Applied Informetrics for Digital Libraries: An Overview of Foundations, Problems and Current Approaches

Philipp Schaer[*]

Abstract: *»Zur Anwendung von Informetrie für Digitale Bibliotheken. Ein Überblick über Grundlagen, Probleme und aktuelle Zugänge«*. The foundation of every research project is a comprehensive literature review. The search for scientific literature in information systems is a discipline at the intersection of information retrieval and digital libraries; recent user studies in both fields show two typical weaknesses of the classical IR approach: ranking of retrieved and maybe relevant documents and the language problem during the query formulation phase. At the same time the performance of traditional retrieval systems that rely primarily on textual document and query features have been stagnating for years, as could be observed in IR evaluation campaigns such as TREC or CLEF. Therefore alternative approaches to surpass these two problem fields are needed. Recent developments in the area of applied informetrics show very promising effects by using long-known informetric and bibliometric methods like the analysis of power-law distributions described by Lotka's, Zipf's or Bradford's laws, or the application of co-occurrences analysis for entities like authors, journals or references of scientific literature. This work will concentrate on the description of the open problems and the current approaches to surpass these by using applied informetrics methodologies.

Keywords: Digital libraries, informetrics, Power Law, Bradford's Law, Lotka's Law, Zipf's Law, information retrieval, co-occurrence analysis.

1. Introduction

The foundation of every research project is a comprehensive literature review. It is used to extract or to establish a research problem, to get an overview of the state of the art in a specific field or to get a feeling for loose ends in current research methodologies or approaches. While in former times, the entrance point for every literature review were private, public or institutional libraries, for more than two decades these physical libraries have become more often complemented by their digital counterparts. The young generation of scholars almost completely changed their information behaviour during these two dec-

[*] Philipp Schaer, GESIS – Leibniz Institute for the Social Sciences, Unter Sachsenhausen 6-8, 50667 Cologne, Germany; philipp.schaer@gesis.org.

ades compared to established scholars before them. Especially the search for scientific literature is determined by the so-called "now or never" mentality where everything that is not available in real time is discarded or ignored. This might be related to the availability of electronic entrance points to scientific documents in the form of web search engines or more general digital libraries (DL). These systems and entrance points are a competitive counterpart to the established physical libraries since they can satisfy the information need of their users faster and more accurately.

Digital library systems share typical structures and characteristics. They are in between the highly chaotic and unstructured characteristics of the Web and the highly structured attributes of specialised information systems. As coined by Fox et al.:

> One dichotomy often posed about DLs is Managed vs. Comprehensive. Thus, a library is managed while the WWW is unmanaged (but closer to being comprehensive). [...] we generally use the term structure. We argue that DLs must be organized, thus having a moderate degree of structure" (Fox, Goncalves and Shen 2012, 6).

By offering this moderate degree of structure, digital libraries try to model a pragmatic and user-centered way to access (specialized or scientific) digital information. User studies in the field of information access and information behaviour in DLs show that, although more and more academics use these kinds of systems, they are unsure about the possibilities and limitations and are in need of a helping hand to deal with the problem of information overload. Traditional information retrieval techniques are not able to cope with these kinds of problems and to satisfy the users' needs.

Another discipline that is using the structured information and metadata available in DL systems is informetrics, which is "the study of the quantitative aspects of information in any form, not just records or bibliographies, and in any social group, not just scientists" (Tague-Sutcliffe 1992, 1). This quantitative research discipline is making heavy use of the available data to study science of information processes themselves but not to support users in information systems. Today the most common use of informetrics and its neighbouring discipline bibliometrics is to rank journals, authors, institutions by different indicators – mostly with the aid of citation analysis. The rankings derived from these indicators and their usage e.g. by administration or governmental institutions are highly controversial. Further, one can say that informetrics and one of the most prominent measures of it – the h-index – are harming the scientific community more than helping it. As a very drastic example, the German Sociological Association (Deutsche Gesellschaft für Soziologie, DGS) refused to contribute to the CHE Ranking that uses different indicators from research and teaching to rank universities[1]. In this work we will not focus on the social implications of numbers like the h-index but on the mathematical foundations and their implications for information systems, and how to make a difference and actually help the users.

The following first part of this work will give an overview on open retrieval problems in digital libraries such as the so-called language problem or the problem sets of ranking and relevance judgements. The focus will be on state-of-the-art methodologies in this area and the advantages and disadvantages of each approach. The second part consists of an overview on informetric methods that are used to analyse the structure, development or impact of scientific work. The last part will give an outlook to the possible combination of these two disciplines to show the mutual benefits for digital libraries, their users and the discipline of informetrics itself.

2. Open Problems in Digital Libraries

Although usage rates of DL systems are growing rapidly, current user studies show that scientific DL do not support users in their actual information-seeking behaviour, which is often shaped through their experience with web search engines (Siegfried and Flieger 2011). Web search engines and specialized information databases are the most common retrieval tools to access scientific information, whereas their specific usage context is mostly different. Free and exploratory search types are more frequent in specialized systems due to their more advanced indexation and data quality. In the case of web search engines, they are commonly more used for retrieval tasks when it comes to a known item search, i.e. the search for a specific paper or data set where title and author are known to the searcher. The actual user needs of DL system users show conflicting goals: on the one hand, academics want to and actually do more and more retrieval work on their own in the whole range of systems available. On the other hand, they are increasingly unable to cope with the variety of systems and possibilities of these systems with their vast amount of information. In the previously cited study by Siegfried and Flieger more than 53% of the interviewed academics had serious problems in judging the relevance of the results produced from a retrieval session. Most interestingly, the criteria to judge on the relevance are all features that are not included in actual information retrieval models. The three most important document features to judge relevance are: (1) currentness of the work, (2) reputation of the journal in which the work was published, and (3) the reputation of the authors. Modern information retrieval

1 <http://www.soziologie.de/fileadmin/user_upload/DGS_Redaktion_BE_FM/Aktuelles/GSA_Statement_on_CHE-Ranking.pdf> (accessed October 12, 2013).

engines, which are the foundation of every digital library system, do not include these (sometimes very abstract and vague) document features to rank their documents.

2.1 The Language Problem

One problem present in every search process that becomes very visible in scientific searches is the so-called language problem: During the formulization of an information need, a searcher can (in theory) use the unlimited possibilities of human language to express himself. This is especially true when expressing a scientific information need using domain-specific expressions that are very unique and context-sensitive. Every scientific community and discipline has developed its own special vocabulary that is not commonly used by other researchers from other domains. With regards to digital libraries this problem becomes even worse as the problem to express something is also valid for the group of indexers that try to address keywords or categories to each information object. Although professional indexers use specialized toolkits like thesauri or classification systems to allow a consistent documentation of scientific texts, the human factor in this process cannot be ignored. Tools like thesauri and classifications are known as knowledge organization systems (KOS). These KOSs try to control the vagueness of human language by defining a strict rule set and controlled vocabularies but are applied (after all) by humans.

When we think of retrieval processes (and therefore academic searches) as "fundamentally a linguistic process" (Blair 2003, 3), these toolkits have to be understood and used by searchers. This is all too often not the case and therefore searchers tend to surpass the language problem by pragmatic approaches as expressed in an interview study by Aula et al: "I choose search terms based not specifically on the information I want, but rather on how I could imagine someone wording […] that information" (Aula, Jhaveri and Käki 2005, 589).

Classics methods to surpass these query formulation problems and to actively support users in expressing their information needs are generally known as query expansion (QE). A wide range of possible query expansion and search term recommender systems are known in the information retrieval community (Efthimiadis 1996). In modern web search engines these recommenders are a common tool and became omnipresent the situation is different in digital libraries. Only few DL systems implement interactive variations of query expansion, of which two different subtypes can be defined: term suggestions and query suggestions systems. In a term suggestion system only single words or phrases are presented to the user while a query suggestion system can suggest a whole query string for which they often make use of query log analysis; however it can also be implemented to suggest queries based on document corpora. Here the structure of digital libraries and their metadata sets can be used as training set for these recommender systems. The main task of a term suggestion systems is to support users in the process of expressing their information need and supporting them in the formulation of a useful query. These systems try to suggest terms that are closely related both to the user's initial query term as well as to the semantic background that is encoded in the KOS of the digital library.

Besides the described language problem and the vagueness of language, other factors come into play when trying to formulate an information need. As previously shown other information entities like non-textual information on the reputation of authors or journals might be used to describe an information need. This is included in the principle of polyrepresentation (Ingwersen 1996; Ingwersen and Järvelin 2005) in which the diversity in both the information needs and in the information space itself is used to increase retrieval performance. In the polyrepresentation principle the cognitive overlaps of different aspects of an information object are in focus to represent the diversity. In the following section we will follow on these diversity aspects.

2.2 Term-Based Relevance Ranking

Traditional document retrieval (in the form of ad-hoc retrieval) has reached a point where no significant performance boosts can be measured in typical retrieval evaluation setups. The usual ranking methods like TF*IDF, BM25 or statistical language models are not generally superior to each other when retrieving documents and have all reached a high level of effectiveness. This is shown in retrieval campaigns like TREC or CLEF where no large increases are recognizable and the retrieval performance of the top systems are stagnating (Armstrong et al. 2009). Therefore other ranking and retrieval approaches are needed.

As described in the earlier sections, traditional digital libraries are widely filled with curated and structured information on diverse documents but lack the features to really make use of all the metadata to support users in the satisfaction of their information need. This is not only true for aspects like the language problem, but also for ranking techniques that are the heart of every retrieval engine. In the last years web search has brought up a variety of alternative (non-textual) ranking methods that extended the common text-based ranking mechanisms used in nearly every digital library system today. All these ranking methods try to calculate a relevance score that can be used to sort a result set of documents in such a way that the most relevant document (in regard to the initial query) are on the top of the result list.

According to Borlund (2003) relevance is one of the fundamental concepts in information science and its application in information retrieval. The definition of relevance is hard to come by since it consists of different classes, types, grades, criteria and levels of what relevance is. In her very broad literature study on the different aspects Borlund therefore chooses to coin the concept of "multidimensionality of relevance". The most general way to differentiate relevance is to distinguish between objective/system-related and subjective/human-related relevance criteria. In relation to Schamber, Eisenberg, and Nilan (1990), Borlund (2003, 913) concludes with three central key features of relevance: (1) relevance is a multidimensional and cognitive concept that is dependent on the personal information need of a user; (2) relevance is a dynamic concept that is dependent on the judgement of the relation between the information itself and the information need to a specific time; (3) relevance is a complex concept that can be measured in a systematic way as long as the user perspective is considered.

All three aspects share the relation to an actual user's information need and a user-centered approach. Since a retrieval system cannot guess the real users' needs from a more technical view, four different types of relevance are used: (1) situated relevance, the usefulness of an information in a specific situation; (2) subjective relevance (sometimes called pertinence), the usefulness of an information to a specific person; (3) objective relevance which in comparison to subjective relevance has to be of usefulness for more than one person; (4) system relevance, the algorithmically computed relevance estimation of a retrieval system. This estimation score is called retrieval status value (RSV) and is used by retrieval systems to sort the result set according to this score. To measure the performance of a retrieval system, these RSVs are compared to the relevance judgements of human assessors (sometimes called raters or judges).

Usually digital libraries are full of curated documents and rich sets of metadata to describe these documents, but only little of the information available is actually used to support the retrieval or ranking process. The methods used in today's systems can be divided into (1) term-based/textual and (2) structural/non-textual ranking methods.

The most traditional methods to establish term-based rankings used in so-called ad-hoc retrieval setups are the extended Boolean model, the vector space model, probabilistic models and statistical language models (Manning, Raghavan and Schütze 2008). These all try to estimate the relevance of a document in regard to their connection to the initial query terms. Therefore we call these the term-based ranking methods. Rankings methods that use the structure of a document or group of documents like the PageRank method use other feature to estimate the relevance. Next to the popular PageRank, methods like link analysis or centrality calculation are used. We call these the structural or non-textual ranking methods although non-textual is often mixed-up with multimedia feature. In this work we think of non-textual features like the centrality of an author in a co-authorship network or the degree to which a document is linked to other documents in the result set. Metzler (2011, 5) compiled an overview on these non-textual features that are used for ranking like: the PageRank, inlink counts from a network analysis, readability of a text, the probability of spam. Other features might be the productivity of authors and journals or the semantic relatedness of concepts within the documents. Some of these might be already be useable in the form of query terms (like author names) but concepts like centrality, productivity or reputation are not included in the ranking models presented before. Some of these concepts and feature are the result of informetric analyses and will be presented in a later section 4.

2.3 Measuring Retrieval Effectiveness

To measure the performance of retrieval systems the information retrieval community developed and refined a special evaluation setup that is commonly known as the Cranfield paradigm or as the IR lab evaluation. These lab evaluations make use of three specific artifacts: a fixed set of documents (the so-called corpus), a list of retrieval problems (called topics) and relevance assessments that judge documents in regard to these topics. All three pieces together form a test collection that can be used to evaluate and to compare different IR systems. One of the major advantages of these test collections and the Cranfield paradigm is the possibility to rerun previous evaluation settings something that is impossible in i.e. an interactive setting where actual users are involved.

One of the most criticised parts of this evaluation setup is the problem that lies in the relevance judgements. The problem with these judgements is that there is obviously no "golden standard." Given the possibility that more than person is asked to judge on the relevance of documents according to a specific topic, there can be a high amount of disagreement (Schaer 2012). Usually these disagreements are avoided by only having one assessor judge the relevance. A wide range of different of evaluation measures were developed during the last two decades that aim at very different parts of the retrieval process. The most well-known measures are precision and recall (for unsorted result sets) and cut off precision, mean average precision, normalized discounted cumulative gain and many others that will be discussed in Schaer (2013b).

3. Informetrics, Scientometrics, Bibliometrics

In the previous chapter, two open problems for academic retrieval systems and digital libraries were presented that involve the two key stages in every retrieval process: the query formulation phase and the ranking of the retrieved document set. While a lot of effort was put into the processing of textual queries and their relations to the documents, these common implementations are not very innovative as they ignore a most of the documents features from a structural/non-textual view.

According to Hobohm (2012) digital libraries should be enabled to derive knowledge out of them, in contrast to just retrieving documents or information carrying objects. He argues that besides the steps of processing, describing and storing documents that are well developed in today's DL systems, other aspects are not treated as equally important and are therefore underdeveloped. He mentions neglecting understanding and social knowledge as two examples for these underdeveloped aspects. An initial step would be to open up these systems to an additional level of analysis, like it is common practice in informetrics to derive additional knowledge out of document collections. Informetrics as the study of quantitative aspects of information offers a wide range of measures, models and tools to analyse bibliographic data and to derive information and knowledge out of them.

Besides informetrics, other terms like bibliometrics, scientometrics, and webometrics are used to describe the analysis of these quantitative aspects. Bibliometrics focus on the analysis of published documents like books and articles while scientometrics focus on general outcomes of scientific research in order to understand the dynamics of science and research processes. In webometrics similar methods are applied to analyse and understand information processes in the web. Although it might be obvious to analyse the available documents in digital libraries with these informetrics methods to get a deeper understanding and to generate value-added features for the users this is rarely done. Havemann (2009, 58) points out that the practical application of bibliometrics methods in the field of information retrieval are rare and don't exceed a Bradford analysis to decide on journal subscriptions for an institute.

According to Glänzel (2003, 9) only three practical fields of application for bibliometrical methods are known. The first one would be called bibliometrics for bibliometricans, where the main focus lies in methods and fundamental research interests in the field. Practical use cases are not central and the statistical, mathematical or methodological approaches are studied to deepen the understanding of the field itself. The second use case would be bibliometrics for other scientific disciplines. This is the most common use case and the field in which most bibliometric studies would be applied. Bibliometrics here is a central tool to answer questions or to solve problems in the area under consideration. The third use case is the most controversial one: bibliometrics for science policies. Here bibliometrics is used to do academic rankings for individuals, institutions, documents, journals and many more entities in the scientific and academic field. The rankings are most often crucial for funding or career options and are therefore criticised a lot. For politics the benefit is to allow comparison on a national or international level (e.g. to measure the influence or impact of German academics in Europe). Of course there are other use cases like for example the use of bibliometric studies in the field of history of science (Scharnhorst and Garfield 2010), but these kinds of studies are rare and are most often just special cases of the three applications mentioned before.

It is surprising to see so few real world applications of bibliometrics and informetrics methods because one of the most central key ideas behind digital libraries, documentation and the whole discipline of information science is to make actual use of the curated metadata and the insights that might be gained from them (Hjørland 2000, 510-2). With regards to the argument of Hobohm: Before you can extract knowledge out of digital libraries you need to find the information you want. Of course it is a crucial part to actually build and curate data sets for digital libraries but the main goal should be to make all these information objects available, findable and interpretable. Hjørland especially focuses on retrieval tasks and names this as the most useful domain to apply information science and informetrics insights to.

Figure 1: Overview of the Relationship between Academic Papers, Journals Authors, Terms and References

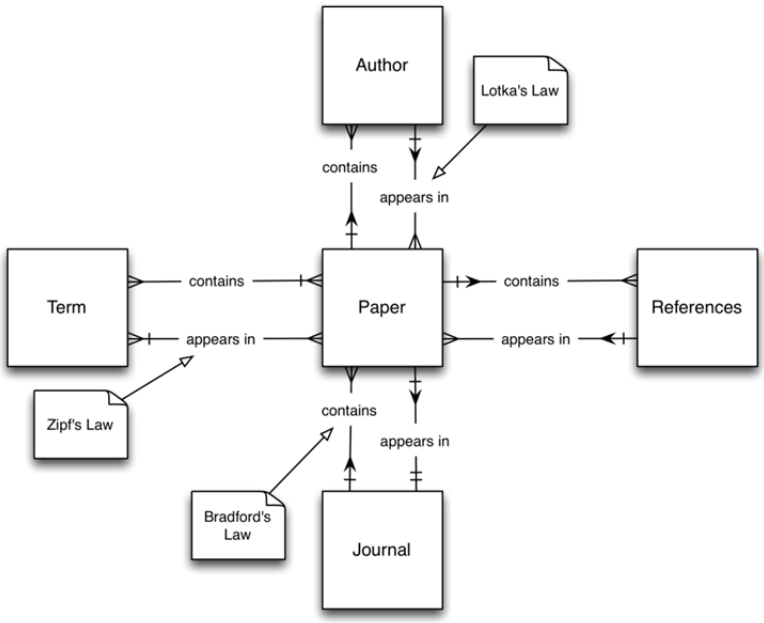

This figure was inspired by a graphic from Morris and Yen 2004, 5292.

One of the central principles in informetrics and in the analysis of the so-called information production process (IPP) is the so-called power-law distribution. This statistical distribution can be seen in bibliometric analyses of journal publications, the use of language and term frequencies (Newman 2005). Many man-made or naturally occurring phenomena show typical features that follow the equation

$$f(x) = cx^{-\alpha}.$$

Here x is the rank, c is most often constant and α is usually between 1 and 2 and the whole power-law is monotone. These laws describe a typical pattern where large events are rare and small events are common. When you apply this to the use of terms in a text this principle becomes clear. There are only few words, like “and” and “the” that occur very frequently, but many which occur rarely. When you plot these distributions on two logarithmic scales you see a straight line.

These power-laws are so fundamental for informetrics that they were (re-)discovered for many different entity relationships in the information production process (Egghe 2005). In figure 1 the most prominent informetric laws are marked in an overview of the typical bibliometric entities. In this figure we see the most central object in every analysis: the (academic) paper or document. This paper has a relationship to its author(s), its terms that are used to describe it, the journal or book it appears in and the references it might contain or might be contained in. As in many other man-made systems these relationships follow power-laws. Here the names that are best known are Zipf’s law (usage of words and terms in texts), Bradford’s law of scattering (centralisation of topical related publications in so core journals) and Lotka’s law (productivity of authors).

Another method that is used frequently in informetrics is co-occurrence analysis. Here the simultaneous appearance of two entities in one document or in a group of documents is analysed. In informetric studies these are typically co-citations, co-authorships, co-classifications or co-words. More generally one can explore item-item, entity-entity or item-entity relationships (Tijssen and Van Raan 1994). An item is usually a document and an entity are the corresponding attributes like journals, authors, and references (see figure 1).

4. Applied Informetrics for Digital Libraries

The direct application of informetric methods for the use in retrieval systems and digital libraries was first proposed by White (1981). He suggested reranking result sets according to the productivity of their corresponding journal (according to Bradford’s law of scattering). This way articles that were published in core journals (the most productive ones in a field) would be ranked before articles that were published in journals from the periphery (the less productive ones). Later this approach was evaluated and it was shown that indeed the core journals have a higher concentration of relevant documents that the ones in the periphery (Mayr 2013). Other works tried to use co-occurrence methods like bibliographic coupling and co-citation for document retrieval (Bichteler and Eaton 1980) and showed the general feasibility and potential of these approaches.

More recently authors like White (2007), Mutschke et al. (2011), Ingwersen (2012), and Schaer (2013a) try to push these ideas to a new level by studying the connections between the two fields of informetrics and information retrieval. For too long both disciplines were neighbours but did not actively cooperate with each other. Information retrieval was focusing on term-based topical relevance for too long and informetrics was seeing information retrieval as an auxiliary science.

Two examples should demonstrate the usefulness of cooperation between the two fields.

Schaer (2013a) presented two different search support systems and evaluated them with a lab evaluation using the IR test collection GIRT and iSearch with 150 and 65 topics, respectively. These two systems are (1) a query expansion that is based on the analysis of co-occurrences of document attributes and (2) a ranking mechanism that applies informetric analysis of the productivity of information producers in the information production process. Both systems were compared to a baseline system using the Solr search engine. Both methods showed positive effects when applying additional document attributes like author names, ISSN codes and controlled terms. The query expansion showed an improvement in precision (bpref +12%) and in recall (R +22%). The alternative ranking methods were able to compete with the baseline for author names and ISSN codes and were able to beat the baseline by using controlled terms (MAP +14%). A clear negative influence was seen when using entities like publishers or locations. Both methods were able to generate a substantially different sorting of the result set, measured using Kendall's τ ($\geq$0.8). So, additional to the improved relevance in the result list, the user can get a new and different view on the document set. Query expansion using author names, ISSN codes and thesaurus terms showed great potential that lies within the rich metadata sets of digital library systems. The proposed ranking methods could outperform standard relevance ranking methods after they were filtered by the existence of a so-called power law. This showed that the proposed ranking methods cannot be used universally in any case but require specific frequency distributions in the metadata. A connection between the underlying informetric laws of Bradford, Lotka and Zipf is made clear.

Strohman, Croft, and Jensen (2007) tried a new approach that went beyond the traditional short query paradigm and suggested that academic searchers might already know something about their information need. Maybe they already wrote a page or two about this specific information need. The main idea from Strohman et al. was to use these documents as query to not only retrieve topically similar documents to the query but to retrieve documents that might cite the query document. They tested their approach using the Rexa collection with 946,977 literature references (from which 105,601 contain full text information) and a total of 1.46 million citations. They used different document features to test their approach like publication year, text similarity, co-citation coupling, information on same authors, the Katz graph distance measure, and raw citation counts. Different combinations of these features where compared to a system that only worked on text similarity (like most retrieval systems). All experimental systems outperform text similarity (Wilcoxon, $p = 0.01$) using the MAP measure.

5. Conclusion

In this paper only a very short overview on the intersections of digital libraries, information retrieval and informetrics could be outlined. Only two examples of recent systems have been outlined, but they show the general feasibility and performance of alternative retrieval techniques that are based in informetric methods. The following paper will more go into the (methodical and mathematical) details. Although only few examples were presented in this short paper the momentum is for more cooperation is there. The information retrieval and informetrics communities move more closely together with combined workshops like "Computational Scientometrics"[2] (held at one of the most important information retrieval conferences CIKM 2013) and "Combining Bibliometrics and Information Retrieval"[3] (held at the biggest informetrics conference ISSI 2013).

With the existence of new information retrieval test collections that contain citation and bibliographic information like the iSearch collection (Lykke et al. 2010) or the ACL collection (Ritchie, Teufel and Robertson 2006) one might hope that the IR community gets an interest in these kind of systems again and that new approaches based on informetrics get developed and evaluated. The goals for all these approaches remain the same: Surpass the obvious drawbacks of existing systems by applying the methods and insight of informetric research.

References

Armstrong, Timothy G., Alistair Moffat, William Webber, and Justin Zobel. 2009. Improvements That Don't Add up: Ad-hoc Retrieval Results Since 1998. In *Proceeding of the 18th ACM Conference on Information and Knowledge Management*, 601-10. CIKM '09. Hong Kong, China: ACM <doi:10.1145/1645953.1646031>.

Aula, Anne, Natalie Jhaveri, and Mika Käki. 2005. Information Search and Reaccess Strategies of Experienced Web Users. In *Proceedings of the 14th Interna-*

[2] <http://www.cse.unt.edu/~ccaragea/CIKM-WS-13.html> (accessed October 12, 2013).
[3] <http://www.gesis.org/en/events/conferences/issiworkshop2013> (accessed October 12, 2013).

tional Conference on World Wide Web, 583-92. WWW '05. New York, NY, USA: ACM. doi:10.1145/1060745.1060831 <http://doi.acm.org/10.1145/1060745.1060831>.
Bichteler, Julie, and Edward A. Eaton. 1980. The Combined Use of Bibliographic Coupling and Cocitation for Document Retrieval. *Journal of the American Society for Information Science* 31 (4): 278-82 <doi:10.1002/asi.4630310408>.
Blair, David C. 2003. Information Retrieval and the Philosophy of Language. *Annual Review of Information Science and Technology* 37 (1): 3-50 <doi:10.1002/aris.1440370102>.
Borlund, Pia. 2003. The Concept of Relevance in IR. *Journal of the American Society for Information Science and Technology* 54 (10): 913-25 <doi:10.1002/asi.10286>.
Efthimiadis, Efthimis N. 1996. Query Expansion, ed. Martha E. Williams. *Annual Review of Information Systems and Technology (ARIST)* 31: 121-87.
Egghe, Leo. 2005. *Power Laws in the Information Production Process: Lotkaian Infometrics*. Library and Information Science Series. Oxford: Elsevier.
Fox, Edward A., Marcos Andre Goncalves, and Rao Shen. 2012. *Theoretical Foundations for Digital Libraries: The 5S Approach*. Morgan & Claypool Publishers.
Glänzel, Wolfgang. 2003. Bibliometrics as a Research Field <http://nsdl.niscair.res.in/bitstream/123456789/968/1/Bib_Module_KUL.pdf>.
Havemann, Frank. 2009. *Einführung in Die Bibliometrie*. Berlin: Gesellschaft für Wissenschaftsforschung <http://edoc.hu-berlin.de/oa/books/reMKADKkid1Wk/PDF/20uf7RZtM6ZJk.pdf>.
Hjørland, Birger. 2000. Library and Information Science: Practice, Theory, and Philosophical Basis. *Information Processing and Management* 36 (3): 501-31 <doi:10.1016/S0306-4573(99)00038-2>.
Hobohm, Hans-Christoph. 2012. Can Digital Libraries Generate Knowledge? *Historical Social Research* 37 (3): 218-29.
Ingwersen, Peter. 1996. Cognitive Perspectives of Information Retrieval Interaction: Elements of a Cognitive IR Theory. *Journal of Documentation* 52 (1): 3-50 <doi:10.1108/eb026960>.
Ingwersen, Peter. 2012. *Scientometric Indicators and Webometrics and the Polyrepresentation Principle in Information Retrieval*. Ess Ess Publication.
Ingwersen, Peter, and Kalervo Järvelin. 2005. *The Turn - Integration of Information Seeking and Retrieval in Context*. Dordrecht: Springer.
Lykke, Marianne, Birger Larsen, Haakon Lund, and Peter Ingwersen. 2010. Developing a Test Collection for the Evaluation of Integrated Search. In *Advances in Information Retrieval*, ed. Cathal Gurrin et al., 5993: 627-30. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer.
Manning, Christopher D., Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge, UK: Cambridge University Press <http://nlp.stanford.edu/IR-book/information-retrieval-book.html>.
Mayr, Philipp. 2013. Relevance Distributions Across Bradford Zones: Can Bradfordizing Improve Search? In *Proceedings of ISSI 2013*, ed. Juan Gorraiz et al., 2: 1493-505. Vienna, Austria.
Metzler, Donald. 2011. *A Feature-Centric View of Information Retrieval*, Vol. 27. The Information Retrieval Series. Berlin, Heidelberg: Springer.

Morris, Steven. A., and Gary G. Yen. 2004. Crossmaps: Visualization of Overlapping Relationships in Collections of Journal Papers. *Proceedings of the National Academy of Sciences* 101 (Suppl 1) (January 23): 5291-6 <doi:10.1073/pnas.0307604100>.

Mutschke, Peter, Philipp Mayr, Philipp Schaer, and York Sure. 2011. Science Models as Value-added Services for Scholarly Information Systems. *Scientometrics* 89 (1) (June): 349-64 <doi:10.1007/s11192-011-0430-x>.

Newman, M. E. J. 2005. Power Laws, Pareto Distributions and Zipf's Law. *Contemporary Physics* 46: 323-51 <doi:10.1080/00107510500052444>.

Ritchie, Anna, Simone Teufel, and Stephen Robertson. 2006. Creating a Test Collection for Citation-based IR Experiments. In *Proceedings of the Main Conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, 391-98. HLT-NAACL '06. Stroudsburg, PA, USA: Association for Computational Linguistics <doi: 10.3115/1220835.1220885>,. <http://dx.doi.org/10.3115/1220835.1220885>.

Schaer, Philipp. 2012. Better Than Their Reputation? On the Reliability of Relevance Assessments with Students. In *Information Access Evaluation. Multilinguality, Multimodality, and Visual Analytics*, ed. Tiziana Catarci et al., 7488: 124-35. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer.

Schaer, Philipp. 2013a. Der Nutzen informetrischer Analysen und nicht-textueller Dokumentattribute für das Information Retrieval in digitalen Bibliotheken. Dissertation, Koblenz: Universität Koblenz-Landau <http://kola.opus.hbz-nrw.de/frontdoor.php?source_opus=896&la=de>.

Schaer, Philipp. 2013b. Information Retrieval und Informetrie: Zur Anwendung informetrischer Methoden in digitalen Bibliotheken. *Historical Social Research* 38 (3): 282-354.

Schamber, Linda, Michael B. Eisenberg, and Michael S. Nilan. 1990. A Reexamination of Relevance: Toward a Dynamic, Situational Definition. *Information Processing & Management* 26 (6): 755-76 <doi:10.1016/0306-4573(90)90050-C>.

Scharnhorst, Andrea, and Eugene Garfield. 2010. Tracing Scientific Influence. *Journal of Dynamics of Socio-Economic Systems* 2 (1): 1-33.

Siegfried, Doreen, and Elisabeth Flieger. 2011. World Wide Wissenschaft – Wie Professionell Forschende Im Internet Arbeiten. ZBW – Leibniz-Informationszentrum Wirtschaft. <http://www.zbw.eu/presse/pressemitteilungen/docs/world_wide_wissenschaft_zbw_studie.pdf>.

Strohman, Trevor, W. Bruce Croft, and David Jensen. 2007. Recommending Citations for Academic Papers. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 705-6. SIGIR '07. New York, NY, USA: ACM <doi:10.1145/1277741. 1277868>, <http://doi.acm.org/10.1145/1277741.1277868>.

Tague-Sutcliffe, Jean. 1992. An Introduction to Informetrics. *Information Processing & Management* 28 (1): 1-3 <doi:10.1016/0306-4573(92)90087-G>.

Tijssen, Robert J. W., and Anthony F. J. Van Raan. 1994. Mapping Changes in Science and Technology: Bibliometric Co-Occurrence Analysis of the R&D Literature. *Evaluation Review* 18 (1): 98-115.

White, Howard D. 2007. Combining Bibliometrics, Information Retrieval, and Relevance Theory, Part 1: First Examples of a Synthesis: Research Articles.

Journal of the American Society for Information Science and Technology 58 (February): 536-59 <doi:10.1002/asi.v58:4>.

White, Howard D. 1981. 'Bradfordizing' Search Output: How It Would Help Online Users. *Online Information Review* 5 (1) (December 31): 47-54 <doi:10. 1108/eb024050>.