the operator receives money, e.g. for clicks on recommendations.

Online evaluations are not without drawbacks. *Zheng et al.* showed that CTR and relevance do not always correlate and concluded that "CTR may not be the optimal metric for online evaluation of recommender systems" and "CTR should be used with precaution" [274]. In addition, conducting online evaluations requires significantly more time than offline evaluations, they are more expensive, and they can only be conducted by researchers who have access to a real-world recommender system.

Among the 75 approaches that included some form of evaluation, only six (8%) were evaluated using an online evaluation [6], [92], [94]. Despite the active experimentation in the field with a large number or evaluations being performed on research-paper recommender systems, we observed that many researchers have no access to real-world systems to evaluate their approaches. Interestingly, the researchers who *do* have access to real-world recommender systems often do not make use of this resource to conduct online evaluations, but rather perform offline evaluations or lab user studies. For instance, *Giles* and his co-authors, who are some of the largest contributors in the field, could have conducted online evaluations with their academic search engine CiteSeer. However, they chose primarily to use offline evaluations. The reason for this might be that offline evaluations are more convenient than conducting online evaluations or user studies. Results are available within minutes or hours and not within days or weeks as is the case for user studies and online evaluations. However, offline evaluations have a set of serious drawbacks, as shown in the next section.

4.1.3   Offline Evaluations

Offline evaluations typically measure the *accuracy* of a recommender system based on a ground-truth. To measure accuracy, precision at position *n* (P@n) is often used to express how many items of the ground-truth are recommended within the top *n* recommendations. Other common evaluation metrics include recall, F-measure, mean reciprocal rank (MRR), normalized discounted cumulative gain (nDCG), mean absolute error, and root mean square error. Offline evaluations are also sometimes used to evaluate aspects such as novelty or serendipity of recommendations [227]. For a comprehensive overview of offline evaluations refer to [275], [276].

Offline evaluations were originally meant to identify a number of promising recommendation approaches [225], [264], [277], [278]. These approaches should then be evaluated in detail with a user study or online evaluation to identify the most effective approaches. However, criticism has been raised on the assumption that offline evaluation could predict an algorithm's effectiveness in online evaluations or user studies. More precisely, several researchers have shown that results from offline evaluations do not necessarily correlate with results from user studies or online evaluations [93], [270], [271], [279–282]. This means that approaches that are effective in offline evaluations are not necessarily effective in real-world recommender systems. *McNee et al.* observed that

> "the research community's dependence on offline experiments [has] created a disconnect between algorithms that score well on accuracy metrics and algorithms that users will find useful." [87]

Other researchers also voiced criticism of offline evaluations. *Jannach et al.* stated that "the results of offline [evaluations] may remain inconclusive or even misleading" and "real-world evaluations and, to some extent, lab studies represent probably the best methods to evaluate systems" [283]. *Knijnenburg et al.* reported that "the presumed link between algorithm accuracy […] and user experience […] is all but evident" [284]. *Said et al.* consider "on-line evaluation [as] the only technique able to measure the true user satisfaction" [269]. *Rashid et al.* observed that biases in the offline datasets may cause bias in the evaluation [278]. The main reason for the criticism in the literature is that offline evaluations focus on accuracy yet ignore human factors; however, human factors strongly affect overall user satisfaction for recommendations. Despite the criticism, offline evaluations are the predominant evaluation method in the recommender community [285] and "surprisingly few studies [evaluate] algorithms in live experiments with real users" [284].

Our review indicates that the voiced criticism of offline evaluations also applies to the field of research-paper recommender systems. Some of the approaches were evaluated using both an offline evaluation and a user study. In two evaluations, results from the offline evaluations were indeed similar to results of the user studies [26], [84]. However,

the user studies had five and 19 participants respectively, which led to statistically insignificant results. Three other studies reported contradicting results for offline evaluations and user studies (two of these studies had more than 100 participants) [57], [93], [117]. This means, offline evaluations could not reliably predict the effectiveness in the real-world use case. Interestingly, the three studies with the most participants were all conducted by the authors of TechLens [26], [93], [117], who are also the only authors in the field of research-paper recommender systems who discuss the potential shortcomings of offline evaluations [87]. It seems that other researchers in this field are not aware of – or chose not to address – problems associated with offline evaluations, although there has been quite a discussion outside the research-paper recommender-system community [93], [270], [271], [279–282].

4.2   The Operator's Perspective

It is commonly assumed that the objective of a recommender system is to make users "happy" [227] by satisfying their needs [87]. However, there is another important stakeholder who is often ignored: the operator of a recommender system [225]. It is often assumed that operators of recommender systems are satisfied when their users are satisfied, but this is not always the case. Operators may also want to keep down costs of labor, disk storage, memory, computing power, and data transfer [264]. Therefore, for operators, an effective recommender system may be one that can be developed, operated, and maintained at a low cost. Operators may also want to generate a profit from the recommender system [225]. Such operators might prefer to recommend items with higher profit margins, even if user satisfaction is not optimal. For instance, publishers might be more interested in recommending papers the user must pay for than papers the user can freely download.

The operator's perspective has been widely ignored in the reviewed articles. Costs of *building* a recommender system, or implementing an approach were not reported in any article. Costs to run a recommender system were reported by Jack from Mendeley [59]. He stated that the costs on Amazon's S3 were $66 a month plus $30 to update the recommender system that served 20 requests per second generated by 2 million users.

Runtime information is crucial to estimate costs, and hence to estimate how feasible an approach will be to apply in practice. In one paper, the runtimes of two approaches differed by a factor of 600 [56]. For many operators, an approach that requires 600 times more computing power than another would probably not be an option. While this example is extreme, other runtime comparisons showed differences by a factor of five or more, which can also affect algorithm selection. However, information on runtime was provided only for 10% of the approaches.

Reporting on computational complexity is also important. For operators who want to offer their system to a large number of users, computational complexity is important for estimating the long-term suitability of an approach. An approach may perform well enough for a few users, but it might not scale well. Approaches with exponentially increasing complexity most likely will not be applicable in practice. However, computational complexity was reported for even fewer approaches than runtime.

4.3   Coverage

Coverage describes how many papers of those in the recommender's database might potentially be recommended [286], [287]. As such, coverage is an important metric to judge the usefulness of a recommender system. For text-based approaches, coverage is usually 100%. For other approaches, coverage is typically lower. For instance, in collaborative filtering not all items are rated by users. Although the unrated items might be relevant, they cannot be recommended. High coverage is important because it increases the number of recommendations a user can receive. Of the reviewed articles, few considered coverage in their evaluations. *He et al.* judge the effectiveness of their approaches based on which approach provides the best tradeoff between accuracy and coverage [51]. The BibTip developers report that 80% of all documents have been co-viewed and can be used for generating recommendations [92]. *Pohl et al.* report that co-download coverage on arXiv is close to 100% while co-citation coverage is only around 30% [110]. The TechLens authors report that all of their hybrid and CBF approaches have 100% coverage, except pure CF which has a coverage of 93% [117].

4.4   Baselines

Another important factor in evaluating recommender systems is the baseline against which an algorithm is compared. For instance, knowing that a certain approach has a particular CTR is not useful if the CTRs of alternative approaches are unknown. Therefore, novel approaches should be compared against a baseline representative of the state-of-the-art approaches. This way it is possible to quantify whether, and when, a novel approach is more effective than the state-of-the-art and by what margin.

Of the 75 evaluated approaches, 15 (20%) were not compared against a baseline [27], [86], [185]. Another 53 (71%) approaches were compared against trivial baselines, such as simple content-based filtering without any sophisticated adjustments. These trivial baselines do not represent the state-of-the-art and are not helpful for deciding whether a novel approach is promising. This is particularly troublesome since the reviewed approaches were not evaluated against the *same* trivial baselines. Even for a simple CBF

baseline, there are many variables, such as whether stop words are filtered, which stemmer is applied, or from which document field the text is extracted. This means that almost all reviewed approaches were compared against different baselines, and results cannot be compared with each other. Seven approaches (9%) were evaluated against approaches proposed by other researchers in the field. Only these evaluations allow drawing some conclusions on which approaches may be most effective.

It is interesting to note that in all evaluations, at least one of the novel approaches performed better than the baseline (if the approach was evaluated against a baseline). No article reported on a non-effective approach. We can just speculate about the reasons: First, authors may intentionally select baselines such that their approaches appear favorable. Second, the simple baselines used in most evaluations achieve relatively poor results, so that any alternative easily performs better. Third, authors do not report failures. Lastly, journals and conferences might not accept publications that report on failures. Whatever the reasons are, we advocate that reporting failures is desirable since it could prevent other researchers from doing the same experiments, and hence wasting time.

4.5   Offline Evaluation Metrics

Precision was used as an evaluation metric in 38 offline evaluations (72%) (Table 5). Recall was used in 12 evaluations (23%); F-measure in six evaluations (11%); nDCG in 11 evaluations (21%); MRR in 10 evaluations (19%); and other measures in 12 evaluations (23%). Overall, results of the different measures highly correlated. That is, an algorithm that performed well measured by precision tended to perform well measured by nDCG, for instance. However, there were exceptions. Zarrinkalam and Kahani tested the effectiveness of abstract and title against abstract, title, and citation context [125]. When *co-citation probability* was used as an evaluation metric, title and abstract were most effective. Based on recall, the most effective field combination was abstract, title, and citation context. With the nDCG measure, results varied depending on how the candidate set was generated and which ranking approach was used.

Table 5: Evaluation metrics of reviewed recommendation approaches[18]

	Precision	Recall	F-Measure	nDCG	MRR	Other
Absolute	38	12	6	11	10	12
Relative	72%	23%	11%	21%	19%	23%

---

[18] Some approaches' effectiveness was measured with multiple metrics. Hence, numbers do not add up to 100%

4.6   Datasets and Architectures

Researchers and developers in the field of recommender systems can benefit from publicly available architectures and datasets[19]. *Architectures* help with the understanding and building of recommender systems, and are available in various recommendation domains, such as e-commerce [288], marketing [289], and engineering [290]. *Datasets* enable the evaluation of recommender systems by allowing researchers to evaluate their systems with the same data. Datasets are available in several recommendation domains, including movies[20], music[21], and baby names[22]. Notable are also the various TREC datasets that facilitated and standardized evaluations in several domains[23].

Architectures of research-paper recommender systems were published by few authors. The developers of *CiteSeer(x)* published an architecture that focused on crawling and searching academic PDFs [14], [108]. This architecture has some relevance for recommender systems, since many tasks in academic search are related to recommender systems (e.g. crawling and indexing PDFs, and matching user models or search-queries with research papers). *Bollen and van de Sompel* published an architecture that later served as the foundation for the research-paper recommender system *bX* [15]. This architecture focuses on recording, processing, and exchanging scholarly usage data. The developers of *BibTiP* [33] also published an architecture that is similar to the architecture of bX (both bX and BibTip exploit usage data to generate recommendations).

Several academic services published datasets that eased the process of researching and developing research-paper recommender systems. *CiteULike*[24] and *Bibsonomy*[25] published datasets containing the social tags that their users added to research articles. The datasets were not originally intended for recommender-system research but are frequently used for this purpose [56], [62], [112]. *CiteSeer* made its corpus of research papers public[26], as well as the citation graph of the articles, data for author name disambiguation, and the co-author network [291]. CiteSeer's dataset has been frequently used by researchers for evaluating research-paper recommender systems [20], [24], [51], [56], [65], [106], [112], [117], [125]. *Jack et al.* compiled a dataset based on the reference management software *Mendeley* [292]. The dataset includes 50,000 randomly selected personal libraries from 1.5 million users. These 50,000 libraries contain 4.4 million

---

[19] Recommendation frameworks such as *LensKit or Mahout* may also be helpful for researchers and developers, but frameworks are not the topic of this paper.
[20] http://grouplens.org/datasets/movielens/
[21] http://labrosa.ee.columbia.edu/millionsong/
[22] http://www.kde.cs.uni-kassel.de/ws/dc13/
[23] http://trec.nist.gov/data.html

[24] http://www.citeulike.org/faq/data.adp
[25] https://www.kde.cs.uni-kassel.de/bibsonomy/dumps/
[26] http://csxstatic.ist.psu.edu/about/data

articles of which 3.6 million are unique. Due to privacy concerns, *Jack et al.* publish only the unique IDs of articles and no title or author names. Additionally, only those libraries with at least 20 articles were included in the dataset. *Sugiyama and Kan* released two small datasets[27], which they created for their academic recommender system [115]. The datasets include some research papers, and the interests of 50 researchers. The CORE project released a dataset[28] with enriched metadata and full-texts of academic articles that could be helpful in building a recommendation candidate corpus.

Of the 53 reviewed offline evaluations, 17 (32%) were evaluated using data from CiteSeer and six (11%) were evaluated using data from CiteULike (Table 6). Other data sources included ACM (9%), DBLP (8%), and a variety of others, often not publicly available datasets (51%). Even when data originated from the same sources, it did not guarantee that the same datasets were used. For instance, 32% of the approaches were evaluated with data from CiteSeer but no single CiteSeer dataset exists. Authors collected CiteSeer data at different times and pruned datasets differently. Some authors removed documents with fewer than two citations from the CiteSeer corpus [24], others with fewer than three citations [117], and others with fewer than four citations [137]. Other datasets were pruned even more heavily. *Caragea et al.* removed papers with fewer than ten and more than 100 citations, as well as papers citing fewer than 15 and more than 50 papers [20]. From 1.3 million papers in the corpus, around 16,000 remained (1.2%). *Pennock et al.* removed documents from the corpus with fewer than 15 implicit ratings [106]: from originally 270,000 papers, 1,575 remained (0.58%). It is therefore safe to say that no two studies, performed by different authors, used the same dataset. This raises the question to what extent results based on different datasets are comparable.

Table 6: Source of datasets for reviewed recommendation approaches that performed offline evaluations

	CiteSeer	CiteULike	ACM	DBLP	Others
Absolute	17	6	5	4	27
Relative	32%	11%	9%	8%	51%

Naturally, recommendation approaches perform differently on different datasets [225], [293], [294]. This is particularly true for the absolute effectiveness of recommendation approaches. For instance, an algorithm that achieved a recall of 4% on an IEEE dataset, achieved a recall of 12% on an ACM dataset [101]. The *relative* effectiveness of two approaches is also not necessarily the same with different datasets. For instance, because approach A is more effective than approach B on dataset I, does not mean that A is also more effective than B on dataset II. However, among the few reviewed approaches that were evaluated on different datasets, the effectiveness was surprisingly consistent.

---

[27] http://www.comp.nus.edu.sg/~sugiyama/SchPaperRecData.html
[28] http://core.kmi.open.ac.uk/intro/data_dumps

Of the evaluated approaches, seven were evaluated on multiple offline datasets. Dataset combinations included CiteSeer and some blogs [100], CiteSeer and Web-kd [65], CiteSeer and CiteULike [56], CiteSeer and Eachmovie [106], and IEEE, ACM and ScienceDirect [101]. Results differed notably among the different datasets only in one study. However, the absolute ranking of the approaches remained the same [56] (Table 7). In that article, the proposed approach (CTM) performed best on two datasets (CiteULike and CiteSeer), with a MRR of 0.529 and 0.467 respectively. Three of the four baselines performed similarly on the CiteSeer dataset (all with a MRR between 0.238 and 0.288). However, for the CiteULike dataset the TM approach performed four times as well as CRM. Consequently, if TM had been compared with CRM, rankings would have been similar on the CiteSeer dataset but different on the CiteULike dataset.

Table 7: MRR on different datasets used for offline evaluations

		Dataset	
Rank	Approach	CiteSeer	CiteULike
1	CTM	0.529	0.467
2	TM	0.288	0.285
3	cite-LDA	0.285	0.143
4	CRM	0.238	0.072
5	link-LDA	0.028	0.013

Overall, a sample size of seven is small, but it gives at least some indication that the impact of the chosen dataset is rather low in the domain of research-paper recommender systems. This finding is interesting because in other fields it has been observed that different datasets lead to different results [225], [293]. Nevertheless, we doubt that pruning datasets drastically should be considered good practice, especially if just a fraction of the original data remains.

4.7   Reproducibility and the Butterfly Effect

The reproducibility of experimental results is the "fundamental assumption" in science [295], and the "cornerstone" that allows drawing meaningful conclusions about the generalizability of ideas [296]. Reproducibility describes the situation when (slightly) different ideas, scenarios, and evaluations lead to similar experimental results [295], where we define "similar results" as results that allow the same conclusions to be drawn. *Reproducibility* should not be confused with *replicability*. Replicability describes an exact copy of an experiment that uses the same tools, follows the same steps, and produces the same results [297]. Therefore, replicability is important when analyzing whether the original experiment was conducted thoroughly and whether the results can be trusted.

Conversely, if changes in the ideas, scenarios, or evaluations cause dissimilar results, i.e. results that do not allow the same conclusions to be drawn, we speak of non-reproducibility. Non-reproducibility is expected when significant changes are made to the

ideas, scenarios, or evaluations. However, if minor changes are made but results are unexpectedly dissimilar, then we speak of what we term the "butterfly effect".

During the review, we found several examples of this butterfly effect, i.e. variations in experimental results that we considered unexpected and non-reproducible. For instance, the developers of the recommender system *bx* report that the effectiveness of their recommender system varied by a factor of three at different institutions although the same recommendation approach was used [116]. *Lu et al.* reported that the *translation* model had twice the accuracy of the *language* model [86], but in another evaluation, accuracy was only 18% higher [49]. *Huang et al.* report that the *Context-aware Relevance Model* (CRM) and *cite-LDA* performed similarly, but in another evaluation by the same authors, CRM performed significantly worse than cite-LDA [56]. *Lu et al.* found that sometimes terms from the abstract performed better than terms from the body-text, while sometimes the opposite was true [86]. *Zarrinkalam and Kahani* found that sometimes terms from the title *and* abstract were most effective, while sometimes terms from the title, abstract, *and* citation context were most effective [125]. *Bethard and Jurafsky* reported that citation counts *strongly* increased the effectiveness of their recommendation approach [12], while *He et al.* reported that citation counts *slightly* increased the effectiveness of their approach [51].

Most interesting with respect to the butterfly effect, were some evaluations by the TechLens team (Table 8). The TechLens team evaluated several content-based (CBF) and collaborative filtering (CF) approaches for research-paper recommendations. In 2002, *McNee et al.* conducted an offline evaluation in which CF and CBF performed similarly [93]. However, their additional user study led to a different result – CBF outperformed CF. A user study by *Torres et al.* in 2004 reports results similar to the user study by *McNee et al.* (CBF outperformed CF) [117]. However, the offline evaluation from *Torres et al.* contradicted the previous results – this time, CF outperformed CBF. In 2006, another user study by *McNee et al.* indicated that CF (slightly) outperforms CBF [87], which showed the opposite of the previous user studies. In 2009, *Dong et al.*, who are not affiliated with TechLens, evaluated the approaches of *Torres et al.* with an offline evaluation [24]. In this evaluation, CBF outperformed CF, contradicting the previous offline-results from *Torres et al*. In 2010, *Ekstrand et al.* found that CBF performed worse than CF in both an offline evaluation and a user study, which again did not align with the previous findings [26].

Table 8: Results of different CBF and CF evaluations

	McNee et al. 2002		Torres et al. 2004		McNee et al. 2006		Dong et al. 2009		Ekstrand et al. 2010	
	Offline	User Std.	Offline	User Std.	Offline	User Std.	Offline	User Std.	Offline	User Std.
CBF	Draw	Win	Lose	Win	--	Lose	Win	--	Lose	Lose
CF	Draw	Lose	Win	Lose	--	Win	Lose	--	Win	Win

The authors of the studies provide some potential reasons for the variations, such as different datasets (as discussed in Section 4.6) differences in user populations, and variations in the implementations. However, these reasons can only explain *some* of the variations. Overall, we consider most of the different outcomes to be unexpected. We view this as a problem, since we see the primary purpose of evaluations in aiding developers and researchers to identify the most effective recommendation approaches (for a given scenario). Consequently, a developer looking for an effective recommendation approach, or a researcher needing an appropriate baseline to compare a novel approach against, would not find much guidance in the existing evaluations. Similarly, the currently existing evaluations do not help to identify whether CF or CBF is more promising for research-paper recommender systems.

Interestingly, reproducibility is widely ignored by the (research-paper) recommender-system community, even by researchers focusing on recommender-systems evaluation. For instance, *Al-Maskari et al.* analyzed how well classic IR evaluation metrics correlated with user satisfaction in recommender systems [298]. *Gunawardana and Shani* published a survey about accuracy metrics [225]. *Herlocker et al.* wrote an article on how to evaluate collaborative filtering approaches [226]. Various authors showed that offline and online evaluations often provide contradictory results [93], [270], [281]. Many papers about various aspects of recommender-system evaluation have been published [227], [269], [276], [281], [284], [299–301]. However, while many of the findings in these papers are important with respect to reproducibility, none of the authors mentioned or discussed their findings in the context of reproducibility.

The neglect of reproducibility in recommender-systems evaluation is also observed by *Ekstrand et al.* and *Konstan and Adomavicius*. They state that "it is currently difficult to reproduce and extend recommender systems research results," evaluations are "not handled consistently" [302], and many research papers "contribute little to collective knowledge," primarily due to non-reproducibility of the results [303]. They concluded:

> *"[T]he Recommender Systems research community is facing a crisis where a significant number of papers present results that contribute little to collective knowledge […] often because the research lacks the […] evaluation to be properly judged and, hence, to provide meaningful contributions."*

Not all researchers agree that the primary purpose of evaluations is to aid developers and researchers identify the most effective recommendation approaches. When we submitted a paper on reproducibility to the ACM RecSys conference observing that the evaluations were largely non-reproducible, one reviewer commented:

> *"I think that it is widely agreed in the community that this [non-reproducibility] is just the way things are – if you want a recsys for a specific application, there is no better way than just test and optimize a number of alternatives. This probably cannot be avoided – there will never be a sufficient set of experiments that would allow "practitioners" to make*

> *decisions without running through this optimization process for their specific app and dataset."*

When it comes to the lack of reproducibility, we hope that the research community can move beyond the observation that "*this is just the way things are*", given that reproducibility is a "bedrock principle in the conduct and validation of experimental science" [295]. The view of the reviewer above also leads to the question: How should the "number of [promising] alternatives" be determined? At least for research-paper recommender systems, there is no small number of promising alternatives; there are only *all* alternatives, because nearly all of the evaluated approaches were most effective in at least one evaluation. Practitioners could hardly implement all approaches to find the most effective approach for their scenario. Even if a few promising approaches were identified, how should they be optimized? There is no list of parameters that might be worth optimizing, and even if there were, there would probably be dozens of parameters, each with dozens or even hundreds of possible values, that would require testing. Again, this would hardly be feasible for someone who wanted to implement a recommender system. In addition, datasets and "specific features" of recommender systems change over time. What does this mean for the operation of a recommender system? Would the operator have to reevaluate the "number of alternatives" every time documents are added to the recommendation corpus, or whenever minor features of the recommender system were changed?

5 Survey of the Recommendation Classes

Aside from collaborative and content-based filtering, which were briefly introduced in Section 2, there are feature-based, knowledge-based, behavior-based, citation-based, context-based, ruse-based, and many more recommendations classes [133], [187], [301], [304–307]. We consider the following seven classes to be most appropriate for distinguishing the approaches in the field of research-paper recommender systems:

1. Stereotyping
2. Content-based Filtering
3. Collaborative Filtering
4. Co-Occurrence
5. Graph-based
6. Global Relevance
7. Hybrid

Originally, we planned to review the most promising approach or approaches of each recommendation class. However, as the review of the evaluations showed, most approaches were evaluated in ways making them nearly impossible to compare. Therefore, the most promising approaches could not be determined. Instead, we provide an overview of the most important aspects and techniques that have been used in the field. The analysis is based on the "in-depth" dataset, i.e. the 127 articles on 62 recommendation approaches that we classified as significant.

5.1 Stereotyping

Stereotyping is one of the earliest user modeling and recommendation classes. It was introduced by *Rich* in the recommender system *Grundy*, which recommended novels to its users [308]. *Rich* was inspired by stereotypes from psychology that allowed psychologists to quickly judge people based on a few characteristics. *Rich* defined stereotypes – which she called "facets" – as collections of characteristics. For instance, Grundy assumed that male 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" [308]. Consequently, Grundy recommended books that had been manually classified to match the facets.

One major problem with stereotypes is that they can pigeonhole users. While many men may have a negative interest in romance, this is not true for *all* men. In addition, building stereotypes is often labor intensive, since the items typically need to be manually

classified for each facet. This limits the number of items, for example, books that can reasonably be personalized [309].

Advocates of stereotype approaches argue that once the stereotypes are created the recommender system needs little computing power and may perform quite well in practice. For instance, *Weber and Castillo* observed that female users were usually searching for the composer Richard Wagner when they entered the search query 'Wagner' on *Yahoo!* [310]. In contrast, male users entering the same query were usually looking for the Wagner paint sprayer. *Weber and Castillo* modified the 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 [311]. Consequently, Orbitz assigned their 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, only *Beel et al.* applied stereotypes [312], [313]. The authors assume that all users of their reference-management software *Docear* are researchers or students. Hence, papers and books are recommended that are potentially interesting for researchers and students (for example, papers about optimizing scholarly literature for Google Scholar [314]). *Beel et al.* used stereotypes as a fallback model when other recommendation approaches could not deliver recommendations. They report mediocre performance of the stereotype approach with click-through rates (CTR) around 4%, while their content-based filtering approaches achieved CTRs over 6%.

5.2   Content-based Filtering

Content-based filtering (CBF) is one of the most widely used and researched recommendation approaches [263]. One central component of CBF is the user modeling process, in which the interests of users are inferred from the items that users interacted with. "Items" are usually textual, for instance emails [315] or webpages [316]. "Interaction" is typically established through actions, such as downloading, buying, authoring, or tagging an item. Items are represented by a content model containing the items' features. Features are typically word-based, i.e. single words, phrases, or n-grams. Some recommender systems also use non-textual features, such as writing style [317], [318], layout information [319], [320], and XML tags [321]. Typically, only the most descriptive features are used to model an item and users and these features are commonly weighted. Once the most discriminative features are identified, they are stored, often as a

vector that contains the features and their weights. The user model typically consists of the features of a user's items. To generate recommendations, the user model and recommendation candidates are compared, for example using the vector space model and the cosine similarity coefficient.

In the research-paper recommender-system community, CBF is the predominant approach: of the 62 reviewed approaches, 34 (55%) applied the idea of CBF [6], [51], [91]. "Interaction" between users and items was typically established through authorship [1], [115], [118], having papers in one's personal collection [6], [19], [60], adding social tags [29], or downloading [106], reading [124], and browsing papers [15], [48], [92].

Most of the reviewed approaches use plain words as features, although some use *n-grams* [29], [101], *topics* (words and word combinations that occurred as social tags on CiteULike) [61], and *concepts* that were inferred from the Anthology Reference Corpus (ACL ARC) via Latent Dirichlet Allocation [12], and assigned to papers through machine learning. A few approaches utilize non-textual features, and if they did then these non-textual features were typically utilized *in addition* to words. *Giles et al.* used citations in the same way as words were used and weighted the citations with the standard TF-IDF measure (they called this method *CC-IDF*) [43]. Others adopted the idea of CC-IDF or used it as a baseline [6], [14], [27]. However, *Beel* recently provided some initial evidence that CC-IDF might not be an ideal weighting scheme [312]. *Zarrinkalam and Kahani* considered *authors* as features and determined similarities by the number of authors two items share [125].

The approaches extracted words from the title [79], [85], [112], abstract [26], [51], [61], header [43], introduction [57], foreword [57], author-provided keywords [26], [57], [58], and bibliography [27], as well as from the papers' body text [65], [101], [112]. The approaches further extracted words from external sources, such as social tags [58], [62], ACM classification tree and DMOZ categories [91], [94], and citation context [51], [55], [65]. Using citation context is similar to the way search engines use anchor analysis for webpage indexing since the 1990's [322], [323]. Citation context analysis was also used in academic search before it was used by research-paper recommender systems [324].

Words from different document fields have different discriminative powers [325]. For instance, a word occurring in the title is usually more meaningful than a word occurring in the body text. *Nascimento et al.* accounted for this and weighted terms from the title three times stronger than terms from the text body, and text from the abstract twice as strong [101]. This weighting scheme was arbitrarily selected and not based on empirical evidence. *Huang et al.* experimented with different weights for papers' content and citation context [55]. They found that an equal weight for both fields achieved the highest precision. The other reviewed approaches that used text from different fields did not report on any field weighting.

The most popular model to store item representations and user models was the Vector Space Model (VSM), which was used by 9 (64%) of those 14 approaches that reported how they stored the user and item models. Other approaches modeled their users as a graph [102], [103], [122], as a list of topics that were assigned through machine learning [91], or as an ACM hierarchy [66]. Of those who used the VSM, all but one used the cosine measure to calculate similarities between user models and recommendation candidates. In 1998, *Giles et al.* compared headers of documents with a string distance measure [14], but neither they nor others mentioned this technique again, which might imply that string edit distance was not effective.

TF-IDF was the most popular weighting scheme (70%) among those approaches for which the scheme was specified. Other weighting schemes included plain term frequency (TF) [29], [101], [115], and techniques that the authors called "phrase depth" and "life span" [29].

CBF has a number of advantages compared to stereotypes. CBF allows a user-based personalization so the recommender system can determine the best recommendations for each user individually, rather than being limited by stereotypes. CBF also requires less up-front classification work, since user models can be created automatically.

On the downside, content-based filtering requires more computing power than stereotyping. Each item must be analyzed for its features, user models must be built, and similarity calculations must be performed. If there are many users and many items, these calculations require significant resources. The weakness of content-based filtering is its low serendipity and overspecialization leading it to recommend items as similar as possible to the ones a user already knows [263]. Content-based filtering also ignores quality and popularity of items [24]. For instance, two research papers may be considered equally relevant by a CBF recommender system if the papers share the same terms with the user model. This relevance might not always be justified, for example if one paper was written by an authority in the field and presents original results, while another paper was written by a student who paraphrases the results of other research papers. Ideally, a recommender system should recommend only the first paper but a CBF system would fail to do so. Another criticism of content-based filtering is that it is dependent on access to the item's features [24]. For research-paper recommendations, usually PDFs must be processed and converted to text, document fields must be identified, and features, such as terms must be extracted. None of these tasks are trivial and they may introduce errors into the recommendations [257], [326], [327].

5.3   Collaborative Filtering

The term "collaborative filtering" (CF) was coined in 1992 by *Goldberg et al.*, who proposed that "information filtering can be more effective when humans are involved in the filtering process" [328]. The concept of collaborative filtering as it is understood today was introduced two years later by *Resnick et al.* [222]. Their theory was that users like what like-minded users like, where two users were considered like-minded when they rated items alike. When like-minded users were identified, items that one user rated positively were recommended to the other user, and vice versa. Compared to CBF, CF offers three advantages. First, CF is content independent, i.e. no error-prone item processing is required [117], [265], [287]. Second, because humans do the ratings, CF takes into account real quality assessments [24]. Finally, CF is supposed to provide serendipitous recommendations because recommendations are not based on *item* similarity but on *user* similarity [87], [287].

From the reviewed approaches, only 11 (18%) applied collaborative filtering [93], [106], [119]. *Yang et al.* intended to let users rate research papers, but users were "too lazy to provide ratings" [124]. *Naak et al.* faced the same problem and created artificial ratings for their evaluation [98]. This illustrates one of the main problems of CF: CF requires user participation, but often the motivation to participate is low. This problem is referred to as the "cold-start" problem, which may occur in three situations [265]: new users, new items, and new communities or disciplines. If a new user rates few or no items, the system cannot find like-minded users and therefore cannot provide recommendations. If an item is new in the system and has not been rated yet by at least one user, it cannot be recommended. In a new community, no users have rated items, so no recommendations can be made and as a result, the incentive for users to rate items is low.

To overcome the cold-start problem, implicit ratings may be inferred from the interactions between users and items. *Yang et al.* inferred implicit ratings from the number of pages the users read: the more pages users read, the more the users were assumed to like the documents [124]. *Pennock et al.* interpreted interactions, such as downloading a paper, adding it to ones' profile, editing paper details, and viewing its bibliography as positive votes [106]. *McNee et al.* assumed that an author's citations indicate a positive vote for a paper [93]. They postulated that when two authors cite the same papers, they are like-minded. Similar, if a user reads or cites a paper the citations of the cited paper are supposed to be liked by the user.

Using *inferred* ratings voids CF's advantage of being based on real user quality assessments. This criticism applies to citations as well as to other types of implicit ratings [329–331]. For example, we cite papers in this survey that had inadequate evaluations, or were written in barely understandable English. Thus, interpreting citations always as a positive vote can be misguiding. Similarly, when a user spends a lot of time reading a

paper this *could* mean that the paper contains interesting information that the user would rate positively; but it could also mean that the paper is difficult to understand and requires a lot of effort to read. Consequently, CF's advantage of explicit human quality assessments mostly vanishes when implicit ratings are used.

Using citations as inferred ratings might also void CF's second advantage of being content-independent. Typically, reliable citation data is not widely available. Therefore, access to the paper's content is required to build a citation network, but this process is even more fault-prone than word extraction in CBF. In CBF, the text of the papers must be extracted, and maybe fields such as the title or abstract must be identified. For citation-based CF, the text must also be extracted but in the text, the bibliography and its individual references must be identified, including their various fields including title and author. This is an error-prone task [326].

A general problem of collaborative filtering in the domain of research-paper recommender systems is sparsity. *Vellino* compared the implicit ratings on Mendeley (research papers) and Netflix (movies), and found that sparsity on Netflix was three orders of magnitude lower than on Mendeley [212]. This is caused by the different ratio of users and items. In domains like movie recommendations, there are typically few items and many users. For instance, the movie recommender MovieLens has 65,000 users and 5,000 movies [226]. Typically, many users watched the same movies. Therefore, like-minded users can be found for most users and recommendations can be given effectively. Similarly, most movies have been watched by at least some users and hence most movies can be recommended. The situation is different in the domain of research papers. There are typically few users but millions of papers, and very few users have rated the same papers. Finding like-minded users is often not possible. In addition, many papers are not rated by any users and therefore cannot be recommended.

There are further critiques of CF. Computing time for CF tends to be higher than for content-based filtering [265]. Collaborative filtering is generally less scalable and requires more offline data processing than CBF [332]. *Torres et al.* note that collaborative filtering creates similar users [117] and *Sundar et al.* observe that collaborative filtering dictates opinions [333]. *Lops* makes the criticism that collaborative filtering systems are black boxes that cannot explain why an item is recommended except that other users liked it [263]. Manipulation is also a problem: since collaborative filtering is based on user opinions, blackguards might try to manipulate ratings to promote their products so they are recommended more often [334–336].

5.4   Co-occurrence recommendations

To give *co-occurrence* recommendations, those items are recommended that frequently co-occur with some source items. One of the first applications of co-occurrence was co-citation analysis introduced by *Small* in 1973 [234]. *Small* proposed that two papers are more related to each other, the more often they are co-cited. Many others adopted this concept, the most popular example being Amazon's "*Customers Who Bought This Item Also Bought....*" Amazon analyzes which items are frequently bought together, and when a customer browses a product, items frequently bought with that item are recommended.

One advantage of co-occurrence recommendations is the focus on relatedness instead of similarity. Similarity expresses how many features two items have in common. Recommending similar items, as CBF is doing, is often not ideal because similar items are not serendipitous [337]. In contrast, *relatedness* expresses how closely coupled two items are, not necessarily dependent on their features. For instance, two papers sharing the same features (words) are similar. In contrast, paper and pen are not similar but related, because both are required for writing letters. Hence, co-occurrence recommendations provide more serendipitous recommendations, and in this way are comparable to collaborative filtering. In addition, no access to content is needed and complexity is rather low. It is also rather easy to generate anonymous recommendations, and hence to assure users' privacy. On the downside, recommendations are not highly personalized and items can only be recommended if they co-occur at least once with another item.

Six of the reviewed approaches are based on co-occurrences (10%). Three of those approaches analyze how often papers are *co-viewed* during a browsing session [15], [48], [92]. Whenever a user views a paper, those papers that were frequently co-viewed with the browsed paper are recommended. Another approach uses proximity of co-citations to calculate document relatedness [44]: the closer the proximity of two references within a paper, the more related the cited papers are assumed to be. *Pohl et al.* compared the effectiveness of co-citations and co-downloads and found that co-downloads are more effective than co-citations only in the first two years after a paper is published [110].

Calculating co-occurrence recommendations is not always feasible. For instance, on arXiv.org, two thirds of all papers have no co-citations, and those that do usually have no more than one or two [110]. Despite its limitations, co-occurrence recommendations seem to perform quite well. Two popular research-paper recommender systems, bX and BibTip, both rely on co-occurrence recommendations and deliver millions of recommendations every month [15], [92].

5.5   Graph Based

Ten of the reviewed approaches utilize the inherent connections that exist in academia (16%). Based on these connections, the approaches build graph networks that typically show how papers are connected through citations [9], [72], [84]. Sometimes, graphs include authors [3], [79], [127], users/customers [57], venues [9], [79], [127], genes and proteins [3], [79], and the years the papers were published [79]. *Lao et al.* even included terms from the papers' titles in the graph, which makes their approach a mixture of the graph and content based approach [79]. Depending on the entities in the graph, connections can be citations [9], [79], [84], purchases [57], "published in" relations, [9], [79], [127], authorship [3], [9], [127], relatedness between genes[29] [3], or occurrences of genes in papers [79]. Some authors connected entities based on non-inherent relations. For instance, *Huang et al.* and *Woodruff et al.* calculated text similarities between items and used the text similarity to connect papers in the graph [57], [123]. Other connections were based on attribute similarity[30], bibliographic coupling, co-citation strength [57], [123], [127], or demographic similarity [57]. Once a graph was built, graph metrics were used to find recommendation candidates. Typically one or several input papers were given and from this input random walks with restarts were conducted to find the most popular items in the graph [46], [72], [79].

5.6   Global Relevance

In its simplest form, a recommender system adopts a one-fits-all approach and recommends items that have the highest global relevance. In this case, the relevance is not calculated specific to a user, for example based on the similarity of user models and recommendation candidates. Instead, some global measures are used, such as overall popularity. For instance, a movie-rental system could recommend those movies that were most often rented or that had the highest average rating over all users. In this case, the basic assumption would be that users like what most other users like.

From the reviewed approaches, none used global relevance exclusively but many used it as an additional ranking factor. For instance, five CBF approaches used global popularity metrics in their rankings [12], [51], [125]. They first determined a list of recommendation candidates with a user-specific CBF approach. Then, the recommendation candidates were re-ranked based on the global relevance metrics. Popular metrics were PageRank [12], HITS [51], Katz metric [51], citation counts [12], [51], [112], venues' citation counts [12], [112], citation counts of the authors' affiliations [112], authors' citation count [12],

---

[29] Relatedness between genes was retrieved from an external data source that maintained information about gene relatedness.

[30] Attribute similarity was calculated, e.g., based on the number of pages.

[112], h-index [12], recency of articles [12], title length [112], number of co-authors [112], number of affiliations [112], and venue type [112].

Strohman et al. report that the Katz metric, which quantifies relevance as a function of the paths between two nodes (the shorter the paths the higher the relevance), strongly improved precision [114]. All variations that included Katz were about twice as good as those variations without. *Bethard and Jurafsky* report that a simple citation count was the most important factor, and age (recency) and h-index were even counterproductive [12]. They also report that considering these simple metrics doubled mean average precision compared to a standard content-based filtering approach.

5.7   Hybrid Recommendation Approaches

Approaches of the previously introduced recommendation classes may be combined in hybrid approaches. Many of the reviewed approaches have some hybrid characteristics. For instance, several of the CBF approaches use global relevance attributes to rank the candidates, or graph methods are used to extend or restrict potential recommendation candidates. This type of hybrid recommendation technique is called "feature augmentation" [304]. It is a weak form of hybrid recommendation technique, since the primary technique is still dominant. In true hybrids, the combined concepts are similarly important [304], [338]. From the reviewed approaches, only some of the TechLens approaches may be considered true hybrid approaches.

TechLens [26], [63], [64], [87], [93], [117] is one of the most influential research-paper recommender systems, although it was not the first like some have claimed (e.g. [109]). TechLens was developed by the GroupLens[31] team. Currently TechLens is not publicly available, although the GroupLens team is still active in the development and research of recommender systems in other fields. Between 2002 and 2010, Konstan, Riedel, McNee, Torres, and several others published six articles related to research-paper recommender systems. Often, *McNee et al.'s* article from 2002 is considered to be the original TechLens article [93]. However, the 2002 article 'only' introduced some algorithms for recommending citations, which severed as a foundation for TechLens, which was introduced in 2004 by *Torres et al.* [117]. Two articles about TechLens followed in 2005 and 2007 but added nothing new with respect to recommendations [63], [64]. In 2006, *McNee et al.* analyzed potential pitfalls of recommender systems [87]. In 2010, *Ekstrand et al.* published another article on the approaches of TechLens and suggested enhancements for the approaches [26].

---

[31] http://grouplens.org/

TechLens' algorithms were adopted from Robin Burke [304] and consisted of three CBF variations, two CF variations, and five hybrid approaches.

> Content-Based Filtering: *Pure-CBF* served as a baseline in the form of standard CBF in which a term-based user model was compared with the recommendation candidates. In the case of TechLens, terms from a single input paper were used. In *CBF-Separated*, for each paper being cited by the input paper, similar papers are determined separately and at the end the different recommendation lists are merged and presented to the user. In *CBF-Combined*, terms of the input paper and terms of all papers being cited by the input paper are combined in the user model. Then the papers most similar to this user model are recommended.

> Collaborative Filtering: *Pure-CF* served as another baseline and represented the collaborative filtering approach from *McNee et al.*, in which papers were interpreted as users and citations were interpreted as votes [93]. In *Denser-CF*, citations of the input paper were additionally included in the user model.

> Hybrid: With *Pure-CF->CBF Separated*, recommendations were first created with Pure-CF. These recommendations were then used as input documents for CBF-Separated. Similarly, *Pure-CF->CBF Combined*, *CBF Separated->Pure-CF*, and *CBF-Combined->Pure-CF* were used to generate recommendations. *Fusion* created recommendations with both CBF and CF independently and then merged the recommendation lists.

Despite various evaluations of the approaches, it remains unclear which are most promising (refer to the explanations in Section 4.7).

6 Survey of the Research Field and Shortcomings

Our survey already revealed that there are some shortcomings and challenges in the field of research-paper recommender systems. This is especially the case when it comes to evaluations, which are often non-reproducible and incomparable. However, during the review, we identified more challenges. In the next sections, we introduce these challenges, and hope to stimulate a discussion about future research directions to enhance research-paper recommender systems.

6.1 Neglect of User Modeling

A fundamental part of generating recommendations is the user modeling process that identifies a user's information needs [264]. Ideally, a recommender system identifies the needs automatically by inferring the needs from the user's item interactions. Alternatively, the recommender system asks users to specify their needs by providing a list of keywords or through some other method. However, in this case a recommender system becomes very much like a search engine and loses one of its main features, namely the capability to recommend items even if users do not know exactly what they need.

Of the 62 reviewed approaches, 50 (81%) required users to either explicitly provide keywords [52], [98], [119], or to provide text snippets (e.g. an abstract) [12], [111], [125], or to provide a single paper as input [26], [101], [114], or the approaches ignored the user modeling process entirely. These approaches neglect one of the most important parts of a recommender system, which makes the approaches very similar to classic search, or related document search [339–341], where users provide search terms or an input paper, and receive a list of search results or similar papers. Of course, neither classic search nor related-document search are trivial tasks in themselves, but they neglect the user modeling process and we see little reason to label such systems as recommender systems.

Of the reviewed approaches, 12 (19%) inferred information from the items the users interacted with. However, most approaches that inferred information automatically used *all* papers that a user authored, downloaded, etc. [62], [94], [122]. This is not ideal. When inferring information automatically, a recommender system should determine the items that are currently relevant for the user-modeling process [342]. For example, ten-year-old user-authored papers are probably not suitable to describe a user's current information needs. This aspect is called "concept drift" and it is important for creating meaningful user models. In the research-paper recommender systems community, concept drift is widely ignored: a mere three approaches considered concept drift in detail. *Middleton et al.* weight papers by the number of days since the user last accessed them [91]. *Watanabe et al.* use a similar approach [122]. *Sugiyama and Kan*, who use the user's authored

papers, weight each paper based on the difference between a paper's publication year, and the year of the most recently authored paper [115]. In addition, they found that it makes sense to include only those papers that the user authored in the past three years [115].

Another important question in user modeling is the user-model size. While in search, user models (i.e. search queries) typically consist of a few words; user models in recommender systems may consist of hundreds or even thousands of words. Of the reviewed CBF approaches, 30 (88%) did not report the user-model size. The few authors who reported the user-model size, usually stored fewer than 100 terms. For instance, *Giles et al.* made us of the top 20 words of the papers [43].

6.2   Focus on Accuracy

The research-paper recommender-system community places a strong focus on accuracy, and seems to assume that an accurate recommender system will lead to high user satisfaction. However, outside the research-paper recommender-system community it is agreed that many aspects beyond accuracy affect user satisfaction. For instance, users might become dissatisfied with accurate recommendations when they have no trust in the recommender system's operator [343], their privacy is not ensured [301], they need to wait too long for recommendations [301], or they find the user interfaces unappealing [344]. Other factors that affect user satisfaction are confidence in a recommender system [264], data security [345], diversity [346], user tasks [87], item's lifespan [347] and novelty [348], risk of accepting recommendations [349], robustness against spam and fraud [350], transparency and explanations [351], time to first recommendation [226], and interoperability [352].

Among the reviewed articles, a few authors considered aspects beyond accuracy, as shown in the following sections.

6.2.1   Users' Tasks

Torres et al. from TechLens' considered a user's current task in the recommendation process. The authors distinguished between users who wanted to receive authoritative recommendations and novel recommendations [117]. They showed that different recommendation approaches were differently effective for these tasks. The developers of *TheAdvisor* let users specify whether they are interested in classic or recent papers [73]. *Uchiyama et al.* found that students are typically not interested in finding papers that are "similar" to their input paper [118]. This finding is interesting because content-based filtering is based on the assumption that user want similar papers. However, the study by

Uchiyama et al. was based on 16 participants. As such, it remains uncertain how significant the results are.

6.2.2 Diversity

Diversity of recommendations was mentioned in a few articles, but considered in depth only by two author groups (*Vellino et al.* and *Küçüktunç et al.*). *Vellino et al.* measured diversity as the number of journals from which articles were recommended [119]. If recommendations were all from the same journals, diversity was zero. They compared diversity of a CF approach with the co-occurrence approach from bX and found that CF had a diversity of 60% while diversity of bX was 34%. *Küçüktunç et al.* from TheAdvisor published two articles about diversity in research-paper recommender systems [70], [71]. They provided a survey on diversification techniques in graphs and proposed some new techniques to measure diversity.

6.2.3 Layout

Farooq et al. from CiteSeer analyzed which information users wanted to see when receiving recommendations in RSS feeds [138]. They found that the information to display varies for the type of recommendation. In one approach, *Farooq et al.* recommended papers that cited the user's papers. In this case, users preferred to see the citing paper's bibliographic data (e.g. title, author) and the context of the citation – the sentence in which the citation appeared. When papers were recommended that were co-cited with the users' papers, citation context was not that important. Rather, the users preferred to see the bibliographic data and abstract of the co-cited paper. When papers were recommended that had a similar content to the users' papers, users preferred to see bibliographic data and abstract. These findings are interesting because from the reviewed recommender systems the majority display only the title and not the abstract.

As part of our work, we researched the impact of labeling and found that papers labeled as 'sponsored recommendation' performed worse than recommendations with a label that indicated that the recommendations were 'organic,' though the recommended papers were identical [4]. It also made a difference if paper recommendations were labeled as 'Sponsored' or 'Advertisement' although both labels indicate that recommendations are displayed for commercial reasons.

6.2.4 User Characteristics

For our own recommender system Docear [6], [10], [11], [353], we found that researchers who registered tended to have higher click-through rates than unregistered users (6.95% vs. 4.97%) [7], [354]. In addition, older users seem to have higher average click-through rates (40-44 years: 8.46%) than younger users (20-24 years: 2.73%) [7]. *Middleton et al.*

also report differences for different user groups. Click-through rates in their recommender system, Quickstep, was around 9%, but only around 3.5% for Foxtrot, although both systems applied very similar approaches. However, Quickstep users were recruited from a computer science laboratory, while Foxtrot was a real-world system offered to 260 faculty members and students (although only 14% of users used Foxtrot at least three times).

Click-through rates from the bX recommender are also interesting [116]. They varied between 3% and 10% depending on the university in which recommendations were shown (bX provides more than 1,000 institutions with recommendations) [25]. This could have been caused by different layouts of the recommendations, but it might also have been caused by the different backgrounds of students.

6.2.5 Usage Duration

Middleton et al. reported that the longer someone used the recommender system, the lower click-through rates became [94]. *Jack* reports the opposite, namely that precision increased over time (p=0.025 in the beginning, p=0.4 after six months) and depended on a user's library size (p=0.08 for 20 articles and p=0.40 for 140 articles) [58]. We showed that it might make sense to be "persistent" and show the same recommendations to the same users multiple times – even recommendations that users had clicked before were often clicked again [5].

6.2.6 Recommendation Medium

User satisfaction also depends on the medium through which recommendations are made. *Middleton et al.* report that recommendations via email received half the click-through rate as the same recommendations delivered via a website [94]. Of the reviewed recommender systems, only Docear [6] and Mendeley [58] provide recommendations through a desktop software; CiteSeer provided recommendations in a news feed [138]; and all others deliver their recommendations through websites. If and how click rates differ when recommendations are delivered via desktop software compared to a website remains unknown.

6.2.7 Relevance and Profile Feedback

Relevance feedback is a common technique to improve recommendations [264] but it is widely ignored in the research-paper recommender-system community. *Middleton et al.* showed that profile feedback is better than relevance feedback: allowing users to edit their user models is more effective than just learning from relevance feedback [94]. *Bollacker et al.* from CiteSeer allowed their users to edit their profiles but conducted no research on the effectiveness of this feature [80].

6.3   Translating Research into Practice

Translating research into practice is a current challenge in the research paper recommender system community. Out of the large number of proposed approaches in the field, 24 research-paper recommender systems could be used by users in practice (Table 9)[32]. Of these 24 recommender systems, eight (33%) never left the prototyping stage – and today only one of the prototypes is still publicly available. Of the remaining recommender systems, four are offline (25%), five are no longer actively maintained (31%)[33], while seven are running and actively maintained (44%). Of the seven active recommender systems, four operators are involved with the recommender-system research community[32] and publish information about their systems.

Table 9: Overview of recommender systems surveyed

Status	Name	Maturity	Research Oriented	Type	Presentation
Active	**BibTip**	Real System	No[1]	Stand-Alone	Webpage
Active	**bx**	Real System	No	Stand-Alone	Webpage
Active	**Docear**	Real System	Yes	On-Top	Software
Active	**Mendeley**	--	--	--	--
	Related Papers	Real System	Yes	On-Top	Software
	Suggest	Real System	Yes	On-Top	Software
Active	**RefSeer**	Real System	Yes	Stand-Alone	Webpage
Active	**Scholar Update**	Real System	No	On-Top	Webpage
Idle	**CiteULike**	--	--	--	--
	CF	Real System	No	On-Top	Webpage
	Item-Centric	Real System	No	On-Top	Webpage
Idle	**PubMed PRMA**	Real System	No	On-Top	Webpage
Idle	**ResearchGate**	Real System	No	On-Top	Webpage
Idle	**TheAdvisor**	Real System	Yes	Stand-Alone	Webpage
Idle	**Who Should I Cite?**	Prototype	Yes	Stand-Alone	Webpage
Offline	**CiteSeer**	--	--	--	--
	Alert	Real System	Yes	On-Top	Feed
	Related Documents	Real System	Yes	On-Top	Webpage
Offline	**Foxtrot**	Real System	Yes	Stand-Alone	Webpage, Email
Offline	**TechLens**	Real System	Yes	Stand-Alone	Webpage
Offline	**NSYSU-ETD**	Prototype	Yes	On-Top	Webpage
Offline	**OSUSUME**	Prototype	Yes	On-Top	?
Offline	**Papits**	Prototype	Yes	On-Top	Webpage
Offline	**Papyres**	Prototype	Yes	On-Top	?
Offline	**Pirates**	Prototype	Yes	Stand-Alone	?
Offline	**Quickstep**	Prototype	Yes	Stand-Alone	Webpage
Offline	**Sarkanto & Synthese**	Prototype	Yes	Stand-Alone	Webpage

---

[32] The recommender systems of Mendeley, CiteULike, and CiteSeer are counted twice because they offer or offered two independent recommender systems.

[33] We classified a recommender system as not actively maintained if no article was published or no changes were made to the system for a year.

Most of the real-world recommender systems apply simple recommendation approaches that are not based on recent research. For instance, as far as we could tell, PubMed was still using an approach introduced in 2007; ResearchGate is using a simple content-based filtering approach similar to classic search[34]; CiteULike apparently uses two approaches from 2008/2009; and BibTip and bX are using simple co-occurrence approaches. Whether the RefSeer system applies all the results from their research remains unclear. In other words, the reviewed research typically did not affect real-world recommender systems.

6.4   Persistence and Authorities

One reason the research is not transferred directly into practice might be a lack of persistence and authorities in the field. Of the 276 authors who authored the 185 articles, 201 (73%) published just one article (Figure 5). 15 authors (5%) published five or more articles, but of these authors, several were co-authors publishing the same articles. This means, there are only a few groups that consistently publish research in the field of research-paper recommender systems.
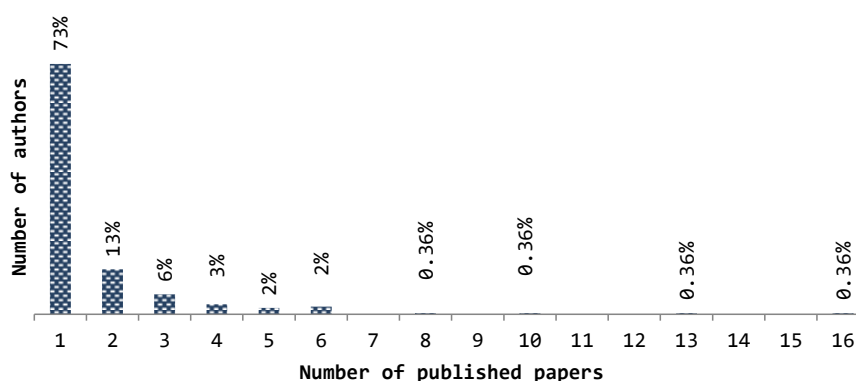
Figure 5: Papers published per author in the research paper recommender field

The most productive authors are C. Lee Giles and his co-author P. Mitra from CiteSeer/RefSeer (Table 10 and Table 11). No other authors have published as many articles (17) over as long a period of time (16 years) on as many different aspects of research-paper recommender systems. Other productive authors are A. Geyer-Schulz and his co-authors M. Hahsler, and M. Jahn from BibTip. They published 13 articles, but these were less often cited in the community than those of *Giles et al*. The articles are also narrower in scope than those of the CiteSeer authors. Our research group, i.e. J. Beel, S. Langer, M. Genzmehr, and B. Gipp from Docear, authored eight papers between 2009 and 2013, including posters or short papers. The research concentrated on aspects beyond

---

[34] ResearchGate also applied other recommender systems, e.g. for people or news, and it seems that these approaches are more sophisticated.

accuracy, such as the impact of labeling recommendations and the impact of demographics on click-through rates. O. Küçüktunç and his co-authors E. Saule and K. Kaya from TheAdvisor published six articles focusing on diversity and graph-based recommendations. J. A. Konstan, S. M. McNee, R. Torres, and J.T. Riedel, who are highly recognized authors in the field of recommender systems, developed TechLens and authored six articles relating to research-paper recommender systems during 2002 and 2010. Two of their articles influenced the work of several others and are among the most cited articles we reviewed [93], [117]. W. W. Cohen, and his PhD student N. Lao, are also two productive authors. They authored six articles from 2008 to 2012 (some of which are unpublished). SE. Middleton and his co-authors published five articles. It is interesting to note that five of the six most productive research groups have access to real-world recommender systems.

Table 10: Most productive authors in the research paper recommender field

Author	Papers
C. Lee Giles	16
A. Geyer-Schulz	13
M. Hahsler	10
J. Beel	8
O. Küçüktunç	6
E. Saule	6
K. Kaya	6
P. Mitra	6
J. A. Konstan	6
W. W. Cohen	6
S. Langer	5
M. Genzmehr	5
M. Jahn	5
S.E. Middleton	5
D.C. De Raure	5

Table 11: Most productive author-groups

Groups	Papers
C. Lee Giles; P. Mitra (CiteSeer/RefSeer)	17
A. Geyer-Schulz; M. Hahsler; M. Jahn (BibTip)	13
J. Beel; S. Langer, M. Genzmehr, B. Gipp (Docear)	8
J. A. Konstan; S.M. McNee; R. Torres, J.T. Riedl (TechLens)	6
O. Küçüktunç; E. Saule; K. Kaya (TheAdvisor)	6
W. W. Cohen; N. Lao	6
S.E. Middleton; D.C. De Raure, N.R. Shadbolt (Quickstep & Foxtrot)	5

6.5 Cooperation

Most articles were authored by multiple authors: the majority of articles had two (26%), three (26%) or four authors (18%) (Figure 6)[35]. 17% of articles were authored by a single researcher. On first glance, these numbers indicate a high degree of collaboration. However, we noticed that between different co-author groups little cooperation took place. The closest cooperation we could identify was that *Giles* was part of a committee for a thesis that *Cohen* supervised [74]. Leading authors of different research groups did not co-author articles together.

Figure 6: Number of authors of the reviewed papers

Co-author groups frequently seemed to work alone and did not always build on the results of the work done by their peers. Among the reviewed articles, it rarely happened that authors reported to have built their novel approach based upon an existing approach. This lack of cooperation also becomes apparent when looking at the citations. Although some of the reviewed articles gained many citations, these citations usually resulted from articles outside the research-paper recommender domain. For instance, the paper "Learning multiple graphs for document recommendations" attracted 63 citations since 2008 [127]. From these citations, only three were made by the reviewed articles. Another article, from the BibTiP developers, gained 24 citations since 2002 [32]. From the 24 citations, ten were self-citations and none was from the reviewed articles. Both examples are typical for most of the reviewed articles. One of the few articles that is constantly cited in the research-paper recommender community is an article about TechLens, which accumulated more than 100 citations [117]. However, many authors cited the article for

---

[35] Median author count was three, maximum count eleven.

authoritative reasons. In the citing papers, TechLens is mentioned but, with few exceptions, its approaches are neither adopted nor used as a baseline.

6.6   Information Scarcity

Most authors provided little detail about their approaches, which makes a re-implementation difficult, if not impossible. For instance, for 24 of the 34 content-based filtering approaches (71%), the authors did not report the weighting scheme they used (e.g. TF-IDF). The feature representation model (e.g. vector space model) was not reported for 20 approaches (59%). Whether stop words were removed was not reported for 23 approaches (68%). For 16 approaches (47%) no information was given on the fields the terms were extracted from (e.g. title or abstract). This information scarcity means, when an evaluation reports promising results for an approach, other researchers would not know how to re-implement the approach in detail. If they tried, and guessed the specifics of an approach, the outcome would probably differ significantly from the original. This might cause problems in replicating evaluations, and reproducing research results and hinders the re-implementation and application of promising approaches in real-word recommender systems.

7 Summary and Outlook

In the 16 years from 1998 to 2013 more than 200 research articles were published in the field of research-paper recommender systems. The articles consisted primarily of peer-reviewed conference papers (59%), journal articles (16%), pre-prints (5%), and other documents such as presentations and web pages (15%). The few existing literature surveys in this field cover only a fraction of these articles, which is why we conducted a comprehensive survey of research-paper recommender systems. The review revealed the following information.

Content-based filtering (CBF) is the predominant recommendation class. From 62 reviewed approaches, 34 used CBF (55%). From these CBF approaches, the majority utilized plain terms contained in the documents. Some used n-grams, or topics based on LDA. A few approaches also utilized non-textual features, such as citations or authors. The most popular model to store item representations was the Vector Space Model. Other approaches modeled their users as graphs, as lists with topics that were assigned through machine learning, or as ACM classification hierarchies. The reviewed approaches extracted text from the title, abstract, header, introduction, foreword, author-provided keywords, bibliography, body text, social tags, and citation context.

Only eleven approaches applied collaborative filtering, and none of them successfully used explicit ratings. *Yang et al.* intended to develop such a system, but their users were "too lazy to provide ratings" [124]. Hence, implicit instead of explicit ratings were used. Implicit ratings were inferred from the number of pages the users read, users' interaction with the papers (e.g. downloads, edits, views) and citations. The main problem of collaborative filtering for research papers seems to be scarcity. *Vellino* compared implicit ratings on Mendeley (research papers) and Netflix (movies), and found that scarcity on Mendeley differs from Netflix by a magnitude of three.

Although stereotyping is one of the oldest user modeling approaches, and is successfully applied by services such as Yahoo!, just one of the reviewed approaches applied stereotypes as a fallback model when other recommendation approaches could not deliver recommendations. The authors reported reasonable performance of the stereotype approach with click-through rates (CTR) around 4% while the CBF approach achieves CTRs around 6%. Given these results, we believe that a more thorough analysis of stereotype recommendations could be interesting.

Six of the reviewed approaches were co-occurrence based. Three approaches analyzed how often papers were *co-viewed* during a browsing session: whenever users browsed a paper, the system recommended those papers that had frequently been co-viewed with the browsed paper in the past. Another approach used co-citations to calculate document relatedness. The higher the proximity of two references within a paper, the more related

they are assumed to be. *Pohl et al.* compared the effectiveness of co-citations and co-downloads and found that co-downloads are more effective than co-citations only in the first two years after a paper was published [110].

Ten recommendation approaches built graphs to generate recommendations. Such graphs typically included papers that were connected via citations. Some graphs included authors, users/customers, venues, genes and proteins, and publishing years of the papers. *Lao et al.* even included terms from the papers' titles in the graph. Depending on the entities in the graph, connections were citations, purchases, "published in" relations, authorship, relatedness between genes, and occurrences of genes in papers. Some authors connected the entities based on non-inherent relations. For instance, *Huang et al.* and *Woodruff et al.* calculated text similarities between items and used the text similarity to connect papers in the graph. Other connections were based on attribute similarity, bibliographic coupling, co-citation strength, and demographic similarity.

Despite a lot of research about research-paper recommender systems, we identified two main problems in the research field.

First, it is currently not possible to determine the most effective recommendation approaches. If we were asked which recommendation approach to apply in practice or to use as baseline, there is no definite answer. We do not even have a clue as to which of the approaches might be most promising. This problem mainly relates to poor experimental design and lack of information, which includes inadequate evaluations and too little information given by the authors: 22% of the approaches were not evaluated. Of the remaining approaches, 20% were not evaluated against a baseline; the majority of the remaining approaches were compared to simple baselines but not to approaches of other researchers in the field. The majority (71%) of approaches were evaluated using offline evaluations, which are subject to various shortcomings. Some claim that offline evaluations should not be used for evaluating research-paper recommender systems [192]. If that is true, most of the reviewed evaluations would be of little significance. Even if this criticism is unjustified, some problems remain. Many authors pruned their offline datasets in drastic ways. For instance, *Pennock et al.* removed all documents with fewer than 15 implicit ratings from the corpus. Therefore, 1,575 papers remained from the original 270,000 (0.58%). Results based on such datasets do not allow drawing reliable conclusions about how the approaches might perform in real-world recommender systems. In addition, the majority of the user studies (58%) had 15 or less participants, which questions the significance of these evaluations. Only 8% of the approaches were evaluated with online evaluations in real-world recommender systems with real users. Additionally, many authors provided little information about their approaches, which hinders re-implementation. For instance, most authors did not report on the text fields they utilized, or which weighting schemes were used.

To solve this problem, we believe it is crucial that the community discusses and develops frameworks and best-practice guidelines for the evaluation of research-paper recommender-systems. This should include an analysis and discussion of how suitable offline evaluations are; to what extent datasets should be pruned; the creation of datasets comparable to existing TREC datasets; the minimum number of participants in user studies; and which factors influence the results of evaluations (e.g. user demographics). Ideally, a set of reference approaches would be implemented that could be used as baselines. In addition, more details on implementations are needed, based on a discussion of the information needed in research articles.

It is crucial to find out why seemingly minor differences in algorithms or evaluations lead to major variations in the evaluation results. As long as the reasons for these variations are not found, scholars cannot rely on existing research results because it is not clear whether the results can be reproduced in a new recommendation scenario.

Second, we identified unused potential in recommender systems research. This problem has two root causes.

A) Research results are often not transferred into practice, or considered by peers. Despite the large number of research articles, just a handful of active recommender systems exist, and most of them apply simple recommendation approaches that are not based on recent research results. As such, the extensive research conducted from 1998 to 2013 apparently had a rather minor impact on research-paper recommender systems in the real world. Additionally, developers of several of the active recommender systems do not engage in the research community or publish information about their systems. Some researchers also seem to be unaware of developments in related research domains, such as user modeling, scientometrics, and the reviewer-assignment problem. In addition, the major co-author groups in the domain of research-paper recommender systems do not cooperate much with each other. One reason for some of these problems might be a relatively short-lived interest in the research field. Most authors (73%) published only a single paper on research-paper recommender systems.

B) The majority of authors did not take into account that user satisfaction might depend not only on accuracy but also on factors such as privacy, data security, diversity, serendipity, labeling, and presentation. The operator perspective was widely neglected. Information about runtime was provided for 10% of the approaches. Complexity was covered by very few authors and the costs of running a recommender system was reported by a single article. We also observed that many authors neglect the user-modeling process: 81% of the approaches made their users provide some keyword, text snippets, or a single input paper to represent their information need. Few approaches automatically inferred information needs from the users' authored, tagged, or otherwise connected papers.

One step towards using the full potential of research-paper recommender systems could be to establish a platform for researchers to collaborate, work on joint publications, communicate ideas, or to establish conferences or workshops focusing solely on research-paper recommender systems. An open-source recommender framework containing the most promising approaches could help transfer the research results into practice. Such a framework would also help new researchers in the field access a number of baselines they could compare their own approaches with. A framework could either be built from scratch, or be based on existing frameworks such as MyMediaLite[36], LensKit[37], Mahout[38], Duine[39], RecLab Core[40], easyrec[41], or Recommender101[42]. Finally, the community could benefit from considering research results from related disciplines. In particular, research in the area of user modeling and scientometrics appears promising, as well as research from the general recommender-systems community about aspects beyond accuracy.

---

[36] http://www.mymedialite.net/

[37] http://lenskit.grouplens.org/

[38] http://mahout.apache.org/

[39] http://www.duineframework.org/

[40] http://code.richrelevance.com/reclab-core/

[41] http://easyrec.org/

[42] http://ls13-www.cs.uni-dortmund.de/homepage/recommender101/index.shtml

References

[1]   Google Scholar, "Scholar Update: Making New Connections," *Google Scholar Blog, http://googlescholar.blogspot.de/2012/08/scholar-updates-making-new-connections.html*, 2012.

[2]   E. M. Airoldi, D. M. Blei, S. E. Fienberg, E. P. Xing, and T. Jaakkola, "Mixed membership stochastic block models for relational data with application to protein-protein interactions," in *Proceedings of the International Biometrics Society Annual Meeting*, 2006, pp. 1–34.

[3]   A. Arnold and W. W. Cohen, "Information extraction as link prediction: Using curated citation networks to improve gene detection," in *Proceedings of the 4th International Conference on Wireless Algorithms, Systems, and Applications*, 2009, pp. 541–550.

[4]   J. Beel, S. Langer, and M. Genzmehr, "Sponsored vs. Organic (Research Paper) Recommendations and the Impact of Labeling," in *Proceedings of the 17th International Conference on Theory and Practice of Digital Libraries (TPDL 2013)*, 2013, pp. 395–399.

[5]   J. Beel, S. Langer, M. Genzmehr, and A. Nürnberger, "Persistence in Recommender Systems: Giving the Same Recommendations to the Same Users Multiple Times," in *Proceedings of the 17th International Conference on Theory and Practice of Digital Libraries (TPDL 2013)*, 2013, vol. 8092, pp. 390–394.

[6]   J. Beel, S. Langer, M. Genzmehr, and A. Nürnberger, "Introducing Docear's Research Paper Recommender System," in *Proceedings of the 13th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL'13)*, 2013, pp. 459–460.

[7]   J. Beel, S. Langer, A. Nürnberger, and M. Genzmehr, "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)*, 2013, pp. 400–404.

[8]   W. Böhm, A. Geyer-schulz, M. Hahsler, and M. Jahn, "Repeat-Buying Theory and Its Application for Recommender Services," in *Proceedings of the 25th Annual Conference of the Gesellschaft für Klassifikation e.V.*, 2003, pp. 229–239.

[9]   M. Baez, D. Mirylenka, and C. Parra, "Understanding and supporting search for scholarly knowledge," in *Proceeding of the 7th European Computer Science Summit*, 2011, pp. 1–8.

[10]   J. Beel, B. Gipp, S. Langer, and M. Genzmehr, "Docear: An Academic Literature Suite for Searching, Organizing and Creating Academic Literature," in *Proceedings of the 11th Annual International ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, 2011, pp. 465–466.

[11]   J. Beel, B. Gipp, and C. Mueller, "SciPlore MindMapping' - A Tool for Creating Mind Maps Combined with PDF and Reference Management," *D-Lib Magazine*, vol. 15, no. 11, Nov. 2009.

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

[13]   T. Bogers and A. van den Bosch, "Recommending scientific articles using citeulike," in *Proceedings of the 2008 ACM conference on Recommender systems*, 2008, pp. 287–290.

[14] K. D. Bollacker, S. Lawrence, and C. L. Giles, "CiteSeer: An autonomous web agent for automatic retrieval and identification of interesting publications," in *Proceedings of the 2nd international conference on Autonomous agents*, 1998, pp. 116–123.

[15] J. Bollen and H. Van de Sompel, "An architecture for the aggregation and analysis of scholarly usage data," in *Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries*, 2006, pp. 298–307.

[16] T. CiteSeerX, "About RefSeer," *http://refseer.ist.psu.edu/about*, 2012.

[17] CiteULike, "My Top Recommendations," *Website, http://www.citeulike.org/profile/<username>/recommendations*. 2011.

[18] CiteULike, "Science papers that interest you," *Blog, http://blog.citeulike.org/?p=11*. Oct-2009.

[19] CiteULike, "Data from CiteULike's new article recommender," *Blog, http://blog.citeulike.org/?p=136*. Nov-2009.

[20] C. Caragea, A. Silvescu, P. Mitra, and C. L. Giles, "Can't See the Forest for the Trees? A Citation Recommendation System," in *iConference 2013 Proceedings*, 2013, pp. 849–851.

[21] K. Chandrasekaran, S. Gauch, P. Lakkaraju, and H. Luong, "Concept-based document recommendations for citeseer authors," in *Proceedings of the 5th international conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, 2008, pp. 83–92.

[22] W. Choochaiwattana, "Usage of tagging for research paper recommendation," in *Proceedings of the 3rd International Conference on Advanced Computer Theory and Engineering (ICACTE)*, 2010, vol. 2, pp. 439–442.

[23] I. Councill, C. Giles, E. Di Iorio, M. Gori, M. Maggini, and A. Pucci, "Towards next generation CiteSeer: A flexible architecture for digital library deployment," in *Research and Advanced Technology for Digital Libraries*, 2006, pp. 111–122.

[24] R. Dong, L. Tokarchuk, and A. Ma, "Digging Friendship: Paper Recommendation in Social Network," in *Proceedings of Networking & Electronic Commerce Research Conference (NAEC 2009)*, 2009, pp. 21–28.

[25] EXLibris, "bx Recommender Service," *http://www.exlibrisgroup.com/?catid=7098DEDA-0C18-41C0-B9E0-FD36B705EE5D*. 2013.

[26] M. D. Ekstrand, P. Kannan, J. A. Stemper, J. T. Butler, J. A. Konstan, and J. T. Riedl, "Automatically building research reading lists," in *Proceedings of the fourth ACM conference on Recommender systems*, 2010, pp. 159–166.

[27] E. Erosheva, S. Fienberg, and J. Lafferty, "Mixed-membership models of scientific publications," in *Proceedings of the National Academy of Sciences of the United States of America*, 2004, vol. 101, no. Suppl 1, pp. 5220–5227.

[28] M. Franke and A. Geyer-Schulz, "Using restricted random walks for library recommendations and knowledge space exploration," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 21, no. 02, pp. 355–373, 2007.

[29] F. Ferrara, N. Pudota, and C. Tasso, "A Keyphrase-Based Paper Recommender System," in *Proceedings of the IRCDL'11*, 2011, pp. 14–25.

[30] A. Geyer-Schulz and M. Hahsler, "Comparing two recommender algorithms with the help of recommendations by peers," in *Proceedings of the WEBKDD 2002 - Mining Web Data for Discovering Usage Patterns and Profiles*, 2003, pp. 137–158.

[31] A. Geyer-Schulz and M. Hahsler, "Evaluation of recommender algorithms for an internet information broker based on simple association rules and on the repeat-

buying theory," in *Proceedings of the 4th WebKDD Workshop: Web Mining for Usage Patterns & User Profiles*, 2002, pp. 100–114.

[32] A. Geyer-Schulz, M. Hahsler, and M. Jahn, "A customer purchase incidence model applied to recommender services," in *Proceedings of the Third International Workshop on Mining Web Log Data Across All Customers Touch Points*, 2002, pp. 25–47.

[33] A. Geyer-Schulz, M. Hahsler, and M. Jahn, "Recommendations for virtual universities from observed user behavior," in *Proceedings of the 24th Annual Conference of the Gesellschaft für Klassifikation e.V.*, 2002, pp. 273–280.

[34] A. Geyer-Schulz, M. Hahsler, M. Jahn, and A. Geyer, "Wissenschaftliche Recommendersysteme in Virtuellen Universitäten," in *Proceedings of the Symposiom of Unternehmen Hochschule*, 2001, pp. 101–114.

[35] A. Geyer-Schulz, M. Hahsler, A. Neumann, and A. Thede, "An integration strategy for distributed recommender services in legacy library systems," in *Between Data Science and Applied Data Analysis*, Springer, 2003, pp. 412–420.

[36] A. Geyer-Schulz, M. Hahsler, A. Neumann, and A. Thede, "Behavior-based recommender systems as value-added services for scientific libraries," *Statistical Data Mining & Knowledge Discovery*, pp. 433–454, 2003.

[37] A. Geyer-Schulz, M. Hahsler, and A. Thede, "Comparing Simple Association-Rules and Repeat-Buying Based Recommender Systems in a B2B Environment," in *Proceedings of the 26th Annual Conference of the Gesellschaft für Klassifikation e.V.*, 2003, pp. 421–429.

[38] A. Geyer-Schulz, A. Neumann, and A. Thede, "An architecture for behavior-based library recommender systems," *Information technology and libraries*, vol. 22, no. 4, pp. 165–174, 2003.

[39] A. Geyer-Schulz, A. Neumann, and A. Thede, "Others also use: A robust recommender system for scientific libraries," in *Proceedings of the 7th European Conference on Digital Libraries*, 2003, pp. 113–125.

[40] B. Gillitzer, "Der Empfehlungsdienst BibTip - Ein flächendeckendes Angebot im Bibliotheksverbund Bayern," *http://www.b-i-t-online.de/heft/2010-01/nachrichtenbeitrag3*. pp. 1–4, 2010.

[41] S. Gottwald, "Recommender Systeme fuer den Einsatz in Bibliotheken / Survey on recommender systems," Konrad-Zuse-Zentrum für Informationstechnik Berlin, ZIB-Report 11-30, 2011.

[42] A. Geyer-Schulz, M. Hahsler, and M. Jahn, "Educational and scientific recommender systems: Designing the information channels of the virtual university," *International Journal of Engineering Education*, vol. 17, no. 2, pp. 153–163, 2001.

[43] C. L. Giles, K. D. Bollacker, and S. Lawrence, "CiteSeer: An automatic citation indexing system," in *Proceedings of the 3rd ACM conference on Digital libraries*, 1998, pp. 89–98.

[44] B. Gipp and J. Beel, "Citation Proximity Analysis (CPA) - A new approach for identifying related work based on Co-Citation Analysis," in *Proceedings of the 12th International Conference on Scientometrics and Informetrics (ISSI'09)*, 2009, vol. 2, pp. 571–575.

[45] B. Gipp, J. Beel, and C. Hentschel, "Scienstein: A Research Paper Recommender System," in *Proceedings of the International Conference on Emerging Trends in Computing (ICETiC'09)*, 2009, pp. 309–315.

[46] M. Gori and A. Pucci, "Research paper recommender systems: A random-walk based approach," in *Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence*, 2006, pp. 778–781.

[47] V. Henning and J. Reichelt, "Mendeley-A Last. fm For Research?," in *Proceedings of the IEEE Fourth International Conference on eScience*, 2008, pp. 327–328.

[48] S.-Y. Hwang, W.-C. Hsiung, and W.-S. Yang, "A prototype WWW literature recommendation system for digital libraries," *Online Information Review*, vol. 27, no. 3, pp. 169–182, 2003.

[49] J. He, J.-Y. Nie, Y. Lu, and W. X. Zhao, "Position-Aligned translation model for citation recommendation," in *Proceedings of the 19th international conference on String Processing and Information Retrieval*, 2012, pp. 251–263.

[50] Q. He, D. Kifer, J. Pei, P. Mitra, and C. L. Giles, "Citation recommendation without author supervision," in *Proceedings of the fourth ACM international conference on Web search and data mining*, 2011, pp. 755–764.

[51] Q. He, J. Pei, D. Kifer, P. Mitra, and L. Giles, "Context-aware citation recommendation," in *Proceedings of the 19th international conference on World wide web*, 2010, pp. 421–430.

[52] C. Hess, "Trust-based recommendations in multi-layer networks," IOS Press, 2008.

[53] C. Hess, "Trust-based recommendations for publications: A multi-layer network approach," *TCDL Bulletin*, vol. 2, no. 2, pp. 190–201, 2006.

[54] C. Hess, K. Stein, and C. Schlieder, "Trust-enhanced visibility for personalized document recommendations," in *Proceedings of the 2006 ACM symposium on Applied computing*, 2006, pp. 1865–1869.

[55] S. Huang, G. R. Xue, B. Y. Zhang, Z. Chen, Y. Yu, and W. Y. Ma, "Tssp: A reinforcement algorithm to find related papers," in *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, 2004, pp. 117–123.

[56] W. Huang, S. Kataria, C. Caragea, P. Mitra, C. L. Giles, and L. Rokach, "Recommending citations: translating papers into references," in *Proceedings of the 21st ACM international conference on Information and knowledge management*, 2012, pp. 1910–1914.

[57] Z. Huang, W. Chung, T. H. Ong, and H. Chen, "A graph-based recommender system for digital library," in *Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries*, 2002, pp. 65–73.

[58] K. Jack, "Mendeley: Recommendation Systems for Academic Literature," *Presentation at Technical University of Graz (TUG)*. 2012.

[59] K. Jack, "Mendeley Suggest: Engineering a Personalised Article Recommender System," *Presentation at RecSysChallenge workshop 2012*. 2012.

[60] K. Jack, "Mahout Becomes a Researcher: Large Scale Recommendations at Mendeley," *Presentation at Big Data Week Conferences*. 2012.

[61] Y. Jiang, A. Jia, Y. Feng, and D. Zhao, "Recommending academic papers via users' reading purposes," in *Proceedings of the sixth ACM conference on Recommender systems*, 2012, pp. 241–244.

[62] P. Jomsri, S. Sanguansintukul, and W. Choochaiwattana, "A framework for tag-based research paper recommender system: an IR approach," in *Proceedings of the 24th International Conference on Advanced Information Networking and Applications (WAINA)*, 2010, pp. 103–108.

[63] N. Kapoor, J. Chen, J. T. Butler, G. C. Fouty, J. A. Stemper, J. Riedl, and J. A. Konstan, "Techlens: a researcher's desktop," in *Proceedings of the 2007 ACM conference on Recommender systems*, 2007, pp. 183–184.

[64] J. A. Konstan, N. Kapoor, S. M. McNee, and J. T. Butler, "Techlens: Exploring the use of recommenders to support users of digital libraries," in *Proceedings of the Coalition for Networked Information Fall 2005 Task Force Meeting*, 2005, pp. 111–112.

[65] S. Kataria, P. Mitra, and S. Bhatia, "Utilizing context in generative bayesian models for linked corpus," in *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 2010, pp. 1340–1345.

[66] A. Kodakateri Pudhiyaveetil, S. Gauch, H. Luong, and J. Eno, "Conceptual recommender system for CiteSeerX," in *Proceedings of the third ACM conference on Recommender systems*, 2009, pp. 241–244.

[67] M. Kuberek and M. Mönnich, "Einsatz von Recommendersystemen in Bibliotheken Recommender systems in libraries," *Presentation*. 2012.

[68] O. Küçüktunç, K. Kaya, E. Saule, and U. V. Catalyürek, "Fast Recommendation on Bibliographic Networks," in *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2012, pp. 480–487.

[69] O. Küçüktunç, K. Kaya, E. Saule, and U. V. Catalyürek, "Fast Recommendation on Bibliographic Networks with Sparse-Matrix Ordering and Partitioning," *Social Network Analysis and Mining*, vol. 3, no. 4, pp. 1097–1111, 2013.

[70] O. Küçüktunç, E. Saule, K. Kaya, and Ü. V. Çatalyürek, "Result Diversification in Automatic Citation Recommendation," in *Proceedings of the iConference Workshop on Computational Scientometrics: Theory and Applications*, 2013, pp. 1–4.

[71] O. Küçüktunç, E. Saule, K. Kaya, and Ü. V. Çatalyürek, "Diversifying Citation Recommendations," *arXiv preprint arXiv:1209.5809*, pp. 1–19, 2012.

[72] O. Küçüktunç, E. Saule, K. Kaya, and Ü. V. Çatalyürek, "Recommendation on Academic Networks using Direction Aware Citation Analysis," *arXiv preprint arXiv:1205.1143*, pp. 1–10, 2012.

[73] O. Küçüktunç, E. Saule, K. Kaya, and Ü. V. Çatalyürek, "Direction Awareness in Citation Recommendation," in *Proceedings of DBRank Workshop in conjunction with VLDB'12*, 2012, pp. 161–166.

[74] N. Lao, "Efficient Random Walk Inference with Knowledge Bases," PhD Thesis. The Carnegie Mellon University, 2012.

[75] N. Lao and W. W. Cohen, "Personalized Reading Recommendations for Saccharomyces Genome Database," *Unpublished Paper, http://www.cs.cmu.edu/ nlao/publication/2012/2012.dils.pdf*, pp. 1–15, 2012.

[76] N. Lao and W. W. Cohen, "Personalized Reading Recommendations for Saccharomyces Genome Database," *Unpublished Poster, http://www.cs.cmu.edu/ nlao/publication/2012/2012.dils.poster.portrat.pdf*, 2012.

[77] N. Lao and W. W. Cohen, "Contextual Recommendation with Path Constrained Random Walks," *Unpublished, http://www.cs.cmu.edu/ nlao/doc/2011.cikm.pdf*, pp. 1–9, 2011.

[78] P. Lakkaraju, S. Gauch, and M. Speretta, "Document similarity based on concept tree distance," in *Proceedings of the nineteenth ACM conference on Hypertext and hypermedia*, 2008, pp. 127–132.

[79] N. Lao and W. W. Cohen, "Relational retrieval using a combination of path-constrained random walks," *Machine learning*, vol. 81, no. 1, pp. 53–67, 2010.

[80] K. D. B. S. Lawrence, "A System For Automatic Personalized Tracking of Scientific Literature on the Web," in *Proceedings of the fourth ACM conference on Digital libraries*, 1999, pp. 105–113.

[81] S. R. Lawrence, K. D. Bollacker, and C. L. Giles, "Autonomous citation indexing and literature browsing using citation context," U.S. Patent US 6,738,780 B2Summer-2004.

[82] S. R. Lawrence, C. L. Giles, and K. D. Bollacker, "Autonomous citation indexing and literature browsing using citation context," U.S. Patent US 6,289,342 B1Nov-2001.

[83] H. Li, I. Councill, W.-C. Lee, and C. L. Giles, "CiteSeerx: an architecture and web service design for an academic document search engine," in *Proceedings of the 15th international conference on World Wide Web*, 2006, pp. 883–884.

[84] Y. Liang, Q. Li, and T. Qian, "Finding relevant papers based on citation relations," in *Proceedings of the 12th international conference on Web-age information management*, 2011, pp. 403–414.

[85] J. Lin and W. J. Wilbur, "PubMed Related Articles: a Probabilistic Topic-based Model for Content Similarity," *BMC Bioinformatics*, vol. 8, no. 1, pp. 423–436, 2007.

[86] Y. Lu, J. He, D. Shan, and H. Yan, "Recommending citations with translation model," in *Proceedings of the 20th ACM international conference on Information and knowledge management*, 2011, pp. 2017–2020.

[87] S. M. McNee, N. Kapoor, and J. A. Konstan, "Don't look stupid: avoiding pitfalls when recommending research papers," in *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, 2006, pp. 171–180.

[88] S. E. Middleton, H. Alani, and D. C. De Roure, "Exploiting synergy between ontologies and recommender systems," in *Proceedings of the Semantic Web Workshop*, 2002, pp. 1–10.

[89] S. E. Middleton, D. De Roure, and N. R. Shadbolt, "Ontology-based recommender systems," in *Handbook on Ontologies*, Springer, 2009, pp. 779–796.

[90] S. E. Middleton, D. C. De Roure, and N. R. Shadbolt, "Foxtrot recommender system: User profiling, ontologies and the World Wide Web," in *Proceedings of the WWW Conference*, 2002, pp. 1–3.

[91] S. E. Middleton, D. C. De Roure, and N. R. Shadbolt, "Capturing knowledge of user preferences: ontologies in recommender systems," in *Proceedings of the 1st international conference on Knowledge capture*, 2001, pp. 100–107.

[92] M. Mönnich and M. Spiering, "Adding value to the library catalog by implementing a recommendation system," *D-Lib Magazine*, vol. 14, no. 5, pp. 4–11, 2008.

[93] S. M. McNee, I. Albert, D. Cosley, P. Gopalkrishnan, S. K. Lam, A. M. Rashid, J. A. Konstan, and J. Riedl, "On the Recommending of Citations for Research Papers," in *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 2002, pp. 116–125.

[94] S. E. Middleton, N. R. Shadbolt, and D. C. De Roure, "Ontological user profiling in recommender systems," *ACM Transactions on Information Systems (TOIS)*, vol. 22, no. 1, pp. 54–88, 2004.

[95] M. Monnich and M. Spiering, "Einsatz von BibTip als Recommendersystem im Bibliothekskatalog," *Bibliotheksdienst*, vol. 42, no. 1, pp. 54–54, 2008.

[96] A. Naak, "Papyres: un système de gestion et de recommandation d'articles de recherche," Master Thesis. Université de Montréal, 2009.

[97] A. W. Neumann, "Recommender Systems for Information Providers," Springer, 2009, pp. 91–119.

[98] A. Naak, H. Hage, and E. Aimeur, "A multi-criteria collaborative filtering approach for research paper recommendation in papyres," in *Proceedings of the 4th International Conference MCETECH*, 2009, pp. 25–39.

[99] A. Naak, H. Hage, and E. Aimeur, "Papyres: A Research Paper Management System," in *Proceedings of the 10th E-Commerce Technology Conference on Enterprise Computing, E-Commerce and E-Services*, 2008, pp. 201–208.

[100] R. M. Nallapati, A. Ahmed, E. P. Xing, and W. W. Cohen, "Joint latent topic models for text and citations," in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2008, pp. 542–550.

[101] C. Nascimento, A. H. Laender, A. S. da Silva, and M. A. Gonçalves, "A source independent framework for research paper recommendation," in *Proceedings of the 11th annual international ACM/IEEE joint conference on Digital libraries*, 2011, pp. 297–306.

[102] T. Ozono, S. Goto, N. Fujimaki, and T. Shintani, "P2p based knowledge source discovery on research support system papits," in *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, 2002, pp. 49–50.

[103] T. Ozono and T. Shintani, "P2P based Information Retrieval on Research Support System Papits," in *Proceedngs of the IASTED International Conference on Artificial and Computational Intelligence*, 2002, pp. 136–141.

[104] T. Ozono and T. Shintani, "Paper classification for recommendation on research support system papits," *IJCSNS International Journal of Computer Science and Network Security*, vol. 6, pp. 17–23, 2006.

[105] T. Ozono, T. Shintani, T. Ito, and T. Hasegawa, "A feature selection for text categorization on research support system Papits," in *Proceedings of the 8th Pacific Rim International Conference on Artificial Intelligence*, 2004, pp. 524–533.

[106] D. M. Pennock, E. Horvitz, S. Lawrence, and C. L. Giles, "Collaborative filtering by personality diagnosis: A hybrid memory-and model-based approach," in *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, 2000, pp. 473–480.

[107] Y. Petinot, C. L. Giles, V. Bhatnagar, P. B. Teregowda, and H. Han, "Enabling interoperability for autonomous digital libraries: an API to citeseer services," in *Digital Libraries, 2004. Proceedings of the 2004 Joint ACM/IEEE Conference on*, 2004, pp. 372–373.

[108] Y. Petinot, C. L. Giles, V. Bhatnagar, P. B. Teregowda, H. Han, and I. Councill, "A service-oriented architecture for digital libraries," in *Proceedings of the 2nd international conference on Service oriented computing*, 2004, pp. 263–268.

[109] S. Pohl, "Using Access Data for Paper Recommendations on ArXiv. org," Master Thesis. Technical University of Darmstadt, 2007.

[110] S. Pohl, F. Radlinski, and T. Joachims, "Recommending related papers based on digital library access records," in *Proceedings of the 7th ACM/IEEE-CS joint conference on Digital libraries*, 2007, pp. 417–418.

[111] T. Researchgate, "Researchgate Recommender," *http://www.researchgate.net/directory/publications/*, 2011.

[112] L. Rokach, P. Mitra, S. Kataria, W. Huang, and L. Giles, "A Supervised Learning Method for Context-Aware Citation Recommendation in a Large Corpus," in

Proceedings of the Large-Scale and Distributed Systems for Information Retrieval Workshop (LSDS-IR), 2013, pp. 17–22.

[113] Sarkanto, "About the Sarkanto Recommender Demo," *http://lab.cisti-icist.nrc-cnrc.gc.ca/Sarkanto/about.jsp*. 2013.

[114] T. Strohman, W. B. Croft, and D. Jensen, "Recommending citations for academic papers," in *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 2007, pp. 705–706.

[115] K. Sugiyama and M.-Y. Kan, "Scholarly paper recommendation via user's recent research interests," in *Proceedings of the 10th ACM/IEEE Annual Joint Conference on Digital Libraries (JCDL)*, 2010, pp. 29–38.

[116] D. Thomas, A. Greenberg, and P. Calarco, "Scholarly Usage Based Recommendations: Evaluating bX for a Consortium," *Presentation, http://igelu.org/wp-content/uploads/2011/09/bx_igelu_presentation_updated_september-13.pdf*. 2011.

[117] R. Torres, S. M. McNee, M. Abel, J. A. Konstan, and J. Riedl, "Enhancing digital libraries with TechLens+," in *Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries*, 2004, pp. 228–236.

[118] K. Uchiyama, H. Nanba, A. Aizawa, and T. Sagara, "OSUSUME: cross-lingual recommender system for research papers," in *Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation*, 2011, pp. 39–42.

[119] A. Vellino, "A comparison between usage-based and citation-based methods for recommending scholarly research articles," in *Proceedings of the American Society for Information Science and Technology*, 2010, vol. 47, no. 1, pp. 1–2.

[120] A. Vellino and D. Zeber, "A hybrid, multi-dimensional recommender for journal articles in a scientific digital library," in *Proceedings of the 2007 IEEE/WIC/ACM International Conference on Web Intelligence*, 2007, pp. 111–114.

[121] Y. Wang, E. Zhai, J. Hu, and Z. Chen, "Claper: Recommend classical papers to beginners," in *Seventh International Conference on Fuzzy Systems and Knowledge Discovery*, 2010, vol. 6, pp. 2777–2781.

[122] S. Watanabe, T. Ito, T. Ozono, and T. Shintani, "A paper recommendation mechanism for the research support system papits," in *Proceedings of the International Workshop on Data Engineering Issues in E-Commerce*, 2005, pp. 71–80.

[123] A. Woodruff, R. Gossweiler, J. Pitkow, E. H. Chi, and S. K. Card, "Enhancing a digital book with a reading recommender," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2000, pp. 153–160.

[124] C. Yang, B. Wei, J. Wu, Y. Zhang, and L. Zhang, "CARES: a ranking-oriented CADAL recommender system," in *Proceedings of the 9th ACM/IEEE-CS joint conference on Digital libraries*, 2009, pp. 203–212.

[125] F. Zarrinkalam and M. Kahani, "SemCiR - A citation recommendation system based on a novel semantic distance measure," *Program: electronic library and information systems*, vol. 47, no. 1, pp. 92–112, 2013.

[126] F. Zarrinkalam and M. Kahani, "A New Metric for Measuring Relatedness of Scientific Papers Based on Non-Textual Features," *Intelligent Information Management*, vol. 4, no. 4, pp. 99–107, 2012.

[127] D. Zhou, S. Zhu, K. Yu, X. Song, B. L. Tseng, H. Zha, and C. L. Giles, "Learning multiple graphs for document recommendations," in *Proceedings of the 17th international conference on World Wide Web*, 2008, pp. 141–150.

[128] H. Avancini, L. Candela, and U. Straccia, "Recommenders in a personalized, collaborative digital library environment," *Journal of Intelligent Information Systems*, vol. 28, no. 3, pp. 253–283, 2007.

[129] N. Agarwal, E. Haque, H. Liu, and L. Parsons, "A subspace clustering framework for research group collaboration," *International Journal of Information Technology and Web Engineering*, vol. 1, no. 1, pp. 35–58, 2006.

[130] N. Agarwal, E. Haque, H. Liu, and L. Parsons, "Research Paper Recommender Systems: A Subspace Clustering Approach," in *Proceedings of the 6th international conference on Advances in Web-Age Information Management (WAIM'05)*, 2005, pp. 475–491.

[131] J. Bollen and L. M. Rocha, "An adaptive systems approach to the implementation and evaluation of digital library recommendation systems," in *Proceedings of the 4th European Conference on Digital Libraries*, Springer, 2000, pp. 356–359.

[132] C. Bancu, M. Dagadita, M. Dascalu, C. Dobre, S. Trausan-Matu, and A. M. Florea, "ARSYS-Article Recommender System," in *Proceedings of the 14th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*, 2012, pp. 349–355.

[133] S. C. Cazella and L. O. C. Alvares, "Combining Data Mining Technique and Users' Relevance Opinion to Build an Efficient Recommender System," *Revista Tecnologia da Informação, UCB*, vol. 4, no. 2, 2005.

[134] S. Cazella and L. Alvares, "Modeling user's opinion relevance to recommending research papers," in *Proceedings of the UMAP Conference*, 2005, pp. 150–150.

[135] P. Chirawatkul, "Structured Peer-to-Peer Search to build a Bibliographic Paper Recommendation System," Saarland University, 2006.

[136] A. Dattolo, F. Ferrara, and C. Tasso, "Supporting personalized user concept spaces and recommendations for a publication sharing system," in *Proceedings of the 17th International Conference on User Modeling, Adaptation, and Personalization*, 2009, pp. 325–330.

[137] A. Daud and A. H. Muhammad Akramand Rajpar Shaikh, "Scientific Reference Mining using Semantic Information through Topic Modeling," *Research Journal of Engineering & Technology*, vol. 28, no. 2, pp. 253–262, 2009.

[138] U. Farooq, C. H. Ganoe, J. M. Carroll, I. G. Councill, and C. Lee Giles, "Design and evaluation of awareness mechanisms in CiteSeer," *Information Processing & Management*, vol. 44, no. 2, pp. 596–612, 2008.

[139] L. Fernández, J. A. Sánchez, and A. Garcıa, "Mibiblio: personal spaces in a digital library universe," in *Proceedings of the fifth ACM conference on Digital libraries*, 2000, pp. 232–233.

[140] T. Gross, "CYCLADES: A distributed system for virtual community support based on open archives," in *Proceedings of the 11th Euromicro Conference on Parallel, Distributed and Network-Based Processing*, 2003, pp. 484–491.

[141] G. Geisler, D. McArthur, and S. Giersch, "Developing recommendation services for a digital library with uncertain and changing data," in *Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries*, 2001, pp. 199–200.

[142] K. Hong, H. Jeon, and C. Jeon, "UserProfile-based personalized research paper recommendation system," in *Proceedings of the 8th International Conference on Computing and Networking Technology*, 2012, pp. 134–138.

[143] Y. Huang, "Combining Social Networks and Content for Recommendation in a Literature Digital Library," National Sun Yat-Sen University, Taiwan, 2007.

[144] S. Kang and Y. Cho, "A novel personalized paper search system," in *Proceedings of the International Conference on Intelligent Computing*, 2006, pp. 1257–1262.

[145] G. H. Martin, S. Schockaert, C. Cornelis, and H. Naessens, "Metadata impact on research paper similarity," in *14th European Conference on Digital Libraries*, 2010, pp. 457–460.

[146] J. M. Morales-del-Castillo, E. Peis, and E. Herrera-Viedma, "A filtering and recommender system prototype for scholarly users of digital libraries," in *Proceedings of the Second World Summit on the Knowledge Society*, Springer, 2009, pp. 108–117.

[147] Y. Mao, J. Vassileva, and W. Grassmann, "A system dynamics approach to study virtual communities," in *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, 2007, pp. 178–197.

[148] N. F. Matsatsinis, K. Lakiotaki, and P. Delia, "A system based on multiple criteria analysis for scientific paper recommendation," in *Proceedings of the 11th Panhellenic Conference on Informatics*, 2007, pp. 135–149.

[149] G. Mishra, "Optimised Research Paper Recommender System Using Social Tagging," *International Journal of Engineering Research and Applications*, vol. 2, no. 2, pp. 1503–1507, 2012.

[150] A. Nakagawa and T. Ito, "An implementation of a knowledge recommendation system based on similarity among users' profiles," in *Proceedings of the 41st SICE Annual Conference*, 2002, vol. 1, pp. 326–327.

[151] C. Pan and W. Li, "Research paper recommendation with topic analysis," in *Proceedings of the International Conference on Computer Design and Applications (ICCDA)*, 2010, pp. 264–268.

[152] H.-E. Popa, V. Negru, D. Pop, and I. Muscalagiu, "DL-AgentRecom-A multi-agent based recommendation system for scientific documents," in *Proceedings of the 10th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*, 2008, pp. 320–324.

[153] N. Ratprasartporn and G. Ozsoyoglu, "Finding related papers in literature digital libraries," in *Proceedings of the 11th European Conference on Digital Libraries*, 2007, pp. 271–284.

[154] L. M. Rocha, "TalkMine: a soft computing approach to adaptive knowledge recommendation," *Studies in fuzziness and soft computing*, vol. 75, pp. 89–116, 2001.

[155] L. M. Rocha, "Talkmine and the adaptive recommendation project," in *Proceedings of the fourth ACM conference on Digital libraries*, 1999, pp. 242–243.

[156] K. Stock, A. Robertson, F. Reitsma, T. Stojanovic, M. Bishr, D. Medyckyj-Scott, and J. Ortmann, "eScience for Sea Science: A Semantic Scientific Knowledge Infrastructure for Marine Scientists," in *Proceedings of the 5th IEEE International Conference on e-Science*, 2009, pp. 110–117.

[157] U. Straccia, "Cyclades: An Open Collaborative Virtual Archive Environment," *Poster (http://www.ercim.eu/cyclades/cyclades-fs.pdf)*, 2003.

[158] Z. Shaoping, "ActiveCite: An interactive system for automatic citation suggestion," Master Thesis. National University of Singapore, 2010.

[159] K. Stock, V. Karasova, A. Robertson, G. Roger, M. Small, M. Bishr, J. Ortmann, T. Stojanovic, F. Reitsma, L. Korczynski, B. Brodaric, and Z. Gardner, "Finding Science with Science: Evaluating a Domain and Scientific Ontology User Interface for the Discovery of Scientific Resources," *Transactions in GIS*, vol. 1, pp. 1–28, 2013.

[160] T. Y. Tang and G. McCalla, "Towards Pedagogy-Oriented Paper Recommendations and Adaptive Annotations for a Web-Based Learning System,"

in *Knowledge Representation and Automated Reasoning for E-Learning Systems*, 2003, pp. 72–80.

[161] J. Tang and J. Zhang, "A discriminative approach to Topic-Based citation recommendation," *Advances in Knowledge Discovery and Data Mining*, pp. 572–579, 2009.

[162] T. Tang and G. McCalla, "Utilizing artificial learners to help overcome the cold-start problem in a pedagogically-oriented paper recommendation system," in *Adaptive Hypermedia and Adaptive Web-Based Systems*, 2004, pp. 245–254.

[163] T. Tang and G. McCalla, "Beyond learners' interest: personalized paper recommendation based on their pedagogical features for an e-learning system," in *Proceedings of the 8th Pacific Rim International Conference on Artificial Intelligence*, Springer, 2004, pp. 301–310.

[164] T. Y. Tang and G. McCalla, "Mining implicit ratings for focused collaborative filtering for paper recommendations," in *Proceedings of the Workshop on User and Group Models for Web-based Adaptive Collaborative Environments*, 2003.

[165] T. Y. Tang and G. McCalla, "Smart recommendation for an evolving e-learning system," in *Proceedings of the Workshop on Technologies for Electronic Documents for Supporting Learning, at the International Conference on Artificial Intelligence in Education*, 2003, pp. 699–710.

[166] T. Y. Tang, "The design and study of pedagogical paper recommendation," PhD Thesis. University of Saskatchewan, 2008.

[167] T. Y. Tang and G. McCalla, "A multidimensional paper recommender: Experiments and evaluations," *Internet Computing, IEEE*, vol. 13, no. 4, pp. 34–41, 2009.

[168] T. Y. Tang and G. McCalla, "The pedagogical value of papers: a collaborative-filtering based paper recommender," *Journal of Digital Information*, vol. 10, no. 2, pp. 1–12, 2009.

[169] T. Y. Tang and G. McCalla, "On the pedagogically guided paper recommendation for an evolving web-based learning system," in *Proceedings of the FLAIRS Conference*, 2004, pp. 86–91.

[170] T. Y. Tang and G. McCalla, "The social affordance of a paper," in *Proceedings of the Workshop of Assessment of Group and Individual Learning Through Intelligent Visualization on the 13th International Conference on Artificial Intelligence in Education*, 2007, pp. 34–42.

[171] X. Tang and Q. Zeng, "Keyword clustering for user interest profiling refinement within paper recommender systems," *Journal of Systems and Software*, vol. 85, no. 1, pp. 87–101, 2012.

[172] J. Vassileva, "Harnessing p2p power in the classroom," in *Proceedings of the Conference on Intelligent tutoring systems*, 2004, pp. 305–314.

[173] J. Vassileva, "Supporting peer-to-peer user communities," in *Proceedings of the Conference On the Move to Meaningful Internet Systems*, 2002, pp. 230–247.

[174] J. Vassileva, R. Detters, J. Geer, G. Maccalla, S. Bull, and L. Kettel, "Lessons from deploying I-Help," in *Workshop on Multi-Agent Architectures for Distributed Learning Environments. In Proceedings of International Conference on AI and Education. San Antonio, TX*, 2001, pp. 3–11.

[175] A. S. Vivacqua, J. Oliveira, and J. M. de Souza, "i-ProSE: inferring user profiles in a scientific context," *The Computer Journal*, vol. 52, no. 7, pp. 789–798, 2009.

[176] S.-S. Weng and H.-L. Chang, "Using ontology network analysis for research document recommendation," *Expert Systems with Applications*, vol. 34, no. 3, pp. 1857–1869, 2008.

[177] P. Winoto, T. Y. Tang, and G. I. McCalla, "Contexts in a paper recommendation system with collaborative filtering," *The International Review of Research in Open and Distance Learning*, vol. 13, no. 5, pp. 56–75, 2012.

[178] H. Wu, Y. Hua, B. Li, and Y. Pei, "Enhancing citation recommendation with various evidences," in *Proceedings of the 9th International Conference on Fuzzy Systems and Knowledge Discovery (FSKD)*, 2012, pp. 1160–1165.

[179] H. Xia, J. Li, J. Tang, and M.-F. Moens, "Plink-LDA: using link as prior information in topic modeling," in *Proceedings of the Conference on Database Systems for Advanced Applications (DASFAA)*, 2012, pp. 213–227.

[180] Q. Yang, S. Zhang, and B. Feng, "Research on Personalized Recommendation System of Scientific and Technological Periodical Based on Automatic Summarization," in *Proceedings of the 1st International Symposium on Information Technologies and Applications in Education*, 2007, pp. 34–39.

[181] S.-Y. Yang and C.-L. Hsu, "A New Ontology-Supported and Hybrid Recommending Information System for Scholars," in *Proceedings of the 13th International Conference on Network-Based Information Systems (NBiS)*, 2010, pp. 379–384.

[182] P. Yin, M. Zhang, and X. Li, "Recommending scientific literatures in a collaborative tagging environment," in *Proceedings of the 10th international conference on Asian digital libraries*, Springer, 2007, pp. 478–481.

[183] F. Zarrinkalam and M. Kahani, "A multi-criteria hybrid citation recommendation system based on linked data," in *Proceedings of the 2nd International eConference on Computer and Knowledge Engineering*, 2012, pp. 283–288.

[184] M. Zhang, W. Wang, and X. Li, "A Paper Recommender for Scientific Literatures Based on Semantic Concept Similarity," in *Proceedings of the International Conference on Asian Digital Libraries*, 2008, pp. 359–362.

[185] Z. Zhang and L. Li, "A research paper recommender system based on spreading activation model," in *Proceedings of the 2nd International Conference on Information Science and Engineering (ICISE)*, 2010, pp. 928–931.

[186] S. Gottwald and T. Koch, "Recommender Systems for Libraries," in *Proceedings of the ACM International Conference on Recommender Systems*, 2011, pp. 1–5.

[187] S. Leong, "A survey of recommender systems for scientific papers," *Presentation, http://www.liquidpub.org/mediawiki/upload/f/ff/RecommenderSystems.pdf*. 2012.

[188] A. F. Smeaton and J. Callan, "Personalisation and recommender systems in digital libraries," *International Journal on Digital Libraries*, vol. 5, no. 4, pp. 299–308, 2005.

[189] H. Avancini, L. Candela, and U. Straccia, "Recommenders in a personalized, collaborative digital library environment," *Journal of Intelligent Information Systems*, vol. 28, no. 3, pp. 253–283, 2007.

[190] S. Alotaibi and J. Vassileva, "Trust-Based Recommendations for Scientific Papers Based on the Researcher's Current Interest," in *Artificial Intelligence in Education*, 2013, pp. 717–720.

[191] J. Beel, S. Langer, M. Genzmehr, B. Gipp, C. Breitinger, and A. Nürnberger, "Research Paper Recommender System Evaluation: A Quantitative Literature Survey," in *Proceedings of the Workshop on Reproducibility and Replication in Recommender Systems Evaluation (RepSys) at the ACM Recommender System Conference (RecSys)*, 2013, pp. 15–22.

[192] J. Beel, S. Langer, M. Genzmehr, B. Gipp, and A. Nürnberger, "A Comparative Analysis of Offline and Online Evaluations and Discussion of Research Paper Recommender System Evaluation," in *Proceedings of the Workshop on*

Reproducibility and Replication in Recommender Systems Evaluation (RepSys) at the ACM Recommender System Conference (RecSys), 2013, pp. 7–14.

[193] C. Chen, C. Mao, Y. Tang, G. Chen, and J. Zheng, "Personalized recommendation based on implicit social network of researchers," in *Joint International Conference, ICPCA/SWS*, 2013, pp. 97–107.

[194] D. De Nart, F. Ferrara, and C. Tasso, "Personalized Access to Scientific Publications: from Recommendation to Explanation," in *Proceedings of the International Conference on User Modeling, Adaptation, and Personalization*, 2013, pp. 296–301.

[195] D. De Nart, F. Ferrara, and C. Tasso, "RES: A Personalized Filtering Tool for CiteSeerX Queries Based on Keyphrase Extraction," in *Proceedings of the International Conference on User Modeling, Adaptation, and Personalization (UMAP)*, 2013, pp. 341–343.

[196] M. Franke, A. Geyer-Schulz, and A. Neumann, "Building recommendations from random walks on library opac usage data," in *Data Analysis, Classification and the Forward Search*, Springer, 2006, pp. 235–246.

[197] S. Kim, "iScholar: A mobile research support system," PhD Thesis. University of Regina, 2013.

[198] O. Küçüktunç, "Result Diversication on Spatial, Multidimensional, Opinion, and Bibliographic Data," Ohio State University, 2013.

[199] O. Küçüktunç, E. Saule, K. Kaya, and Ü. V. Çatalyürek, "TheAdvisor: a webservice for academic recommendation," in *Proceedings of the 13th ACM/IEEE-CS joint conference on Digital libraries*, 2013, pp. 433–434.

[200] O. Küçüktunç, E. Saule, K. Kaya, and Ü. V. Çatalyürek, "Towards a personalized, scalable, and exploratory academic recommendation service," in *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2013, pp. 636–641.

[201] Y. Lai and J. Zeng, "A cross-language personalized recommendation model in digital libraries," *The Electronic Library*, vol. 31, no. 3, pp. 164–277, 2013.

[202] Y. Li, M. Yang, and Z. M. Zhang, "Scientific articles recommendation," in *Proceedings of the 22nd ACM International conference on information & knowledge management*, 2013, pp. 1147–1156.

[203] J. Lee, K. Lee, and J. G. Kim, "Personalized Academic Research Paper Recommendation System," *ArXiv Preprint*, vol. http://arxiv.org/abs/1304.5457, pp. 1–8, 2013.

[204] N. Manouselis and K. Verbert, "Layered Evaluation of Multi-Criteria Collaborative Filtering for Scientific Paper Recommendation," in *Procedia Computer Science*, 2013, vol. 18, pp. 1189–1197.

[205] F. Meng, D. Gao, W. Li, X. Sun, and Y. Hou, "A unified graph model for personalized query-oriented reference paper recommendation," in *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, 2013, pp. 1509–1512.

[206] B. P. Nunes, B. Fetahu, and M. A. Casanova, "Cite4Me: Semantic Retrieval and Analysis of Scientific Publications," in *Proceedings of the LAK Data Challenge, co-located at the International Learning Analytics & Knowledge Conference*, 2013, pp. 10–15.

[207] M. S. Pera and Y.-K. Ng, "Exploiting the wisdom of social connections to make personalized recommendations on scholarly articles," *Journal of Intelligent Information Systems*, pp. 1–21, 2013.

[208] K. Sugiyama and M.-Y. Kan, "Exploiting potential citation papers in scholarly paper recommendation," in *Proceedings of the 13th ACM/IEEE-CS joint conference on Digital libraries*, 2013, pp. 153–162.

[209] J. Sun, J. Ma, X. Liu, Z. Liu, G. Wang, H. Jiang, and T. Silva, "A Novel Approach for Personalized Article Recommendation in Online Scientific Communities," in *Proceedings of the 46th Hawaii International Conference on System Sciences (HICSS)*, 2013.

[210] J. Sun, J. Ma, Z. Liu, and Y. Miao, "Leveraging Content and Connections for Scientific Article Recommendation," *The Computer Journal*, pp. 60–71, 2013.

[211] G. Tian and L. Jing, "Recommending scientific articles using bi-relational graph-based iterative RWR," in *Proceedings of the 7th ACM conference on Recommender systems*, 2013, pp. 399–402.

[212] A. Vellino, "Usage-based vs. Citation-based Methods for Recommending Scholarly Research Articles," *Arxiv*, vol. http://arxiv.org/abs/1303.7149, 2013.

[213] R. Yan, H. Yan, and others, "Guess What You Will Cite: Personalized Citation Recommendation Based on Users's Preference," in *Proceedings of the Annual I&R Training and Education Conference*, 2013, pp. 428–439.

[214] W.-S. Yang and Y.-R. Lin, "A task-focused literature recommender system for digital libraries," *Online Information Review*, vol. 37, no. 4, pp. 581–601, 2013.

[215] W. Yao, J. He, G. Huang, J. Cao, and Y. Zhang, "Personalized Recommendation on Multi-Layer Context Graph," in *Web Information Systems Engineering (WISE 2013)*, 2013, pp. 135–148.

[216] L. Yu, J. Yang, D. Yang, and X. Yang, "A Decision Support System for Finding Research Topic based on Paper Recommendation," in *Proceedings of the Pacific Asia Conference on Information Systems*, 2013.

[217] F. Zarrinkalam and M. Kahani, "Using Semantic Relations to Improve Quality of a Citation Recommendation System," *Soft Computing Journal*, vol. 1, no. 2, pp. 36–45, 2013.

[218] Z. P. Zhang, L. N. Li, and H. Y. Yu, "A Hybrid Document Recommender Algorithm Based on Random Walk," *Applied Mechanics and Materials*, vol. 2270, pp. 336–338, 2013.

[219] J. Beel and B. Gipp, "Academic search engine spam and Google Scholar's resilience against it," *Journal of Electronic Publishing*, vol. 13, no. 3, 2010.

[220] J. Bar-Ilan, "Which h-index? - A comparison of WoS, Scopus and Google Scholar," *Scientometrics*, vol. 74, no. 2, pp. 257–271, 2007.

[221] A. Noruzi, "Google Scholar: The new generation of citation indexes," *Libri*, vol. 55, no. 4, pp. 170–180, 2005.

[222] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl, "GroupLens: an open architecture for collaborative filtering of netnews," in *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, 1994, pp. 175–186.

[223] P. Melville, R. J. Mooney, and R. Nagarajan, "Content-boosted collaborative filtering for improved recommendations," in *Proceedings of the National Conference on Artificial Intelligence*, 2002, pp. 187–192.

[224] Y. Shi, M. Larson, and A. Hanjalic, "Collaborative Filtering Beyond the User-Item Matrix: A Survey of the State of the Art and Future Challenges," *ACM Comput. Surv.*, vol. 47, no. 1, pp. 3:1–3:45, 2014.

[225] A. Gunawardana and G. Shani, "A survey of accuracy evaluation metrics of recommendation tasks," *The Journal of Machine Learning Research*, vol. 10, pp. 2935–2962, 2009.

[226] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl, "Evaluating collaborative filtering recommender systems," *ACM Transactions on Information Systems (TOIS)*, vol. 22, no. 1, pp. 5–53, 2004.

[227] M. Ge, C. Delgado-Battenfeld, and D. Jannach, "Beyond accuracy: evaluating recommender systems by coverage and serendipity," in *Proceedings of the fourth ACM conference on Recommender systems*, 2010, pp. 257–260.

[228] S. Bradshaw, "Reference directed indexing: Redeeming relevance for subject search in citation indexes," *Research and Advanced Technology for Digital Libraries*, pp. 499–510, 2003.

[229] A. Ritchie, S. Teufel, and S. Robertson, "Using terms from citations for IR: some first results," in *Advances in Information Retrieval*, Springer, 2008, pp. 211–221.

[230] A. Ritchie, "Citation context analysis for information retrieval," PhD Thesis. University of Cambridge, 2008.

[231] S. T. Dumais and J. Nielsen, "Automating the assignment of submitted manuscripts to reviewers," in *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*, 1992, pp. 233–244.

[232] F. Wang, N. Shi, and B. Chen, "A comprehensive survey of the reviewer assignment problem," *International Journal of Information Technology & Decision Making*, vol. 9, no. 04, pp. 645–668, 2010.

[233] J. E. Hirsch, "An index to quantify an individual's scientific research output," *Proceedings of the National academy of Sciences of the United States of America*, vol. 102, no. 46, p. 16569, 2005.

[234] H. Small, "Co-citation in the Scientific Literature: A New Measure of the Relationship Between Two Documents," *Journal of the American Society for Information Science*, vol. 24, pp. 265–269, 1973.

[235] M. M. Kessler, "Bibliographic Coupling Between Scientific Papers," *American Documentation*, vol. 14, pp. 10–25, 1963.

[236] K. Zyczkowski, "Citation graph, weighted impact factors and performance indices," *Scientometrics*, vol. 85, no. 1, pp. 301–315, 2010.

[237] G. Fischer, "User modeling in human-computer interaction," *User modeling and user-adapted interaction*, vol. 11, no. 1, pp. 65–86, 2001.

[238] M. Eirinaki and M. Vazirgiannis, "Web mining for web personalization," *ACM Transactions on Internet Technology (TOIT)*, vol. 3, no. 1, pp. 1–27, 2003.

[239] D. Pierrakos, G. Paliouras, C. Papatheodorou, and C. D. Spyropoulos, "Web usage mining as a tool for personalization: A survey," *User modeling and user-adapted interaction*, vol. 13, no. 4, pp. 311–372, 2003.

[240] R. J. Mooney and L. Roy, "Content-based book recommending using learning for text categorization," in *Proceedings of the fifth ACM conference on Digital libraries*, 2000, pp. 195–204.

[241] P. Brusilovsky, R. Farzan, and J. Ahn, "Comprehensive personalized information access in an educational digital library," in *Digital Libraries, 2005. JCDL'05. Proceedings of the 5th ACM/IEEE-CS Joint Conference on*, 2005, pp. 9–18.

[242] D. Faensen, L. Faultstich, H. Schweppe, A. Hinze, and A. Steidinger, "Hermes: a notification service for digital libraries," in *Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries*, 2001, pp. 373–380.

[243] S. Das, P. Mitra, and C. L. Giles, "Similar Researcher Search' in Academic Environments," in *Proceedings of the JCDL'12*, 2012, pp. 167–170.

[244] A. Abu-Jbara and D. Radev, "Coherent citation-based summarization of scientific papers," in *Proceedings of the 49th Annual Meeting of the Association for*

Computational Linguistics: Human Language Technologies, 2011, vol. 1, pp. 500–509.

[245] S. Mohammad, B. Dorr, M. Egan, A. Hassan, P. Muthukrishan, V. Qazvinian, D. Radev, and D. Zajic, "Using citations to generate surveys of scientific paradigms," in *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2009, pp. 584–592.

[246] S. Teufel and M. Moens, "Summarizing scientific articles: experiments with relevance and rhetorical status," *Computational linguistics*, vol. 28, no. 4, pp. 409–445, 2002.

[247] L. M. Collins, K. K. Mane, M. L. Martinez, J. A. Hussell, and R. E. Luce, "ScienceSifter: Facilitating activity awareness in collaborative research groups through focused information feeds," in *First International Conference on e-Science and Grid Computing*, 2005, pp. 40–47.

[248] R. Patton, T. Potok, and B. Worley, "Discovery & Refinement of Scientific Information via a Recommender System," in *INFOCOMP 2012, The Second International Conference on Advanced Communications and Computation*, 2012, pp. 31–35.

[249] R. Klamma, P. M. Cuong, and Y. Cao, "You never walk alone: Recommending academic events based on social network analysis," in *Complex Sciences*, Springer, 2009, pp. 657–670.

[250] Z. Yang and B. D. Davison, "Venue Recommendation: Submitting your Paper with Style," in *Machine Learning and Applications (ICMLA), 2012 11th International Conference on*, 2012, vol. 1, pp. 681–686.

[251] S. Oh, Z. Lei, W.-C. Lee, P. Mitra, and J. Yen, "CV-PCR: a context-guided value-driven framework for patent citation recommendation," in *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, 2013, pp. 2291–2296.

[252] A. Singhal, R. Kasturi, V. Sivakumar, and J. Srivastava, "Leveraging Web Intelligence for Finding Interesting Research Datasets," in *Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM International Joint Conferences on*, 2013, vol. 1, pp. 321–328.

[253] B. Gipp and J. Beel, "Citation Based Plagiarism Detection - a New Approach to Identify Plagiarized Work Language Independently," in *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia*, 2010, pp. 273–274.

[254] S. Zhan, A. Byung-Ryul, E. Ki-Yol, K. Min-Koo, K. Jin-Pyung, and K. Moon-Kyun, "Plagiarism Detection Using the Levenshtein Distance and Smith-Waterman Algorithm," in *Proceedings of the 3rd International Conference on Innovative Computing Information and Control*, 2008, pp. 569–569.

[255] M. Zini, M. Fabbri, M. Moneglia, and A. Panunzi, "Plagiarism Detection through Multilevel Text Comparison," in *Proceedings of the 2nd Conference on Automated Production of Cross Media Content for Multi-Channel Distribution*, 2006, pp. 181–185.

[256] M. Ley and P. Reuther, "Maintaining an online bibliographical database: The problem of data quality," *EGC'2006, Actes des sixièmes journées Extraction et Gestion des Connaissances*, pp. 17–20, 2006.

[257] J. Beel, S. Langer, M. Genzmehr, and C. Müller, "Docears PDF Inspector: Title Extraction from PDF files," in *Proceedings of the 13th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL'13)*, 2013, pp. 443–444.

[258] J. Beel, B. Gipp, A. Shaker, and N. Friedrich, "SciPlore Xtract: Extracting Titles from Scientific PDF Documents by Analyzing Style Information (Font Size)," in *Research and Advanced Technology for Digital Libraries, Proceedings of the 14th European Conference on Digital Libraries (ECDL'10)*, 2010, vol. 6273, pp. 413–416.

[259] H. Han, C. L. Giles, E. Manavoglu, H. Zha, Z. Zhang, and E. A. Fox, "Automatic document metadata extraction using support vector machines," in *Proceedings of the 3rd ACM/IEEE-CS Joint Conference on Digital libraries*, 2003, pp. 37–48.

[260] Y. Hu, H. Li, Y. Cao, L. Teng, D. Meyerzon, and Q. Zheng, "Automatic extraction of titles from general documents using machine learning," *Information Processing and Management*, vol. 42, no. 5, pp. 1276–1293, 2006.

[261] F. Peng and A. McCallum, "Information extraction from research papers using conditional random fields," *Information Processing and Management*, vol. 42, no. 4, pp. 963–979, 2006.

[262] S. Lawrence, C. L. Giles, and K. D. Bollacker, "Autonomous citation matching," in *Proceedings of the third annual conference on Autonomous Agents*, 1999, pp. 392–393.

[263] P. Lops, M. Gemmis, and G. Semeraro, "Content-based recommender systems: State of the art and trends," *Recommender Systems Handbook*, pp. 73–105, 2011.

[264] F. Ricci, L. Rokach, B. Shapira, and K. B. Paul, *Recommender systems handbook*. Springer, 2011, pp. 1–35.

[265] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen, "Collaborative filtering recommender systems," *Lecture Notes In Computer Science*, vol. 4321, p. 291, 2007.

[266] P. H. Rossi, M. W. Lipsey, and H. E. Freeman, *Evaluation: A systematic approach*, 7th edition. Sage publications, 2004.

[267] G. Gorrell, N. Ford, A. Madden, P. Holdridge, and B. Eaglestone, "Countering method bias in questionnaire-based user studies," *Journal of Documentation*, vol. 67, no. 3, pp. 507–524, 2011.

[268] G. Leroy, *Designing User Studies in Informatics*. Springer, 2011.

[269] A. Said, D. Tikk, Y. Shi, M. Larson, K. Stumpf, and P. Cremonesi, "Recommender systems evaluation: A 3d benchmark," in *ACM RecSys 2012 Workshop on Recommendation Utility Evaluation: Beyond RMSE, Dublin, Ireland*, 2012, pp. 21–23.

[270] P. Cremonesi, F. Garzotto, and R. Turrin, "Investigating the persuasion potential of recommender systems from a quality perspective: An empirical study," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 2, no. 2, pp. 1–11, 2012.

[271] P. Cremonesi, F. Garzotto, S. Negro, A. V. Papadopoulos, and R. Turrin, "Looking for 'good' recommendations: A comparative evaluation of recommender systems," in *Human-Computer Interaction–INTERACT 2011*, Springer, 2011, pp. 152–168.

[272] A. C. Burns and R. F. Bush, *Marketing Research*, 7th edition. Prentice Hall, 2013.

[273] J. L. Loeppky, J. Sacks, and W. J. Welch, "Choosing the sample size of a computer experiment: A practical guide," *Technometrics*, vol. 51, no. 4, 2009.

[274] H. Zheng, D. Wang, Q. Zhang, H. Li, and T. Yang, "Do clicks measure recommendation relevancy?: an empirical user study," in *Proceedings of the fourth ACM conference on Recommender systems*, 2010, pp. 249–252.

[275] A. Karatzoglou, L. Baltrunas, and Y. Shi, "Learning to Rank for Recommender Systems," in *Proceedings of the 7th ACM Conference on Recommender Systems*, 2013, pp. 493–494.

[276] J. A. Konstan and J. Riedl, "Recommender systems: from algorithms to user experience," *User Modeling and User-Adapted Interaction*, pp. 1–23, 2012.

[277] J. Matejka, W. Li, T. Grossman, and G. Fitzmaurice, "CommunityCommands: command recommendations for software applications," in *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, 2009, pp. 193–202.

[278] A. M. Rashid, I. Albert, D. Cosley, S. K. Lam, S. M. McNee, J. A. Konstan, and J. Riedl, "Getting to know you: learning new user preferences in recommender systems," in *Proceedings of the 7th international conference on Intelligent user interfaces*, 2002, pp. 127–134.

[279] W. Hersh, A. Turpin, S. Price, B. Chan, D. Kramer, L. Sacherek, and D. Olson, "Do batch and user evaluations give the same results?," in *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, 2000, pp. 17–24.

[280] W. R. Hersh, A. Turpin, L. Sacherek, D. Olson, S. Price, B. Chan, and D. Kraemer, "Further Analysis of Whether Batch and User Evaluations Give the Same Results with a Question-Answering Task.," in *In Proceedings of the Ninth Text REtrieval Conference (TREC 9)*, 2000, pp. 16–25.

[281] A. Said, "Evaluating the Accuracy and Utility of Recommender Systems," *PhD Thesis. Technische Universität Berlin*, 2013.

[282] A. H. Turpin and W. Hersh, "Why batch and user evaluations do not give the same results," in *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, 2001, pp. 225–231.

[283] D. Jannach, L. Lerche, F. Gedikli, and G. Bonnin, "What Recommenders Recommend–An Analysis of Accuracy, Popularity, and Sales Diversity Effects," in *User Modeling, Adaptation, and Personalization*, Springer, 2013, pp. 25–37.

[284] B. P. Knijnenburg, M. C. Willemsen, Z. Gantner, H. Soncu, and C. Newell, "Explaining the user experience of recommender systems," *User Modeling and User-Adapted Interaction*, vol. 22, no. 4–5, pp. 441–504, 2012.

[285] D. Jannach, M. Zanker, M. Ge, and M. Gröning, "Recommender Systems in Computer Science and Information Systems – A Landscape of Research," in *Proceedings of the 13th International Conference, EC-Web*, 2012, pp. 76–87.

[286] N. Good, J. B. Schafer, J. A. Konstan, A. Borchers, B. Sarwar, J. Herlocker, and J. Riedl, "Combining collaborative filtering with personal agents for better recommendations," in *Proceedings of the National Conference on Artificial Intelligence*, 1999, pp. 439–446.

[287] J. L. Herlocker, J. A. Konstan, A. Borchers, and J. Riedl, "An algorithmic framework for performing collaborative filtering," in *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, 1999, pp. 230–237.

[288] L. Palopoli, D. Rosaci, and G. M. Sarné, "A Multi-tiered Recommender System Architecture for Supporting E-Commerce," in *Intelligent Distributed Computing VI*, Springer, 2013, pp. 71–81.

[289] Y.-L. Lee and F.-H. Huang, "Recommender system architecture for adaptive green marketing," *Expert Systems with Applications*, vol. 38, no. 8, pp. 9696–9703, 2011.

[290] M. E. Prieto, V. H. Menéndez, A. A. Segura, and C. L. Vidal, "A recommender system architecture for instructional engineering," in *Emerging Technologies and Information Systems for the Knowledge Society*, Springer, 2008, pp. 314–321.

[291] S. Bhatia, C. Caragea, H.-H. Chen, J. Wu, P. Treeratpituk, Z. Wu, M. Khabsa, P. Mitra, and C. L. Giles, "Specialized Research Datasets in the CiteSeerx Digital Library," *D-Lib Magazine*, vol. 18, no. 7/8, 2012.

[292] K. Jack, M. Hristakeva, R. G. de Zuniga, and M. Granitzer, "Mendeley's Open Data for Science and Learning: A Reply to the DataTEL Challenge," *International Journal of Technology Enhanced Learning*, vol. 4, no. 1/2, pp. 31–46, 2012.

[293] J. S. Breese, D. Heckerman, and C. Kadie, "Empirical analysis of predictive algorithms for collaborative filtering," in *Proceedings of the 14th conference on Uncertainty in Artificial Intelligence*, 1998, no. MSR-TR-98–12, pp. 43–52.

[294] G. Karypis, "Evaluation of item-based top-n recommendation algorithms," in *Proceedings of the tenth international conference on Information and knowledge management*, 2001, pp. 247–254.

[295] A. Casadevall and F. C. Fang, "Reproducible science," *Infection and immunity*, vol. 78, no. 12, pp. 4972–4975, 2010.

[296] J. Rehman, "Cancer research in crisis: Are the drugs we count on based on bad science?," *http://www.salon.com/2013/09/01/is_cancer_research_facing_a_crisis/*, 2013.

[297] C. Drummond, "Replicability is not reproducibility: nor is it good science," in *Proc. of the Evaluation Methods for MachineLearning Workshop at the 26th ICML*, 2009.

[298] A. Al-Maskari, M. Sanderson, and P. Clough, "The relationship between IR effectiveness measures and user satisfaction," in *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 2007, pp. 773–774.

[299] B. P. Knijnenburg, M. C. Willemsen, and A. Kobsa, "A pragmatic procedure to support the user-centric evaluation of recommender systems," in *Proceedings of the fifth ACM conference on Recommender systems*, 2011, pp. 321–324.

[300] P. Pu, L. Chen, and R. Hu, "A user-centric evaluation framework for recommender systems," in *Proceedings of the fifth ACM conference on Recommender systems*, 2011, pp. 157–164.

[301] P. Pu, L. Chen, and R. Hu, "Evaluating recommender systems from the user's perspective: survey of the state of the art," *User Modeling and User-Adapted Interaction*, pp. 1–39, 2012.

[302] M. D. Ekstrand, M. Ludwig, J. A. Konstan, and J. T. Riedl, "Rethinking the recommender research ecosystem: reproducibility, openness, and LensKit," in *Proceedings of the fifth ACM conference on Recommender systems*, 2011, pp. 133–140.

[303] J. A. Konstan and G. Adomavicius, "Toward identification and adoption of best practices in algorithmic recommender systems research," in *Proceedings of the International Workshop on Reproducibility and Replication in Recommender Systems Evaluation*, 2013, pp. 23–28.

[304] R. Burke, "Hybrid recommender systems: Survey and experiments," *User modeling and user-adapted interaction*, vol. 12, no. 4, pp. 331–370, 2002.

[305] S. Perugini, M. A. Gonçalves, and E. A. Fox, "Recommender systems research: A connection-centric survey," *Journal of Intelligent Information Systems*, vol. 23, no. 2, pp. 107–143, 2004.

[306] I. Torre, "Adaptive systems in the era of the semantic and social web, a survey," *User Modeling and User-Adapted Interaction*, vol. 19, no. 5, pp. 433–486, 2009.

[307] M. Zanker, M. Jessenitschnig, D. Jannach, and S. Gordea, "Comparing recommendation strategies in a commercial context," *IEEE Intelligent Systems*, vol. 22, no. 3, pp. 69–73, 2007.

[308] E. Rich, "User modeling via stereotypes," *Cognitive science*, vol. 3, no. 4, pp. 329–354, 1979.

[309] M. Barla, "Towards social-based user modeling and personalization," *Information Sciences and Technologies Bulletin of the ACM Slovakia*, vol. 3, pp. 52–60, 2011.

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

[311] D. Mattioli, "On Orbitz, Mac users steered to pricier hotels," *Wall Street Journal*, vol. http://online.wsj.com/news/articles/SB10001424052702304458604577488822667325882, 2012.

[312] J. Beel, "Towards Effective Research-Paper Recommender Systems and User Modeling based on Mind Maps," *PhD Thesis. Otto-von-Guericke Universität Magdeburg*, 2015.

[313] J. Beel, S. Langer, G. M. Kapitsaki, C. Breitinger, and B. Gipp, "Exploring the Potential of User Modeling based on Mind Maps," in *Proceedings of the 23rd Conference on User Modelling, Adaptation and Personalization (UMAP)*, 2015, vol. 9146, pp. 3–17.

[314] J. Beel, B. Gipp, and E. Wilde, "Academic Search Engine Optimization (ASEO): Optimizing Scholarly Literature for Google Scholar and Co.," *Journal of Scholarly Publishing*, vol. 41, no. 2, pp. 176–190, 2010.

[315] W. Paik, S. Yilmazel, E. Brown, M. Poulin, S. Dubon, and C. Amice, "Applying natural language processing (nlp) based metadata extraction to automatically acquire user preferences," in *Proceedings of the 1st international conference on Knowledge capture*, 2001, pp. 116–122.

[316] Y. AlMurtadha, M. N. Sulaiman, N. Mustapha, and N. I. Udzir, "Improved web page recommender system based on web usage mining," in *Proceedings of the 3rd International Conference on Computing and Informatics (ICOCI)*, 2011, pp. 8–9.

[317] Y. Seroussi, "Utilising user texts to improve recommendations," *User Modeling, Adaptation, and Personalization*, pp. 403–406, 2010.

[318] Y. Seroussi, I. Zukerman, and F. Bohnert, "Collaborative inference of sentiments from texts," *User Modeling, Adaptation, and Personalization*, pp. 195–206, 2010.

[319] F. Esposito, S. Ferilli, T. M. A. Basile, and N. D. Mauro, "Machine Learning for Digital Document Processing: from Layout Analysis to Metadata Extraction," *Studies in Computational Intelligence (SCI)*, vol. 90, pp. 105–138, 2008.

[320] C. K. Shin and D. Doermann, "Classification of document page images based on visual similarity of layout structures," in *Proceedings of the SPIE Document Recognition and Retrieval VII*, 2000, pp. 182–190.

[321] D. Buttler, "A short survey of document structure similarity algorithms," in *Proceedings of the 5th International Conference on Internet Computing*, 2004.

[322] S. Brin and L. Page, "The anatomy of a large-scale hypertextual Web search engine," *Computer networks and ISDN systems*, vol. 30, no. 1–7, pp. 107–117, 1998.

[323] O. A. McBryan, "GENVL and WWWW: Tools for Taming the Web," in *Proceedings of the First International World Wide Web Conference*, 1994, vol. 341.

[324] S. Shi, F. Xing, M. Zhu, Z. Nie, and J.-R. Wen, "Anchor Text Extraction for Academic Search," in *Proceedings of the 2009 Workshop on Text and Citation Analysis for Scholarly Digital Libraries (ACL-IJCNLP 2009)*, 2009, pp. 10–18.

[325] C. D. Manning, P. Raghavan, and H. Schütze, *An Introduction to Information Retrieval*, Online Edition. Cambridge University Press, Cambridge, England, 2009.

[326] I. G. Councill, C. L. Giles, and M. Y. Kan, "ParsCit: An open-source CRF reference string parsing package," in *Proceedings of LREC*, 2008, vol. 2008, pp. 661–667.

[327] S. Marinai, "Metadata Extraction from PDF Papers for Digital Library Ingest," *10th International Conference on Document Analysis and Recognition*, 2009.

[328] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry, "Using collaborative filtering to weave an information Tapestry," *Communications of the ACM*, vol. 35, no. 12, pp. 61–70, 1992.

[329] T. A. Brooks, "Private acts and public objects: an investigation of citer motivations," *Journal of the American Society for Information Science*, vol. 36, no. 4, pp. 223–229, 1985.

[330] M. Liu, "Progress in documentation the complexities of citation practice: a review of citation studies," *Journal of Documentation*, vol. 49, no. 4, pp. 370–408, 1993.

[331] M. H. MacRoberts and B. MacRoberts, "Problems of Citation Analysis," *Scientometrics*, vol. 36, pp. 435–444, 1996.

[332] S. Sosnovsky and D. Dicheva, "Ontological Technologies for User Modeling," *International Journal of Metadata, Semantics and Ontologies*, vol. 5, no. 1, pp. 32–71, 2010.

[333] S. S. Sundar, A. Oeldorf-Hirsch, and Q. Xu, "The bandwagon effect of collaborative filtering technology," in *CHI'08 Extended Abstracts on Human Factors in Computing Systems*, 2008, pp. 3453–3458.

[334] B. Mehta, T. Hofmann, and P. Fankhauser, "Lies and propaganda: detecting spam users in collaborative filtering," in *Proceedings of the 12th international conference on Intelligent user interfaces*, 2007, pp. 14–21.

[335] B. Mehta, T. Hofmann, and W. Nejdl, "Robust collaborative filtering," in *Proceedings of the 2007 ACM conference on Recommender systems*, 2007, pp. 49–56.

[336] B. Mehta and W. Nejdl, "Attack resistant collaborative filtering," in *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, 2008, pp. 75–82.

[337] K. Sugiyama and M. Y. Kan, "Serendipitous recommendation for scholarly papers considering relations among researchers," in *Proceeding of the 11th annual international ACM/IEEE joint conference on Digital libraries*, 2011, pp. 307–310.

[338] R. Burke, "Hybrid web recommender systems," *The adaptive web*, pp. 377–408, 2007.

[339] P. Ahlgren and C. Colliander, "Document–Document Similarity Approaches and Science Mapping: Experimental Comparison of Five Approaches," *Journal of Informetrics*, vol. 3, no. 1, pp. 49–63, 2009.

[340] K. M. Hammouda and M. S. Kamel, "Phrase-based document similarity based on an index graph model," in *Data Mining, 2002. ICDM 2003. Proceedings. 2002 IEEE International Conference on*, 2002, pp. 203–210.

[341] M. D. Lee, B. Pincombe, and M. Welsh, "An Empirical Evaluation of Models of Text Document Similarity," in *Proceedings of the 27th Annual Conference of the Cognitive Science Society*, 2005, pp. 1254–1259.

[342] A. Tsymbal, "The problem of concept drift: definitions and related work," *Computer Science Department, Trinity College Dublin*, 2004.

[343] P. Victor, M. De Cock, and C. Cornelis, "Trust and recommendations," in *Recommender systems handbook*, Springer, 2011, pp. 645–675.

[344] K. Verbert, D. Parra, P. Brusilovsky, and E. Duval, "Visualizing recommendations to support exploration, transparency and controllability," in *Proceedings of the 2013 international conference on Intelligent user interfaces*, 2013, pp. 351–362.

[345] S. Lam, D. Frankowski, and J. Riedl, "Do you trust your recommendations? An exploration of security and privacy issues in recommender systems," *Emerging Trends in Information and Communication Security*, pp. 14–29, 2006.

[346] C. N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen, "Improving recommendation lists through topic diversification," in *Proceedings of the 14th international conference on World Wide Web*, 2005, pp. 22–32.

[347] R. Burke and M. Ramezani, "Matching recommendation technologies and domains," in *Recommender Systems Handbook*, Springer, 2011, pp. 367–386.

[348] Y. Zhang, J. Callan, and T. Minka, "Novelty and redundancy detection in adaptive filtering," in *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, 2002, pp. 81–88.

[349] L. Pizzato, T. Rej, K. Yacef, I. Koprinska, and J. Kay, "Finding someone you will like and who won't reject you," *User Modeling, Adaption and Personalization*, pp. 269–280, 2011.

[350] D. Cosley, S. K. Lam, I. Albert, J. A. Konstan, and J. Riedl, "Is seeing believing? How recommender system interfaces affect users' opinions," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2003, pp. 585–592.

[351] J. L. Herlocker, J. A. Konstan, and J. Riedl, "Explaining collaborative filtering recommendations," in *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, 2000, pp. 241–250.

[352] F. Carmagnola, F. Cena, and C. Gena, "User model interoperability: a survey," *User Modeling and User-Adapted Interaction*, pp. 1–47, 2011.

[353] J. Beel, S. Langer, B. Gipp, and A. Nürnberger, "The Architecture and Datasets of Docear's Research Paper Recommender System," *D-Lib Magazine*, vol. 20, no. 11/12, 2014.

[354] S. Langer and J. Beel, "The Comparability of Recommender System Evaluations and Characteristics of Docear's Users," in *Proceedings of the Workshop on Recommender Systems Evaluation: Dimensions and Design (REDD) at the 2014 ACM Conference Series on Recommender Systems (RecSys)*, 2014, pp. 1–6.

Additional Information

Academic Platforms

Authors

Joeran Beel

Bela Gipp

Stefan Langer

Corinna Breitinger

BibTeX

```
@ARTICLE{Beel2015r,
  author = {Joeran Beel and Bela Gipp and Stefan Langer and Corinna Breitinger},
  title = {Research Paper Recommender Systems: A Literature Survey},
  journal = {International Journal on Digital Libraries},
  year = {2015},
  pages = {1-34},
  doi = {10.1007/s00799-015-0156-0},
  issn = {1432-5012},
}
```

RefMan (RIS)

```
TY  - JOUR
T1  - Research-paper recommender systems: a literature survey
A1  - Beel, Joeran
A1  - Gipp, Bela
A1  - Langer, Stefan
A1  - Breitinger, Corinna
JO  - International Journal on Digital Libraries
SP  - 1
EP  - 34
SN  - 1432-5012
Y1  - 2015
PB  - Springer
ER  -
```

Endnote

```
%0 Journal Article
%T Research-paper recommender systems: a literature survey
%A Beel, Joeran
%A Gipp, Bela
%A Langer, Stefan
```