

Running head: NO RESULTS FOUND

Your Search Returned: 0 Results

Improving Digital Library Search Tools

Paul Aumer-Ryan

The University of Texas at Austin

December 13, 2006

INF 397.3

Abstract

“No results found” is a misleading phrase because it masquerades as a definitive answer; in reality, the library being searched may contain content that matches a user’s query. This research seeks to understand the effect of null result sets on search behavior and on the perception of contents in digital libraries. Concepts utilized by social computing frameworks and current-generation eCommerce Web sites will be examined for their efficacy in digital library search interfaces. An experiment is outlined that attempts to flesh out a connection between null result sets and authoritativeness of digital library contents.

Introduction

The poster child for failure in information retrieval is a slew of search results that have nothing to do with the search query presented to the system by the user. If, for example, one were to search a database of household pets for the term “dog” and only cats with the name “dog” were returned, we could easily consider this to be a failure (unless, of course, the database actually did not contain any dogs). Similarly, if one were to perform a Web search on a hypothetical 5-year-old neighbor named Michael Jackson, and the only results returned dealt with the famous musician, there would be limited recourse to explain to the system the intended meaning (see, e.g., McRae-Spencer & Shadbolt, 2006). Much research has been undertaken in the information retrieval (IR) field to address this problem, ranging from the traditional concept of search term vectors, to popularity based on interrelations, to similar word pairings (e.g., synonyms, singular/plural forms of words, and verb tenses). To a large extent these efforts have been tremendously successful; modern search engines—whether they categorize the documents on a private Intranet, the items in a digital library, or the Web at large—are much better at

interpreting individual intention from search queries than their predecessors; most inappropriate and unrelated items do not make it into the search results for a given query, and there is a fair amount of trust in modern search engines that the results returned are a faithful representation of the documents in the collection that relate to the search query.

However, there is often an unintended consequence of our desire to perfectly match search results with the semantic intention behind a search query, especially when dealing with imperfect systems and the wide array of individual differences. In a sense, it is the opposite effect from the previously mentioned poster child of failure (i.e., too many results that do not relate to the search query): when an empty result set is returned to a user's search query. More succinctly, it can be summed up by an often unattractive visual: the "no results found" page returned by a search engine (see Appendix A for sample phrases taken from the "no results found" pages of current academic search tools).

By itself, a "no results found" page is not bad (although it may be poorly designed, or have any number of usability or navigation issues); however, the meaning it conveys may not be the appropriate one. "No results found" is a misleading phrase because it masquerades as a definitive answer; in reality, the collection being searched may actually contain content that matches a user's query.

Let us pause for a moment and examine the various meanings that "no results found" can have for a digital library user. There is the definition from computer science: "The explicit assemblage of characters you submitted does not occur anywhere in our index of items in our collection." There is also the anthropomorphic definition, where it is as if the computer is saying: "We don't understand what you just typed" or "we understand some of the things you typed, but not all of them." Then there are the potentially unintended and far too abrupt

meanings: “We have what you are looking for, but we call it something else”; “We don’t have what you are looking for”; “Go away.” How are we to know how the user interprets these pages? “No results found” seems very authoritative and final; in fact, it is a statement of truth, and it is coming from a computer (and computers tend not to be known for their ability to equivocate).

In examining this problem, it helps to narrow our focus a bit; in this case, we will be restricting our discussion to digital libraries and their search tools. Digital library search routines are presented with a much different and a much more varied subject of focus than their larger siblings, Web search engines. The most obvious difference between the two is collection size, but there are others.

Digital library search engines must be able to appropriately deal with smaller collections. Whereas one of the features that makes Web search engines so appealing is the vast size of the index (thus ensuring, in most cases, that results will be found), digital libraries must return relevant results from a much smaller data set. This makes it all the more important to deal with “no results found” pages, since they are likely to occur more frequently. Additionally, the Web, by its very nature, is interrelated, and this allows for some fascinating conclusions to be drawn about the connections between seemingly unrelated documents (e.g., Google’s PageRank relies in large part on a page’s “popularity” by counting how many other pages link to it). In most digital libraries this is impossible because of the independent nature of the items in the collection. Finally, certain digital libraries—those that are of a single media and document type, say collections of journal articles—have certain benefits compared to their multimedia brethren and the Web at large. Because of their homogeneity, metadata about the items in the collection

becomes more meaningful, and further inferences can be drawn that can influence the appropriateness of search results.

Defining the “Digital Library”

Before proceeding, I believe that the term “digital library” deserves some discussion; while it is safe to say that there is a common root that we can all agree on when invoking the term “digital library,” there is enough variation to warrant a specific definition for the purpose of this paper. Herein, a “library” is a collection of described objects that has been purposefully assembled from a larger, potentially unorganized heap. A “digital library” only adds the distinction that its items be digital or digitized (but be aware that this carries heavy connotations, since digital objects are inherently free from the restraints physical objects have, i.e., they can only exist in one place at a time and hence are not easily duplicated, lent, or categorized). Note that there is no concept of a complete digital library (for in that case, it contains everything and ceases to be a library, for it is no longer representative of anything outside the library). By this definition, the “Web” is not a digital library (it has no inherent organizational scheme, nor is it a subset of a larger collection), but a Web search engine’s (e.g., Google’s) index of Web pages is (it is an organized subset of all potential Web pages). However, because of the massive size of Web search engines and the interrelated nature of the objects they index, it is helpful for the purposes of this paper to consider them in a different category than digital libraries that focus on a smaller and more distinct subset of digital objects.

Models of Search Behavior

In the same way that there are many types of digital libraries, it is important to note that all users of search engines do not behave in the same fashion. Since the “no results found” page relies heavily on the assumption that all searchers are adept at and comfortable with query refinement (e.g., see Velez, Weiss, Sheldon, & Gifford, 1997), it is important to examine if searchers really do fit this mold. The most basic division of searcher behaviors is that between expert and novice, but these are not the only categories. Research by Heinström (2003) develops three categories of searchers: Deep Divers (those who follow one topic and are not easily distracted), Broad Scanners (those who look at a variety of resources, but not in too much depth), and Fast Surfers (those who avoid lengthy paragraphs and glean information through rapidity). In these categories, both Broad Scanners and Fast Surfers are unlikely to do much query refinement (i.e., to change their original search query in small increments to either narrow their search results or broaden them).

The other main difference in searcher behavior lies in the intention of the search process. When searchers are searching, they are either actively seeking information (i.e., information seeking; see Marchionini, 1997) or casually exploring (i.e., information encountering; see Erdelez, 2000; 1997). Those that are performing the directed act of information seeking will exhibit different tendencies than those who are merely encountering (e.g., those exploring will be more likely to chase down unlikely sources of information, with the potential for opening up new realms). Similarly, it is expected that those exploring will be less likely to “try again” after receiving no results from a search query, since they are not particularly invested in that particular library and that particular query.

The overriding implication here is that not only do searcher behaviors impact the way searches are performed, but they also affect the way searchers deal with empty result sets. The conclusion must be that a single purpose, catch-all page that states that no results were found is improper for many styles of search behavior.

Suggestions in Search Results

A common cause of empty result sets is when the user uses a term that the digital library does not recognize—in this case, “no results found” does not necessarily mean that the library does not contain what is being sought, but rather that it does not understand the query. Examples of this type of misunderstanding include misspellings or use of jargon, abbreviations, or acronyms. Dalianis (2002) discusses the implementation of spelling support (while most Web search engines now include spelling support, most digital libraries are behind the curve and do not). The current method of implementing dictionary support is to place a “did you mean to type X” prominently on the (empty) results page, allowing the user to simply click the correct spelling to get the appropriate results. An alternative is to detect common misspellings and automatically search using the correct spelling, especially in the cases where the original, supposedly misspelled query returned no results (e.g., a user searching for “kats” would be able to see results for “cats”).

Similarly, the implementation of a thesaurus (whether field-specific or general purpose) can help alleviate some inappropriate “no results found” pages (see, e.g., Medelyan & Witten, 2006; Song, Song, Allen, & Obradovic, 2006; Weaver, Strickland, & Crane, 2006). Thesaurus-based searching is one way to see relationships between supposedly distinct objects in a digital library, and can help users make connections that they may not have known about. This can be

implemented in much the same way as the spelling correction, with a list of synonyms displayed prominently on the search results page, or by automatically including synonyms in the search results (e.g., a user searching for “cats” could also see results with “kitten” or “feline”).

Another common problem with search results, especially in digital libraries that focus on specialized fields with a hefty amount of jargon, is the recognition of acronyms and abbreviations. For example, most academic journals and conferences have both a full name and a shorter acronym (e.g., JCDL for the Joint Conference on Digital Libraries). When searching a digital library of articles published in journals and conference proceedings, will the same results be returned for both the acronym and the full name? Ideally this should be the case, but it often is not. Many researchers have examined the potential for automatically recognizing acronyms (which would allow us to avoid giving two labels to metadata fields, or to maintain a dictionary that relates acronyms and abbreviations to their full text equivalents), with promising—yet still not widely implemented—results (see, e.g., Dominich, Goth, & Skrop, 2003; Larkey, Ovilvie, Price, & Tamilio, 2000; Yeates, 1999; Taghva & Gilbreth, 1999; Federiuk, 1999).

Finally, a search engine can perform automatic permutations on a search phrase that did not return any results (in other words, the searcher can be offered a list of potential results if certain search terms were removed from the query). At the time of this writing I only know of one “library” that uses this technique, with a fair amount of success: eBay. It is easily applicable to digital library searches, and can help users recognize when the “no results found” page is actually saying, “I understand most of what you typed, but not all of it.”

Social Aspects in Search Results

The final category of research that is applicable to digital library searching that we will examine comes from social computing. With the recent explosion of Web content that relies on social networks (e.g., Myspace.com, Facebook.com, Flickr.com, Del.icio.us, and Craigslist.org) and their success in showing how social relationships can be a strong substratum for information relationships, it seems inevitable that digital libraries will embrace social computing styles.

With that in mind, a wave of recent research has indeed focused on social computing in the realm of digital libraries: collaborative filtering (Huang, Li, & Chen, 2005); sharing information encounters (Marshall & Bly, 2004); and peer verification (Schmalbeck, Stuart-Moore, & Evans, 2006), to name a few. It is a short leap to see the potential that social computing methods can have on empty result pages—if other people searched for the same thing or something similar, and had the same difficulties in finding the documents they were looking for, why not learn from that experience and apply it for future users?

We may also ask: why not allow the community to add their own metadata to objects in the digital library? What is lost in consistency and hierarchy can be made up in robustness and variety (not to mention that the original, consistent, and hierarchical metadata can be left untouched). Similarly, there is the potential for involving the community in the ranking of the objects in digital libraries (e.g., this was helpful, that was not).

The only convincing negative answer to these questions revolves around traditional notions of privacy and confidentiality in the library sphere, especially with acts of personalization or peer recommendation. Both of these require some amount of personal information storage by the digital library, and personal information sharing between patrons of the library, and this can be troublesome. However, it is important to note that not all social

actions demand privacy in the library sphere—after all, physical libraries require us to attend them in person, which is in some sense an invasion of privacy. It seems that the safest methods include explicit contributions by users to the library—tagging, ranking, and sharing—rather than implicit data mining. A traditional library is in many ways a community center, and we as researchers should not be afraid of making their digital counterparts behave in the same way.

Proposed Study

Previous studies that have examined empty result sets (see DeFelice, Kastens, Rinaldo, & Weatherley, 2006; Kan & Poo, 2005; Zhuang, Wagle, & Giles, 2005) have been more focused on the collections themselves rather than the users encountering the empty result sets. In order to probe the effects that null result sets have on digital library users, I propose a study that addresses two broad questions. First, what are the affective implications of encountering a null result set? Measures of affect are an appropriate indicator of many aspects of the user experience, and bear a relationship to many expectations for a digital library: namely, retention of users, ease-of-use, minimizing extraneous cognitive load, satisfaction, and sense of accomplishment (Dillon, 2001). By understanding the affective impact that digital library search tools have on end users, both individually and as groups, we can shed some light on user expectations of digital library interfaces as well as improve their efficacy for everyday use.

Second, what impact does the digital library interface have on the interpretation of its contents? An examination of the effect of a digital library's interface on the perception of its contents may have strong repercussions for digital library designers and administrators. The general assumption is that the content and the interface to the content are separate entities. For example, a research paper can exist in two different digital libraries; one of these libraries can

have a highly usable interface, while the other is a disaster of design. However, users who end up finding the research paper in question will treat it the same way, regardless of which library they found it in. In other words, the content in the research paper (or book, video, or any other item in a collection) is self-contained, and not in any way affected by the way in which it arrived in the hands of a reader (in this case, the digital library and its interface). This study seeks to test this notion, with the belief that this is not the case; I believe that a poorly developed user interface, especially in the case of a digital library (where there is a large cache of information completely mediated by a forced interface that separates the user from his or her goal), will have a strong effect on the way a user sees the content within the collection.

The purpose of this study is twofold. First, by further understanding the affective response to elements of digital library interfaces (i.e., search tools), we can expect future digital libraries to be more in line with the emotional responses of their users. Second, by demonstrating a link (or lack thereof) between the digital library interface and the contents within the library, future efforts into digital library design and development will be more grounded by their actual effect on users; in other words, if highly usable digital library interfaces merely offer aesthetic beauty to users, but no real productivity benefits, their importance can be relegated appropriately. On the other hand, if the interfaces do indeed have a strong effect on productivity, and (even more importantly) on the interpretation and respect of the actual items in the collection, then it is further proof of the vast importance of appropriate and beneficial design for the end users of digital libraries.

Method

In this study, a mock digital library will be created as a standardized test bed for research. Efforts will be made to make the digital library both simplistic and clean (to minimize extraneous variables that may affect participants, and to emphasize the benefits of a clean interface) and believable (so participants will be in the mindset that they are using a real digital library).

Participants will interact with the mock digital library via a simple search tool; they will be under the impression that they are evaluating a new digital library and its overall design and responsiveness. They will be given a topic to search for (this topic will be chosen to be academic in nature, but not a topic that is intimately familiar to potential participants) and several comprehension questions to answer regarding the topic. The mock digital library will contain a small (30 or less) set of documents pertaining to the topic at hand and another small set of unrelated documents (30 or less); their presentation on the search results screen will be the main manipulation in this study.

Participants will be divided into three groups: the first (control) group will get appropriate results returned to them on the search results pages, regardless of the terms they use in their search query; the second (experimental #1) group will always get an empty result set back in response to their first search query; further refinement of that search query will deliver the appropriate results; finally, the third (experimental #2) group will encounter inappropriate and unrelated results in response to their first search query; further refinement of that query will then deliver the appropriate results. We expect to have at least 50 participants in each group, drawn from a university student population of undergraduates with a small amount of research experience.

The following data will be collected before the experiment is performed: a set of demographic questions (age, gender, college major); a brief affective mood inventory, such as the SAM (Dormann, 2003; Morris, 1995a; 1995b; Lang, 1985); and a survey of the participant's familiarity with technology and research (experience with computers, experience with research, knowledge of the term "digital library" and a usable definition, previous use of libraries and library materials, both digital and physical).

During the experiment, the following data will be collected: overall time-on-task; search queries typed into the search field; number of mouse clicks; and number of back button presses.

After the experiment, the following data will be collected: a set of questions confirming comprehension of the topic searched for (these questions will be drawn from the appropriate items in the collection to ensure that they were examined); an overall impression of the digital library and the search tool, similar to Toms, Dufour, & Hesemeier (2004); a repeat of the affective measure used to determine mood (e.g., the SAM); and finally, an overall rating of the perceived authoritativeness of the contents in the digital library.

This last question that deals with authoritativeness is the central question of the study; I hope to find a significant difference between the groups (using the analysis of variance statistical test, F). Measures of authority will be similar to Chesney's (2006) examination of Wikipedia's credibility by novices and experts.

Expected Findings

I expect the following hypotheses to be supported:

H1: Participants who encounter a null result set after their first search query will take more time to complete their task;

H2: Participants who encounter a null result set after their first search query will rank the results as less authoritative;

H3: Participants who encounter a null result set after their first search query will have a lower opinion of the search tool;

H4: Participants who encounter a null result set after their first search query will exhibit more negative affect (e.g., frustration, anger, distress);

H5: Novice users are more susceptible to H1 – H4 than expert users (Chesney, 2006).

Participants who encounter inappropriate and unrelated results after their first search query are expected to react similarly.

This study is intended to elucidate the dangers of “no results found” responses (especially when they are misleading) by showing their actual effect on potential digital library users. Further, if participants do indeed see the results following a “no results found” page as less authoritative, then this implies that the contents of a digital library are being evaluated not on their own merit, but by the interface’s effect on them. It is also important to note that if users have a lower opinion of a digital library (because of a “no results found” page, or any other reason), they are likely to turn elsewhere for their information needs (in other words, the assumption of the captive user is not necessarily true).

Conclusion

Since the focus here has been so strongly on empty search result pages, a topic which gets little attention both in the design of digital libraries and in the research literature, I thought I would close with a brief discussion of their ephemeral existence. “No results found” pages often

disappear during the design and testing phase for a digital library; evidence is readily available by examining any number of null result sets in digital library search tools. Some are grammatically incorrect; others look thrown together; all exhibit terseness unbecoming of a helpful library assistant. Why is this? I would speculate that it is because “no results found” pages are not seen as actual pages or destinations in the digital library; they are merely fleeting error messages that have little impact other than prodding the user to “try again.” I believe there is also the assumption that they occur so infrequently (especially in formulaic testing phases) that it would be a waste of time to make them better or more infrequent. After all, in a world where information overload is so apparent (Levy, 2005) that it has become cliché, why focus on the opposite? But because of their deceptive nature (i.e., no results “found” does not necessarily mean no results “exist”), I believe they deserve our attention until they no longer mislead digital library users.

Appendix A: Sample Phrases from No Results Found Pages

(Note the difference in punctuation, capitalization, and phrase vs. sentence structure.)

- “No results found”
- “No results returned for your criteria.”
- “No results were returned.”
- “Nothing Found”
- “Sorry, no documents were found matching search terms.”
- “There are 0 results”
- “No Results Found.”
- “0 articles with title/keywords/abstract containing *”
- “Your search matched 0 documents.”
- “There are no products that match your search”
- “No videos were found to match your query.”
- “No results were found.”
- “Sorry, your request returned no records.”
- “Results: Not Found”
- “No documents were found for your search.”
- “No Results matching your search term(s) were found.”

References

- Anick, P. G., & Tipirneni, S. (1999). The paraphrase search assistant: terminological feedback for iterative information seeking. *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, 153-159.
- Bertini, E., Catarci, T., Di Bello, L., & Kimani, S. (2005). Visualization in digital libraries. In Hemmje, Niederee, & Risse (Eds.), *Integrated Publication and Information Systems to Information and Knowledge Environments* (pp. 183-196). Berlin, Germany: Springer.
- Bertuca, D. J. (2001). Letting go of the mouse: Using alternative computer input devices to improve productivity and reduce injury. *OCLC Systems & Services*, 17(2), 79-83.
- Chesney, T. (2006). An empirical examination of Wikipedia's credibility. *First Monday*, 11(11).
- Coors, V., & Jung, V. (1998). Using VRML as an interface to the 3D data warehouse. *Proceedings of the third symposium on Virtual reality modeling language*.
- Cruz, I. F., & Lucas, W. T. (1997). A visual approach to multimedia querying and presentation. *Proceedings of the fifth ACM international conference on Multimedia*, 109-120.
- Dalianis, H. (2002). Evaluating a spelling support in a search engine. *Proceedings of NLDB-2002, 7th International Workshop on the Applications of Natural Language to Information Systems, Junio*.
- DeFelice, B., Kastens, K. A., Rinaldo, C., & Weatherley, J. (2006). *Insights into collections gaps through examination of null result searches in DLESE*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries. from <http://doi.acm.org.content.lib.utexas.edu:2048/10.1145/1141753.1141823>.

- Dillon, A. (2001). Beyond usability: process, outcome, and affect in HCI. *Canadian Journal of Information Science*, 26(4), 57–69.
- Dominich, S., Góth, J., & Skrop, A. (2003). A Study of the Usefulness of Institutions' Acronyms as Web Queries. In *Sebastiani, F. (Ed.), Advances in Information Retrieval, Lecture Notes in Computer Science, LNCS (Vol. 2633, pp. 580-587)*.
- Dormann, C. (2003). *Affective experiences in the Home: measuring emotion*.
- Erdelez, S. (1997). Information encountering: a conceptual framework for accidental information discovery. *Proceedings of an international conference on Information seeking in context*, 412-421.
- Erdelez, S. (2000). Towards Understanding Information Encountering on the Web. *Proceedings of the ASIS Annual Meeting*, 37, 363-371.
- Federiuk, C. S. (1999). The Effect of Abbreviations on MEDLINE Searching. *Academic Emergency Medicine*, 6(4), 292.
- Heinström, J. (2003). *Fast surfers, broad scanners and deep divers: personality and information-seeking behaviour*: Åbo Akademi University Press.
- Hsieh-Yee, I., Tang, R., & Zhang, S. (2006). *User perceptions of a federated search system*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Huang, Z., Li, X., & Chen, H. (2005). *Link prediction approach to collaborative filtering*. Paper presented at the Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries.

- Jones, S. (1998). Graphical query specification and dynamic result previews for a digital library. *Proceedings of the 11th annual ACM symposium on User interface software and technology*, 143-151.
- Kan, M.-Y., & Poo, D. C. C. (2005). *Detecting and supporting known item queries in online public access catalogs*. Paper presented at the Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries.
- Kerne, A., Koh, E., Dworaczyk, B., Mistrot, J. M., Choi, H., Smith, S. M., et al. (2006). *combinFormation: a mixed-initiative system for representing collections as compositions of image and text surrogates*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Kim, J. (2005). *Task difficulty in information searching behavior: expected difficulty and experienced difficulty*. Paper presented at the Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries.
- Kules, B., Kustanowitz, J., & Shneiderman, B. (2006). *Categorizing web search results into meaningful and stable categories using fast-feature techniques*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Lang, P. J. (1985). *The cognitive psychophysiology of emotion: Fear and anxiety*: Lawrence Erlbaum Associates, Inc.
- Larkey, L. S., Ogilvie, P., Price, M. A., & Tamilio, B. (2000). Acrophile: an automated acronym extractor and server. *Proceedings of the fifth ACM conference on Digital libraries*, 205-214.

- Levy, D. M. (2005). *To grow in wisdom: Vannevar Bush, information overload, and the life of leisure*. Paper presented at the Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries.
- Marchionini, G. (1997). *Information Seeking in Electronic Environments*: Cambridge University Press.
- Marshall, C. C., & Bly, S. (2004). *Sharing encountered information: digital libraries get a social life*. Paper presented at the Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries.
- McRae-Spencer, D. M., & Shadbolt, N. R. (2006). *Also by the same author: AKTiveAuthor, a citation graph approach to name disambiguation*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Medelyan, O., & Witten, I. H. (2006). *Thesaurus based automatic keyphrase indexing*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Morris, J. D. (1995a). Observations: SAM: The self-assessment manikin. *Journal of Advertising Research, 35*, 63-68.
- Morris, J. D. (1995b). SAM: The Self-Assessment Manikin. an Efficient Cross-Cultural Measurement of Emotional Response. *Journal of Advertising Research, 35*(6).
- Na, J.-C., Khoo, C. S. G., Chan, S., & Hamzah, N. B. (2005). *Sentiment-based search in digital libraries*. Paper presented at the Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries.
- Nakazato, M. (2003). *Toward Flexible User Interaction in Content-Based Multimedia Data Retrieval*.

- Nowell, L. T., France, R. K., Hix, D., Heath, L. S., & Fox, E. A. (1996). Visualizing search results: some alternatives to query-document similarity. *Proceedings of the 19th annual international ACM SIGIR conference on Research and development in information retrieval*, 67-75.
- Perugini, S., McDevitt, K., Richardson, R., Manuel A. Perez-Quinones, Shen, R., Ramakrishnan, N., et al. (2004). *Enhancing usability in CITIDEL: multimodal, multilingual, and interactive visualization interfaces*. Paper presented at the Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries.
- Rao, R., Pedersen, J. O., Hearst, M. A., Mackinlay, J. D., Card, S. K., Masinter, L., et al. (1995). Rich interaction in the digital library. *Communications of the ACM*, 38(4), 29-39.
- Schmalbeck, L., Stuart-Moore, J., & Evans, M. (2006). *Adapting peer verification, validation and accreditation processes for digital libraries*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Shneiderman, B., Byrd, D., & Croft, W. B. (1998). Sorting out searching: a user-interface framework for text searches. *Communications of the ACM*, 41(4), 95-98.
- Song, M., Song, I. Y., Allen, R. B., & Obradovic, Z. (2006). *Keyphrase extraction-based query expansion in digital libraries*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Spoerri, A. (2004). MetaCrystal: A Visual Interface for Meta Searching. *Proceedings of ACM CHI*.
- Taghva, K., & Gilbreth, J. (1999). Recognizing acronyms and their definitions. *International Journal on Document Analysis and Recognition*, 1(4), 191-198.

- Toms, E. G., Dufour, C., & Hesemeier, S. (2004). *Measuring the user's experience with digital libraries*. Paper presented at the Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries.
- Vélez, B., Weiss, R., Sheldon, M. A., & Gifford, D. K. (1997). Fast and effective query refinement. *ACM SIGIR Forum*, 31, 6-15.
- Weaver, G., Strickland, B., & Crane, G. (2006). *Quantifying the accuracy of relational statements in Wikipedia: a methodology*. Paper presented at the Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries.
- Yeates, S. (1999). Automatic extraction of acronyms from text. *Proceedings of the Third New Zealand Computer Science Research Students' Conference*, 117-124.
- Zhuang, Z., Wagle, R., & Giles, C. L. (2005). *What's there and what's not?: Focused crawling for missing documents in digital libraries*. Paper presented at the Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries.