

Measuring Our Relevancy: Comparing Results in a Web-Scale Discovery Tool, Google & Google Scholar

Elizabeth Namei and Christal A. Young

Introduction

The University of Southern California (USC) Libraries' adopted Summon as its discovery solution in 2010. Since installing it as the default search option on the Libraries' homepage usage numbers have steadily increased, reaching over 3 million searches in 2014.¹ Despite this heavy use, there are still many students, faculty and librarians who express various levels of frustration or ambivalence about it. A common complaint is that known item searches often return "unexpected results." Librarians at other institutions have reported similar comments, such as finding "unpredictable," "erratic" results, or even "irrelevant junk" when searching their institution's unified search system.² Although these Google-like library search boxes have been shown to appeal to novice users as an intuitive starting place,³ Lundrigan found that those who were trained to use specialized databases were less likely to give high ratings to discovery services.⁴ Because of their prominent position on library websites, discovery layers need to meet the needs of both novice and power users.⁵ Based on these mostly anecdotal critiques of discovery layers we decided to conduct a quantitative study to gather evidence about how successful Summon was in returning relevant results. Our aim was to better understand when and why USC's discovery service returned "unexpected"

results and whether this is related to user search behavior, Summon's relevance algorithm, or a combination of the two.

The Challenge of Relevancy

One of the reasons why providing relevant results is so vital for any search tool is because most users tend to only look at the first page of results and many only click on the first item in the results list.⁶ Jansen and Spink speculate that the reason behind this behavior may be due to "the increasing ability of Web search engines to retrieve and rank Web documents more effectively."⁷ This theory implies that relevance algorithms shape user behavior, yet two other studies suggest otherwise. After presenting students with a list of ten Google results in reverse order Pan et al. found that students "were heavily influenced by the position of items in the results list"⁸ and Keane et al. found that users are often "misled by the presented order of the items."⁹ Keane et al. speculated that this reliance on top results might be explained by what has been labeled "satisficing": the tendency to choose the easiest and most convenient path, one that leads to "good enough" information rather than the "best" information.¹⁰ Satisficing can also be understood as a coping mechanism for dealing with information abundance and overload.¹¹ Since it would

Elizabeth Namei is Reference & Instruction Librarian, University of Southern California Libraries, e-mail: namei@usc.edu; Christal A. Young is Reference & Instruction Librarian, University of Southern California Libraries, e-mail: youngc@usc.edu

be a monumental task to look through millions of results, users tend to accept what is on the first page, revise their search, or abandon the search tool altogether if useful results are not found near the top of the list.¹² Once users find a search tool that consistently and reliably delivers relevant results at the top of the list they are reluctant to leave this comfort zone.¹³ If a user feels overwhelmed by results, that has been recognized as a sure sign that something is wrong with the relevancy algorithm, since the number of results returned does not matter when the best results are at the top of the list.¹⁴ In other words, the fact that Google is never criticized for returning “too many” results is a testament to its excellent relevance algorithm.¹⁵

Successful search tools, like Google, continuously adapt and adjust their ranking algorithms as they encounter user behaviors that might otherwise impede the provision of relevant results. For instance, numerous studies have shown that users tend to construct simple or ambiguous queries when searching for information online.¹⁶ This simplistic searching, in turn, places enormous pressure on algorithms to produce high quality results,¹⁷ since users provide limited “signals” to be of much help.¹⁸ Asher argues that making search tools easier to use has “counterintuitively” led to lower quality, “unreflective” search behavior.¹⁹ Regardless of these debates over cause and effect, user search behavior and search tools can be best understood as interdependent, working in perpetual concert, reinforcing, influencing, and adapting to the other.

A search tool’s ability to adapt to both technological changes and evolving user expectations in order to consistently provide relevant results can inspire intense loyalty. The inverse is true as well: a search tool that fails to adapt, will be quickly abandoned.²⁰ Libraries are all too familiar with the consequences of failing to adapt our search tools.²¹ This is where web-scale discovery services come in—they have been marketed as a way to “return researchers to the library,” and one way they are trying to accomplish this is by providing more relevant results.²²

Relevancy and Web-Scale Discovery

ProQuest acknowledges that they looked “to open web search technologies that set the bar for user expectations” by providing “highly relevant results” when creating Summon.²³ Yet, soon after library discovery services were developed, concerns arose over the relevancy of results,²⁴ and in the last year known item searching has come under even more scrutiny.²⁵ For example, Roger Schonfeld, in a 2014 Ithaka S+R Report, presented “attitudinal data from library directors” regarding discovery systems. “Respondents indicated that indexed discovery services had broadly improved exploratory searching, but that this was less the case for known-item searching.”²⁶ Marshall Breeding also recently stated that although discovery services have made some improvements in the “handling of known item searching,” they still struggle, especially with items that have one-word and/or common word titles, such as *Nature* or *Time*.²⁷ Brent Cook, the Director of Product Management for Discovery Services at ProQuest, explained, in a recent e-mail exchange, that they are well aware of the problems surrounding known item searching in Summon and are working on solving this “challenging” issue.²⁸ As we wait for vendors to improve the relevancy of their results, this study is an attempt to better understand the extent and scope of the problem, while also getting a clearer picture of how users approach searching for known items in discovery systems.

Related Studies

There have been numerous studies published about discovery services since they were first introduced in 2007. Of particular interest for this study are those that used transaction log analysis to better understand how users interacted with library search tools, those that manually classified search queries or results, and those that rated the relevancy of results.

Transaction Log Analysis

Transaction logs capture data about the actions executed in an online system. Transaction log analysis (TLA) involves extracting this data, and looking

for patterns to better understand how searchers approach and interact with a system, with the overall goal of improving the system's design and functionality.²⁹ The advantages of TLA are that it is inexpensive and unobtrusive and can provide large quantities of "real world data."³⁰ The weaknesses of TLA are that it can be very time consuming and it does not provide contextual information about users: demographics, satisfaction-level, details about the information need or overall search experience.³¹ TLA has been used to better understand user behavior when searching a wide range of search tools, in both libraries and on the open web.³²

McKay and Buchanan compared the initial queries users entered when searching four different search tools from the library's homepage: the library catalog, EBSCOHost, Gale, and Google Scholar. Some of their most interesting findings had to do with the cause of failed searches: typographical errors and the practice of copying and pasting formatted and unedited citations.³³ In 2013 McKay and Buchanan followed up on their previous study, this time examining search queries conducted before and after the implementation of a discovery solution.³⁴ Their intention was to "mine the specific impact that the introduction of web-scale search had on search behavior."³⁵ These two studies depict users' impulse towards entering longer queries when searching for known items. The 2011 study found users entered shorter queries in the library catalog, but longer ones when searching databases and Google Scholar. This suggests queries are formulated differently depending on an understanding of the search tool. In 2013, McKay and Buchanan noticed that users started the year entering longer, more specific known item queries in the discovery service, but by the end of the year shorter queries were more common. They concluded that users were learning and adapting their behavior based on the failures and successes of different search types.³⁶ Users' willingness to adapt to the search tool rather than abandoning it for not behaving as expected is reassuring. Yet, the fact that library discovery systems struggle to process longer, more complex and precise queries is a concern.

Query and Relevance Classification & Judgments

In addition to extracting queries from transaction logs, some researchers also manually classified the searches executed in a library search tool in order to discern user behavior patterns. Meadow and Meadow classified and rated queries from Summon's logs with the aim of understanding "how the single search box model is being utilized."³⁷ Their study investigated search type and quality but did not examine results returned. Chapman et al. categorized website search queries to investigate what users were searching for and "whether the structure of search results pages reflect[ed] those needs."³⁸ In particular they wanted to assess whether their default search option "performed...well for searches in the 'long tail' (i.e., the vast number of searches conducted only a handful of times)."³⁹

Zhang conducted a usability study with graduate students to compare their "subjective assessment of search results" from Primo and Google Scholar. They observed that, although the two search tools returned comparably relevant results, students' complaints about "usability issues" of Primo influenced their perception regarding the relevance of results.⁴⁰ This study illustrates how relevancy is both subjective and contextual. In another study, Singley compared how well different discovery services, and Google, performed in finding known items. Problematic queries were chosen, such as: one-word searches, citation searches copied from a bibliography, ISBN searches, and titles with stop words. Although the results of her study are limited due to the sample size, they do suggest that discovery services have a lot of work to do in terms of providing relevant results for known items and in ensuring that local print content is discoverable.⁴¹

Methods

This study builds upon existing research and uses the methods discussed in the review of literature (TLA, classification of queries, and rating the relevance of results) in order to better understand (1) user search behavior within discovery services and (2) why cer-

tain searches are successful and others are not. A four-step process was developed in order to carry out the study. First, we compiled a random sample of known item queries from USC's Summon transactions logs, and then we replicated those queries in Summon, Google and Google Scholar. We then rated the results returned by each search tool according to a set of relevancy categories. Lastly, in order to address the possible influence of user search behavior on the outcome of a search, we recorded information about the queries, including: (1) the type of search executed, (2) the number of words in each query, (3) the formats being searched for, and (4) the presence of user input errors. In this paper we will be elaborating on how the first category—type of search executed—contributes to the success or failure of known item searches. Types of searches identified were: (1) advanced/field searches, (2) title searches (either copy and pasted or manually entered), (3) searches with two metadata elements, (4) formatted citations copied and pasted (with three or more metadata elements), (5) numeric queries, and (6) periodical/database title searches.

Definitions: Known Items and Relevancy

Following the methodology of our earlier study,⁴² we defined known item searches as being specific enough for us to recognize a definitive match. Searches that resulted in numerous exact matches, but were very general, were not categorized as known items (for example, "Catherine the Great" or "Ancient Civilization"). We also looked for visual clues to help us confirm if a query was for a known item, such as: capitalization of major words, punctuation (colons), phrases like, "Introduction to" or "Third Edition," or formatting, including author initials or dates in parentheses.⁴³

A "systems-oriented relevance" decision-making process was utilized which allowed us to rate results without knowledge of the context of the original query. Ratings were based solely on whether the query found a match. A three-level classification scheme was used to rate the degree of relevancy:⁴⁴

- *Relevant*: Assigned if the first or second result was a match for the known item.

- *Partially Relevant*: Assigned if the known item found a match in the 3rd-10th results.
- *Not relevant*: Assigned if the known item was not listed in the first 10 results or no results were returned.

When determining the relevancy of Google and Google Scholar's results, a full-text match was not required. A result was considered relevant if the full text was freely available or for a fee. In addition, a catalog record alerting the user of the known item's existence and availability was also considered a match. If a result only included a citation or description of a topic related to the known item but did not provide information about how to access/obtain the full text, it was not considered relevant (Wikipedia and course syllabi for instance). The following sources were all considered relevant matches when searching Google and Google Scholar for known items: bookstores (Amazon), university/faculty websites including institutional repositories, journal or publisher websites, library catalogs (WorldCat), free or subscription article databases (JSTOR, Pubmed), digital archives (Internet Archive or Google Books), academic social networking sites (Researchgate, Mendeley), media sites (Youtube or imdb.com), and commercial or for-profit websites.

Similar to McKay and Buchanan,⁴⁵ we examined the initial searches executed using the default settings of each search tool unless the query included advanced search commands. In those cases we re-executed the search using the advanced search page if there was one available with applicable fields (i.e., Google Scholar has an author and date field but no title field; Google's advanced search form does not include the fields used in any of the Summon queries). If a field search option did not exist we stripped the command language out of the query before entering it in the basic search box. We decided not to examine search refinements as we concurred with Lau and Goh⁴⁶ and Niu, Zhang and Chen,⁴⁷ that the first search attempt is of the utmost importance. We also adhered to the logic of only reviewing the first ten results,⁴⁸ as the majority of users select items found on the first page of results.⁴⁹

Inter-rater Reliability

Recognizing that rating relevancy can be subjective and contextual, the authors separately executed and rated the same batch of twenty queries, compared ratings, and then made adjustments to the rubric in order to ensure a high level of agreement. We measured our inter-rater reliability using Cohen's kappa statistic, which was "designed to estimate the degree of consensus between two judges" after correcting for agreement by chance.⁵⁰ We achieved a nearly perfect agreement with a 0.81 Cohen kappa. This level of agreement was likely due to the fact that known item searches are straightforward (you either find a match or you do not).

Sample

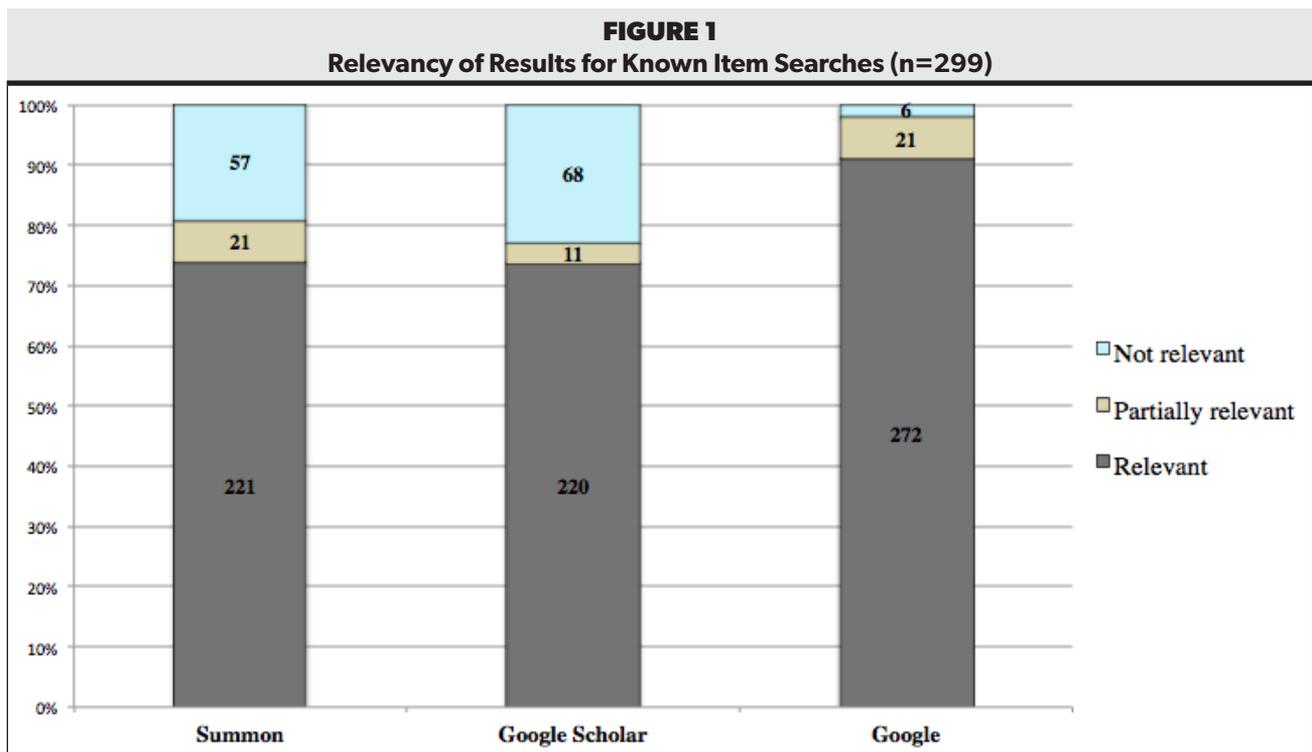
We extracted search queries entered into USC's Summon instance over the course of the Fall 2014 semester. There were a total of 1,194,263 queries conducted during this four-month period, and of those, 433,863 were unique. In order to achieve a 95 percent confidence level and a 5 percent margin of error, we needed to analyze a random sample of 384 queries.⁵¹ Since our study was focused on only

known item searches, we had to conduct a multi-step process to get our final sample. We pulled out a random sample of 1,150 queries, and then manually reviewed them until we found 384 known item queries.⁵² We eliminated 22 queries from our sample that contained unrecognizable characters that would likely negatively influence the outcome of a search. In order to make a comparison about the success or failure of the searches in Summon, we also removed the queries for known items USC did not own (63 items). This left us with a final sample of 299 queries.

Findings

Relevancy of Known Item Queries

Out of the 299 queries, Summon returned relevant results 74 percent of the time and results that were not relevant 19 percent of the time (figure 1). Google Scholar had the same percentage of relevant results as Summon (74 percent) but had the most failed searches of the three search tools (23 percent). In stark contrast, Google delivered relevant results 91 percent of the time and only failed to return relevant results 2 percent of the time.



Relevancy of Scholarly Known Item Queries

This initial comparison included 21 queries that were for “non-scholarly” formats that Google Scholar would, by definition, be less likely to index (these included: videos, journal titles, databases, and sound recordings). With these non-scholarly formats eliminated, Summon returned relevant results 76 percent of the time and results that were not relevant 17 percent of the time (figure 2). With non-scholarly items removed, Google Scholar’s performance also increased slightly (although not a statistically significant increase) to 79 percent relevant results and 17 percent failed searches. Google’s performance remained the same (91 percent of results were relevant). From these findings, it is not surprising that Google is so often the first choice, regardless of the age group, experience level, or information need of the user.⁵³ Finding that a discovery service provides comparably relevant results to Google Scholar confirms Zhang’s findings.⁵⁴

Impact of Search Type on Relevancy

As a first step towards understanding why Summon was successful (or unsuccessful) in returning relevant results when searching for known items, we examined how the type of search executed impacted the relevancy of results. For the sake of comparison across all

search tools we only examined scholarly known items (278), since including all queries (299) would not have made a discernable difference. We categorized all queries by search type (figure 3), and found that the two most commonly utilized were title searches (66 percent) and those with two metadata elements (author/title, title/date, title/URL, and article title/journal title) (25 percent).

Of all the search types, the two that were most successful in returning relevant results in Summon were title searches and advanced/field searches: 90 percent and 89 percent respectively (figure 4). Both Google and Google Scholar also performed well on title searches (95 percent and 86 percent). Despite the high relevancy achieved by advanced/field searches, they only comprised 3 percent of all known item searches (figure 3). It was somewhat surprising to find that Summon outperformed Google on advanced/field searches, with Google only returning relevant results 67 percent (6/9) of the time. This is likely because, as mentioned earlier, Google does not offer precision field searching (comparable to library databases), and thus searches were input via the basic search option. Google’s performance on these searches was also worse than Google Scholar’s, which was 100 percent successful in delivering relevant results

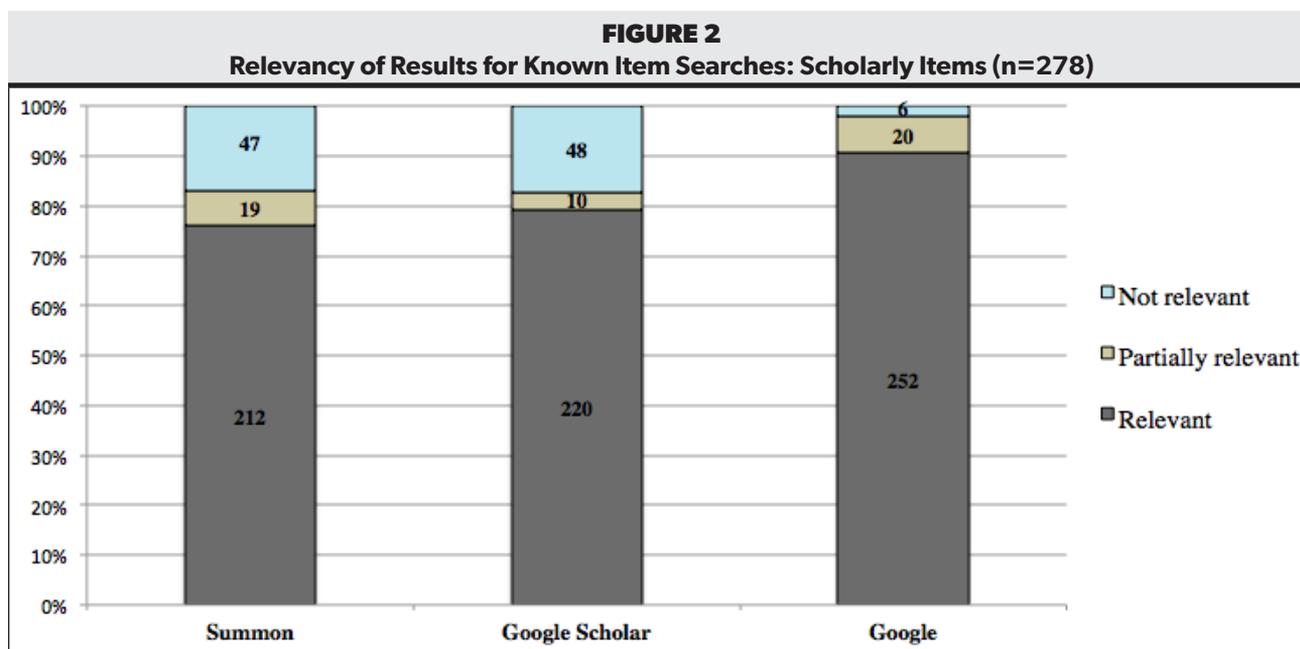
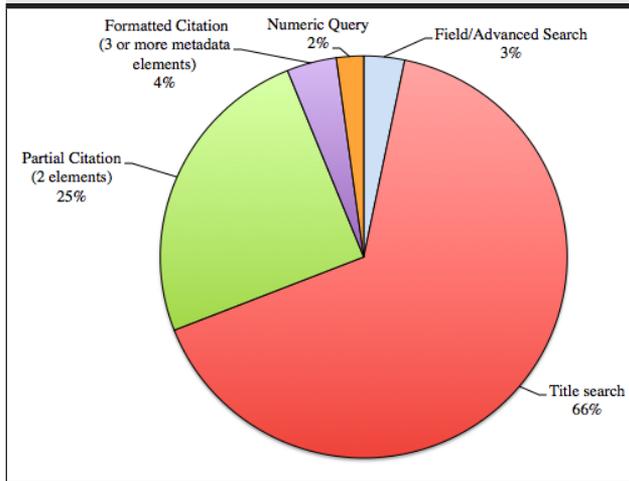


FIGURE 3
Types of Searches Executed: Scholarly Known Items (n=278)



(9/9). These findings highlight the power and usefulness of field searching, something librarians will not be surprised about.

The third most successful search type for Summon were numeric queries, which found relevant results 67 percent (4/6) of the time. This time Summon outperformed Google Scholar, which was only successful for 17 percent (1/6) of these queries. Google came in first (again) for queries with these search types: 83 percent (5/6).

Summon was least successful on searches that included three or more metadata elements, copied and pasted, and often including abbreviations and other formatting (figure 5). 73 percent (8/11) of these searches failed to find relevant results in Summon.

Somewhat surprising was that 32 percent (22/69) of the queries with two metadata elements failed to find relevant results. Of all the queries with two citation elements (69), 91 percent (63/69) were a variation of author/title searches. The author/title combination turns out to be the main cause of failed searches when two metadata elements are included in a search: 96 percent (21/22). On closer examination, queries that include the author’s first and last names resulted in a failed search 42 percent (15/36) of the time. In contrast, queries that only included the author’s last name/title, only had a 22 percent failure rate (6/27). When combined, any query with two or more metadata elements made up 64 percent (30/47) of all failed searches in Summon even though these two search types comprise only 29 percent (80/278) of all search queries (figure 4). These findings corroborate McKay and Buchanan’s findings that “the greater the number of metadata types, the more likely the search was to fail.”⁵⁵

FIGURE 4
Relevant Results for Known Items by Search Type: Scholarly Sources (n=278)

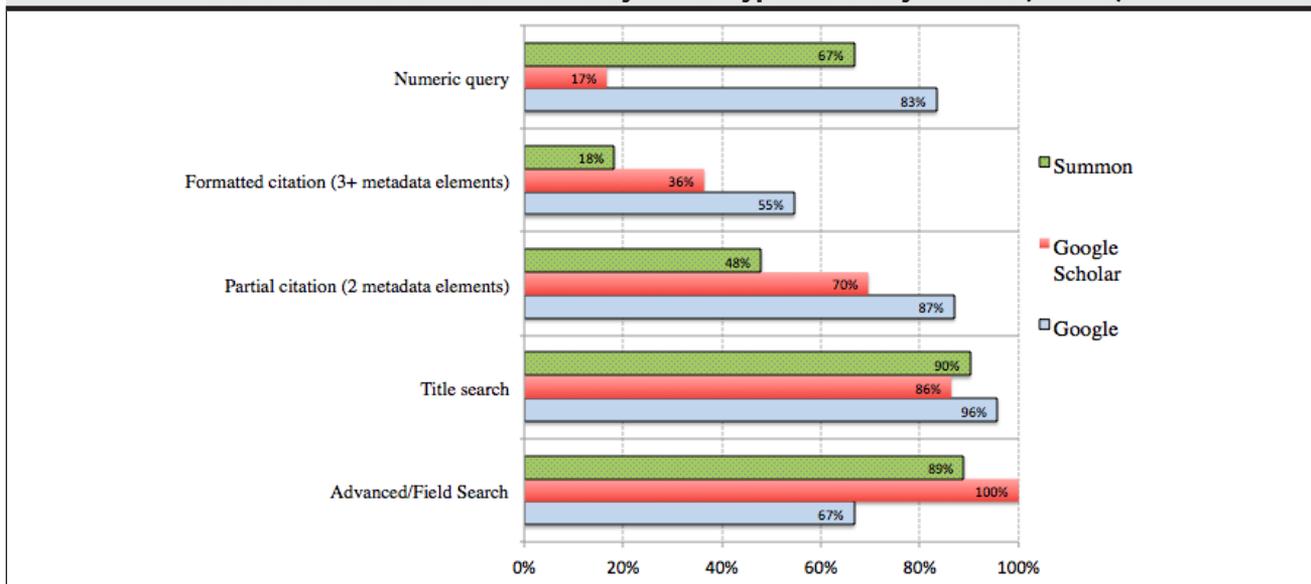


FIGURE 5
Example of a Formatted Citation Query with Three or More Metadata Elements

Brooks M. K. (2010). Hospice services: The dilemmas of technology at the end of life. In T. S. Kerson J. L. M. McCoyd & Associates (Eds.) Social work in health settings: Practice in context (3rd ed. pp. 235-246). New York NY: Routledge.

Discussion

The practice of users copying and pasting entire unedited and formatted citations into a library's unified search box has been documented in several recent studies.⁵⁶ One novel approach to dealing with this user behavior was taken by Mike DeMars who attempted to imitate Google, by "fixing the bad searches" that users were entering into his library's discovery layer. He developed an automated program that looked for errors or problematic queries, and then "scrubbed" them before passing the revised query on to the discovery service. For formatted citations the program looked for "predictable patterns" within APA citations (the most commonly used citation style at his institution). Once the program recognized the APA pattern it stripped out everything but the title of the known item before passing it on to the discovery service. And as our study has shown, title searches are among the most likely search types to retrieve relevant results. DeMars reported that since implementing this "scrubber page" over 1,400 APA formatted citation queries were revised for users "who otherwise may have come to [the library] been frustrated and left and gone to Google."⁵⁷

McKay and Buchanan found that the only search tool that performed well for citation searches was Google Scholar (they did not look at Google).⁵⁸ Yet in our study, Google Scholar's performance was not stellar, failing to return relevant results 64 percent (7/11) of the time. Google on the other hand, only failed to return relevant results for 27 percent (3/11) of the formatted citation searches (this was the search type that Google performed the worst on—figure 3). It is interesting to note that reducing the number of metadata elements in these failed queries (to three) reversed the

outcome: relevant results were found. Thus, it appears that even Google has a limit to how many formatted metadata elements it can handle.

We found that the type of search a user executes does have an impact on the relevancy of results returned, in both discovery systems and in Google and Google Scholar. As a result, a best practice that emerges is to only enter one metadata element (the title) when searching for known items in a discovery layer. Discovery vendors should investigate ways to adapt to searching behavior so that, regardless of the number of metadata elements included in a query (whether formatted or not), the relevancy of results will not suffer. Given the comparable performance of discovery services to Google Scholar, in both this study and Zhang's,⁵⁹ and the daunting task of bridging the relevancy gap with Google, some have begun wondering if libraries should "cede discovery as a function and rely on Google."⁶⁰ While this is not likely to happen, it provides some added motivation for both libraries and vendors to make more significant strides in improving overall discovery and delivery of content.

Relevant Futures

This study set out to find out how well USC's discovery system performed in returning relevant results for known items. Our findings reveal that Summon provides disappointingly average relevancy for users searching for known items (76 percent). In terms of relevancy, the results returned by Summon are, however, an improvement over the results users encounter(ed) when searching traditional library catalogs.⁶¹ It appears that teaching library "machines"⁶² to become better at understanding and adapting to user search behaviors might be as challenging as teaching users' to become better searchers.

Another option to consider, when attempting to improve the relevancy of results provided by library search tools, is incorporating personalization features. Since 2009 Google has been providing customized results based on previous actions taken by users.⁶³ Today, just about every online service, from e-commerce to news and social media sites provide

results enhanced by incorporating various levels of user data.⁶⁴ Yet, personalization, which requires the tracking of user data, is something libraries have been slow to consider, in part due to privacy concerns.⁶⁵ These concerns can be mitigated by carefully and conscientiously leveraging less sensitive usage data and by setting up a transparent opt-in/opt-out system.⁶⁶ Attitudes towards personalization appear to be shifting as more and more librarians and academics are prodding libraries to harness the power of user data to provide more relevant and meaningful results.⁶⁷ One recent proposal, made by David Weinberger from Harvard's Berkman Center for Internet and Society, was to create a "stackscore" which would signify "how relevant an item is to the library's patrons as measured by how they've used it."⁶⁸ He lists numerous datasets that are either already being collected or that could be easily obtained, which could be factored into developing this score, from renewals and recalls to readings listed on a syllabus. Commercial interests motivated the initial push towards offering personalized services, but bringing personalization technologies into libraries holds the promise of enhancing the breadth, depth, and reach of scholarship and scholarly communication in new and exciting ways.

If libraries are to remain relevant in the lives of students and researchers, they must adapt and evolve. When users experience a heightened level of personalized service in almost every other aspect of their lives, and then use a library discovery system, only to encounter "unexpected" or irrelevant results, their impression will likely be that our systems (and services) are out of date or broken.⁶⁹ Part of Google's success is due to its use of personal data to enhance the relevancy of search results for each individual user. Libraries will never succeed in providing a truly Google-like search experience without moving in this direction. By offering personalized search systems, libraries will be better able to serve their users, not just in leading them to relevant content, but in anticipating and meeting their future information needs.

Notes

1. "USC Libraries Online Usage Statistics," last modified January 14, 2015, <http://libguides.usc.edu/usagestats>. In 2011 there were just over 1.5 million searches.
2. Marshall Breeding, "The Future of Library Resource Discovery: A White Paper Commissioned by the NISO Discovery to Delivery (D2D) Topic Committee," *NISO* (February 2015): 30, http://www.niso.org/publications/white_papers/discovery/; Cathy Slaven et al., "From 'I Hate It' to 'It's My New Best Friend!': Making Heads or Tails of Client Feedback to Improve Our New Quick Find Discovery Service," (Australian Library and Information Association Conference & Exhibition, Sydney, N.S.W., February 1, 2011), 6.
3. *Ibid.*, 15-16.
4. Courtney Lundrigan, Kevin Manuel, and May Yan, "Feels Like You've Hit the Lottery: Assessing the Implementation of a Discovery Layer Tool at Ryerson University," *Library and Staff Publications* 19 (2012): 10-11, http://digitalcommons.ryerson.ca/library_pubs/19.
5. Marshall Breeding, "Looking Forward to the Next Generation of Discovery Services," *Computers in Libraries* 32, no. 2 (March 2012): 29.
6. Craig Silverstein et al., "Analysis of a Very Large Web Search Engine Query Log," *ACM SIGIR Forum* 33, no. 1 (September 1, 1999): 12; Laura A. Granka, Thorsten Joachims, and Geri Gay, "Eye-Tracking Analysis of User Behavior in WWW Search," in *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ed. Mark Sanderson, et al. (New York: ACM, 2004), 479; Bing Pan et al., "In Google We Trust: Users' Decisions on Rank, Position, and Relevance," *Journal of Computer-Mediated Communication* 12, no. 3 (April 2007): 803, 813, doi:10.1111/j.1083-6101.2007.00351.x; Bernard J. Jansen, Amanda Spink, and Tefko Saracevic, "Real Life, Real Users, and Real Needs: A Study and Analysis of User Queries on the Web," *Information Processing & Management* 36, no. 2 (March 2000): 212, 214, 225, doi:10.1016/S0306-4573(99)00056-4; Bernard J. Jansen and Amanda Spink, "How Are We Searching the World Wide Web? A Comparison of Nine Search Engine Transaction Logs," *Information Processing & Management* 42 (2006): 260, doi:10.1016/j.ipm.2004.10.007; Andrew D. Asher, Lynda M. Duke, and Suzanne Wilson, "Paths Of Discovery: Comparing the Search Effectiveness of EBSCO Discovery Service, Summon, Google Scholar, and Conventional Library Resources," *College & Research Libraries* 74 (2013): 474, doi:10.5860/crl-374; Helen Georgas, "Google vs. the Library (Part II): Student Search Patterns and Behaviors When Using Google and a Federated Search Tool," *Portal: Libraries and the Academy* 14, no. 4 (2014): 524, doi:10.1353/pla.2014.0034.
7. Jansen and Spink, "How Are We Searching," 260.
8. Pan et al., "In Google We Trust," 813.
9. Mark T. Keane, Maeve O'Brien, and Barry Smyth, "Are People Biased in Their Use of Search Engines?" *Communications of the ACM* 51, no. 2 (February 1, 2008): 51, doi:10.1145/1314215.1314224.
10. *Ibid.*, 52.

11. David Bawden and Lyn Robinson, "The Dark Side of Information: Overload, Anxiety and Other Paradoxes and Pathologies," *Journal of Information Science* 35, no. 2 (November 21, 2008):185, doi:10.1177/0165551508095781.
12. Slaven et al., "From 'I Hate It' to...", 8; Amy Fry and Linda Rich, "Usability Testing for E-Resource Discovery: How Students Find and Choose E-Resources Using Library Web Sites," *The Journal of Academic Librarianship* 37, no. 5 (September 2011): 397, doi:10.1016/j.acalib.2011.06.003; Helen Timpson and Gemma Sansom, "A Student Perspective on E-Resource Discovery: Has the Google Factor Changed Publisher Platform Searching Forever?" *The Serials Librarian* 61, no. 2 (August 2011): 265, doi:10.1080/0361526X.2011.592115.
13. Alison J. Head, "Project Information Literacy: What Can Be Learned About the Information-Seeking Behavior of Today's College Students?" *Association of College and Research Libraries Conference 2013 Proceedings* (Chicago: ACRL, 2013), 475; Sam Brooks, "Discovery: It's About the End User," *Against the Grain* 25, no. 4 (September 2013): 18.
14. Michael Khoo and Catherine Hall, "What Would 'Google' Do? Users' Mental Models of a Digital Library Search Engine," in *Theory and Practice of Digital Libraries*, ed. Panayiotis Zaphiris et al. (Berlin: Springer Berlin Heidelberg, 2012): 9-10.
15. Dana McKay and George Buchanan, "Boxing Clever: How Searchers Use and Adapt to a One-Box Library Search," in *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration* (New York: ACM, 2013), 504, doi:10.1145/2541016.2541031.
16. Jansen and Spink, "How Are We Searching," 256, 259; Asher et al., "Paths of Discovery," 473. Roen Janyk, "Augmenting Discovery Data and Analytics to Enhance Library Services," *Insights* 27 (2014): 264, doi:10.1629/2048-7754.166; McKay and Buchanan, "Boxing Clever: How Searchers," 498.
17. Kristen Calvert, "Maximizing Academic Library Collections: Measuring Changes in Use Patterns Owing to EBSCO Discovery Service," *College & Research Libraries* 76 no.1 (January 2015): 83, doi:10.5860/crl.76.1.81.
18. Corey Davis, "Relevancy Redacted: Web-Scale Discovery and the 'Filter Bubble,'" in *Proceedings of the Charleston Library Conference 2011* (Charleston, SC, 2012), 559, http://dx.doi.org/10.5703/1288284314965.
19. Andrew D. Asher, "Search Magic: Discovering How Undergraduates Locate Information," (paper, American Anthropological Association Annual Meeting, Montreal, Canada, 2011), 6.
20. Janyk, "Augmenting Discovery Data," 266.
21. McKay and Buchanan, "Boxing Clever: How Searchers"; Beth Thomsett-Scott and Patricia E. Reese, "Academic Libraries and Discovery Tools: A Survey of the Literature," *College & Undergraduate Libraries* 19 no.2-4 (2012): 124, doi:10.1080/10691316.2012.697009.
22. Beth Dempsey, "Serials Solutions Advances Library Discovery with Summon 2.0," ProQuest's website, March 18, 2013, http://www.proquest.com/about/news/2013/Serials-Solutions-Advances-Library-Discovery-with-Summon-2-0.html.
23. ProQuest, "Relevance Ranking in the Summon Service," accessed February 18, 2015, 2, http://media2.proquest.com/documents/Summon-RelevanceRanking-Datashet.pdf.
24. "Relevance ranking has been an area of particular investment ...and this area is likely to remain an ongoing development need." Dartmouth College Library, "An Evaluation of Serials Solutions Summon As a Discovery Service for the Dartmouth College Library," November 10, 2009, 3, http://www.dartmouth.edu/~library/admin/docs/Summon_Report.pdf; Margaret G. Grotti and Karen Sobel, "WorldCat Local and Information Literacy Instruction: An Exploration of Emerging Teaching Practice," *Public Services Quarterly* 8, no. 1 (January 2012): 19, doi:10.1080/15228959.2011.563140; Anita K. Foster and Jean B. MacDonald, "A Tale of Two Discoveries: Comparing the Usability of Summon and EBSCO Discovery Service," *Journal of Web Librarianship* 7, no. 1 (2013): 14, doi:10.1080/19322909.2013.757936; "Summon's use of standard information retrieval algorithms to search on a vast corpus of extremely diverse content without the benefit of citation metrics or link-weighting algorithms (like Google's PageRank), often leads to results that aren't well-suited to the research needs of students." Abe Crystal, "NCSU Libraries Summon User Research: Findings and Recommendations," April 13, 2010, NCSU Libraries' website, http://www.lib.ncsu.edu/documents/userstudies/2010summon/Summon-usability.docx.
25. Arta Kabashi, Christine Peterson, and Tim Prather, "Discovery Services: A White Paper For the Texas State Library & Archives Commission," (Amigos Library Services, August 2014), 3, https://www.tsl.texas.gov/sites/default/files/public/tslac/lot/TSLAC_WP_discovery__final_TSLAC_20140912.pdf.
26. Roger C. Schonfeld, "Does Discovery Still Happen in the Library?" *Ithaka S+R* (2014): 7, http://sr.ithaka.org/sites/default/files/files/SR_Briefing_Discovery_20140924_0.pdf
27. Breeding, "Future of Library Resource Discovery," 30.
28. Brent Cook, email message to author, February 19, 2015. (emphasis mine).
29. Stephen Asunka, Hui Soo Chae, Brian Hughes, and Gary Natriell, "Understanding Academic Information Seeking Habits Through Analysis of Web Server Log Files: The Case Of The Teachers College Library Website," *The Journal of Academic Librarianship* 35 (2009): 36, doi:10.1016/j.acalib.2008.10.019; Jansen, "Search Log Analysis," 409.
30. Jansen, "Search Log Analysis," 409.
31. Asunka et al., "Real Life, Real Users," 37.
32. Janset et al., "Real Life, Real Users"; Eng Pwey Lau and Dion Hoe-Lian Goh, "In Search of Query Patterns: A Case Study of a University OPAC," *Information Processing & Management* 42, no. 5 (2006), doi:http://dx.doi.org/10.1016/j.ipm.2006.02.003; Jansen and Spink, "How Are We Searching"; Asunka et al., "Real Life, Real Users"; Dietmar Wolfram, "Search Characteristics in Different Types of Web-Based IR Environments: Are They The Same?" *Information Processing & Management* 44 (2008), doi:10.1016/j.ipm.2007.07.010.
33. Dana McKay and George Buchanan, "One of These Things Is Not Like the Others: How Users Search Different Information Resources," in *Research and Advanced Technology for*

- Digital Libraries*, ed. Stefan Gradmann et al. (Berlin: Lecture Notes in Computer Science, 2011), 267, doi:10.1007/978-3-642-24469-8_28.
34. McKay and Buchanan, "Boxing Clever: How Searchers Use," 497.
 35. *Ibid.*, 499.
 36. *Ibid.*, 505.
 37. Kelly Meadow and James Meadow, "Search Query Quality And Web-Scale Discovery: A Qualitative and Quantitative Analysis," *College & Undergraduate Libraries* 19 (2012): 165, doi:10.1080/10691316.2012.693434.
 38. Suzanne Chapman et al., "Manually Classifying User Search Queries on an Academic Library Web Site," *Journal of Web Librarianship* 7 (2013): 406, doi:10.1080/19322909.2013.842096.
 39. *Ibid.*, 402.
 40. Zhang, Tao, "User-Centered Evaluation of a Discovery Layer System with Google Scholar," in *Design, User Experience, and Usability: Web, Mobile, and Product Design*, ed. Aaron Marcus (Berlin: Springer Berlin Heidelberg, 2013), 321, http://link.springer.com/10.1007/978-3-642-39253-5_34.
 41. Emily Singley, "Discovery Systems—Testing Known Item Searching," *Usable Libraries*, (March 18, 2014), <http://emilysingley.net/discovery-systems-testing-known-item-searching/>.
 42. Elizabeth Namei and Christal A. Young, "Letting User Search Behavior Lead Us: An Analysis Of Search Queries & Relevancy In USC's Web-Scale Discovery Tool," *California Association of Research Libraries Conference Proceedings* (San Jose, CA, April 5, 2014), 3, <http://carl-conference.org/sites/carl-conference.org/files/slides/nameiyoung.pdf>. We followed much of the same methodology we used in our 2014 study.
 43. Chapman et al., "Manually Classifying User Search Queries," 413. A similar approach was used in their study.
 44. Brophy and Bawden, "Is Google Enough?" 504; Maglaughlin and Sonnenwald, "User Perspectives On Relevance," 328.
 45. McKay and Buchanan, "One of These Things," 261.
 46. Lau and Goh, "In Search of Query Patterns," 263.
 47. Xi Niu, Tao Zhang, and Hsin-liang Chen, "Study of User Search Activities With Two Discovery Tools at an Academic Library," *International Journal of Human-Computer Interaction* 30, no. 5 (May 4, 2014): 425, doi:10.1080/10447318.2013.873281.
 48. Jan Brophy and David Bawden, "Is Google Enough? Comparison of an Internet Search Engine with Academic Library Resources," *Aslib Proceedings: New Information Perspectives* (2005): 504; Georgas, "Google vs. the Library," 524, doi:10.1108/00012530510634235.
 49. (Asher et al., "Paths of Discovery," 474; Silverstein et al., "Analysis of a Very Large Web Search Engine," 12; Granka et al., "Eye-Tracking Analysis of User Behavior," 479; Pan et al., "In Google We Trust," 803, 813; Jansen et al., "Real Life, Real Users, and Real Needs," 212, 214, 225; Jansen and Spink, "How Are We Searching," 260.
 50. Steven. E. Stemler, "A Comparison of Consensus, Consistency, and Measurement Approaches to Estimating Interrater Reliability," *Practical Assessment, Research & Evaluation* 9 no. 4 (2004), <http://pareonline.net/getvn.asp?v=9&n=4>.
 51. Robert V. Krejcie and Daryle W. Morgan, "Determining Sample Size For Research Activities," *Educational and Psychological Measurement* 3 (1970): 608.
 52. Namei and Young, "Letting User Search Behavior Lead," 4. Based on this study, we found that 36% (+/- .05) of all searches executed in our Summon instance were for known items. This allowed us to estimate the number of queries we would need to pull from the logs from which we could then extract the requisite number of known item searches. We ended up needing 1,150 queries to get our final sample (This number corroborates our 2014 study's findings for the percentage of known items searches).
 53. Head, "Project Information Literacy," 476, 478; Lisa M. Rose-Wiles and Melissa A. Hofmann, "Still Desperately Seeking Citations: Undergraduate Research in the Age of Web-Scale Discovery," *Journal of Library Administration* 53, no. 2-3 (February 2013): 155, doi:10.1080/01930826.2013.853493; Jessica Mussell and Rosie Croft, "Discovery Layers and the Distance Student: Online Search Habits of Students," *Journal of Library & Information Services in Distance Learning* 7, no. 1-2 (January 2013): 19, doi:10.1080/1533290X.2012.705561.
 54. Zhang, "User-Centered Evaluation," 320.
 55. McKay and Buchanan, "Boxing Clever: How Searchers," 503; "In general, when query length increases, there is a higher probability that users would encounter unsuccessful searches." Lau and Goh, "In Search of Query Patterns," 1324.
 56. Mike DeMars, "Teaching Machines: Creating Better Search Engines" (conference presentation), *Internet Librarian Conference*, October 28, 2013, video, 45:06, <https://vimeo.com/78156437>; "26% of participants entered whole references and 44% entered partial references (e.g. author, year and title) into the [discovery] search box when looking for an item from a reference." Keren Mills, "What Do Students Want from Discovery Tools?" slides presented at the *Internet Librarian International Conference*, London, UK, October 23, 2014, slide 23, <http://www.slideshare.net/mirya/ili14-library-futures-what-do-students-want-from-discovery-tools-slideshare>; McKay and Buchanan, "One of These Things," 264.
 57. DeMars, "Teaching Machines."
 58. McKay and Buchanan, "One of These Things," 267.
 59. Zhang, "User-Centered Evaluation," 320.
 60. Schonfeld, "Does Discovery Still Happen," 11; Herrera suggests that "when facing budgetary crises, subject librarians... may be forced to consider the usefulness of GS over less-used citation databases." Gail Herrera, "Google Scholar Users and User Behaviors: An Exploratory Study," *College & Research Libraries* 72, no. 4 (July 1, 2011): 329, doi:10.5860/crl-125rl.
 61. McKay and Buchanan, "One of These Things," 260; McKay and Buchanan, "Boxing Clever: How Searchers," 498; Lau and Goh, "In Search of Query Patterns," 1316.
 62. DeMars, "Teaching Machines."
 63. Davis, "Relevancy Redacted," 556.
 64. Dartmouth College Library, "An Evaluation of Summon," 10.

65. There are numerous ways to address the privacy concerns. For instance, libraries could by default make these personalized systems opt-in only, and make it extremely transparent how to opt out or to fluctuate between participating and not participating. Ken J. Varnum, "Library Discovery From Ponds to Streams," ed. K. J. Varnum, *The Top Technologies Every Librarian Needs to Know: A LITA Guide* (Chicago, IL: ALA TechSource, 2014), 63. <http://hdl.handle.net/2027.42/107042>; Karen Coombs, "Privacy vs. Personalization," *Library Journal* 28 (2007): 26.
66. Varnum, "Library Discovery From Ponds to Streams," 63-64.
67. Breeding, "Looking Forward to the Next Generation," 30; Varnum, "Top Technologies," 61-3; Sharon Q. Yang, "Charting Discovery System Improvements: (2010-2013)," *Computers in Libraries* 34 (2014): 12; Roger Schonfeld, "Anticipatory Discovery for Current Awareness of New Publications" (conference presentation), LITA Top Tech Trends 2014, American Libraries Association Conference, June 29, 2014, video, 1:28:51, <http://www.ala.org/lita/ttt/2014annual>.
68. David Weinberger, "A Good, Dumb Way to Learn From Libraries," *The Chronicle of Higher Education Blogs: The Conversation*, October 7, 2014, <http://chronicle.com/blogs/conversation/2014/10/07/a-good-dumb-way-to-learn-from-libraries/>.
69. Brian Matthews, "Database vs. Database vs. Web-Scale Discovery Service: Further Thoughts on Search Failure (Or: More Clicks Than Necessary?) (Or: Info-Pushers vs. Pedagogical Partners)," *Ubiquitous Librarian* (blog), *The Chronicle of Higher Education*, August 21, 2013. <http://chronicle.com/blognetwork/theubiquitouslibrarian>.
- Paper Commissioned by the NISO Discovery to Delivery (D2D) Topic Committee." *NISO* (February 2015): 1-53. www.niso.org/publications/white_papers/discovery/.
- Brooks, Sam. "Discovery: It's About the End User." *Against the Grain* 25, no. 4 (September 2013): 18.
- Brophy, Jan and David Bawden. "Is Google Enough? Comparison of an Internet Search Engine with Academic Library Resources." *Aslib Proceedings: New Information Perspectives* (2005): 498-512. doi:10.1108/00012530510634235.
- Calvert, Kristen. "Maximizing Academic Library Collections: Measuring Changes in Use Patterns Owing to EBSCO Discovery Service." *College & Research Libraries* 76 no. 1 (January 2015): 81-99. doi:10.5860/crl.76.1.81.
- Chapman, Suzanne, Shevon Desai, Kat Hagedorn, Ken Varnum, Sonali Mishra, and Julie Piacentine. "Manually Classifying User Search Queries on an Academic Library Web Site," *Journal of Web Librarianship* 7 (2013): 401-421. doi:10.1080/19322909.2013.842096.
- Crystal, Abe. "NCSU Libraries Summon User Research: Findings and Recommendations." April 13, 2010. NCSU Libraries' website. <http://www.lib.ncsu.edu/documents/userstudies/2010summon/Summon-usability.docx>.
- Davis, Corey. "Relevancy Redacted: Web-Scale Discovery and the 'Filter Bubble.'" In *Proceedings of the Charleston Library Conference 2011* (Charleston, SC, 2012), 556-62. <http://dx.doi.org/10.5703/1288284314965>.
- DeMars, Mike. "Teaching Machines: Creating Better Search Engines" (conference presentation). *Internet Librarian Conference*. 2013. Video, 45:06. <https://vimeo.com/78156437>.
- Dempsey, Beth. "Serials Solutions Advances Library Discovery with Summon 2.0." ProQuest's website. March 18, 2013. <http://www.proquest.com/about/news/2013/Serials-Solutions-Advances-Library-Discovery-with-Summon-2-0.html>.
- Dartmouth College Library. "An Evaluation of Serials Solutions Summon As a Discovery Service for the Dartmouth College Library." November 10, 2009. 1-29. http://www.dartmouth.edu/~library/admin/docs/Summon_Report.pdf.
- Foster, Anita K., and Jean B. MacDonald. "A Tale of Two Discoveries: Comparing the Usability of Summon and EBSCO Discovery Service." *Journal of Web Librarianship* 7, no. 1 (2013): 1-19. doi:10.1080/19322909.2013.757936.
- Fry, Amy, and Linda Rich. "Usability Testing for E-Resource Discovery: How Students Find and Choose E-Resources Using Library Web Sites." *The Journal of Academic Librarianship* 37, no. 5 (September 2011): 386-401. doi:10.1016/j.acalib.2011.06.003.
- Georgas, Helen. "Google vs. the Library (Part II): Student Search Patterns and Behaviors When Using Google and a Federated Search Tool." *Portal: Libraries and the Academy* 14, no. 4 (2014): 503-32. doi:10.1353/pla.2014.0034.
- Granka, Laura A., Thorsten Joachims, and Geri Gay. "Eye-Tracking Analysis of User Behavior in WWW Search." In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. Edited by Mark Sanderson, et al. (New York: ACM, 2004), 478-479.
- Grotti, Margaret G., and Karen Sobel. "WorldCat Local and Information Literacy Instruction: An Exploration of Emerg-

Bibliography

Asher, Andrew D. "Search Magic: Discovering How Undergraduates Locate Information." Paper presented at the American Anthropological Association Annual Meeting, Montreal, Canada, 2011.

Asher, Andrew D., Lynda M. Duke, and Suzanne Wilson. "Paths of Discovery: Comparing the Search Effectiveness of EBSCO Discovery Service, Summon, Google Scholar, and Conventional Library Resources." *College & Research Libraries* 74 (2013): 464-88. doi:10.5860/crl-374.

Asunka, Stephen, Hui Soo Chae, Brian Hughes, and Gary Natriell. "Understanding Academic Information Seeking Habits Through Analysis of Web Server Log Files: The Case Of The Teachers College Library Website." *The Journal of Academic Librarianship* 35 (2009): 33-45. doi:10.1016/j.acalib.2008.10.019.

Coombs, Karen. "Privacy vs. Personalization." *Library Journal* 28 (2007): 26.

Bawden, David, and Lyn. Robinson. "The Dark Side of Information: Overload, Anxiety and Other Paradoxes and Pathologies." *Journal of Information Science* 35, no. 2 (November 21, 2008): 180-91. doi:10.1177/0165551508095781.

Breeding, Marshall. "Looking Forward to the Next Generation of Discovery Services." *Computers in Libraries* 32, no. 2 (March 2012): 28-31.

———. "The Future of Library Resource Discovery: A White

- ing Teaching Practice." *Public Services Quarterly* 8, no. 1 (January 2012): 12–25. doi:10.1080/15228959.2011.563140.
- Head, Alison J. "Project Information Literacy: What Can Be Learned About the Information-Seeking Behavior of Today's College Students?" In *Association of College and Research Libraries Conference 2013 Proceedings* (Chicago: ACRL, 2013), 472–482.
- Herrera, Gail. "Google Scholar Users and User Behaviors: An Exploratory Study." *College & Research Libraries* 72, no. 4 (July 1, 2011): 316–30. doi:10.5860/crl-125rl.
- Jansen, Bernard J. "Search Log Analysis: What It Is, What's Been Done, How To Do It." *Library & Information Science Research* 28 (2006): 407–32. doi:10.1016/j.lisr.2006.06.005.
- Jansen, Bernard J., Amanda Spink, and Tefko Saracevic. "Real Life, Real Users, and Real Needs: A Study and Analysis of User Queries on the Web." *Information Processing & Management* 36, no. 2 (March 2000): 207–27. doi:10.1016/S0306-4573(99)00056-4.
- Jansen, Bernard J., and Amanda Spink. "How Are We Searching the World Wide Web? A Comparison of Nine Search Engine Transaction Logs." *Information Processing & Management* 42 (2006): 248–63. doi:10.1016/j.ipm.2004.10.007.
- Janyk, Roen. "Augmenting Discovery Data and Analytics to Enhance Library Services." *Insights* 27 (2014): 262–267. doi:10.1629/2048-7754.166.
- Kabashi, Arta, Christine Peterson, and Tim Prather. "Discovery Services: A White Paper For the Texas State Library & Archives Commission." (Amigos Library Services, August 2014). 1–35. https://www.tsl.texas.gov/sites/default/files/public/tslac/lot/TSLAC_WP_discovery__final_TSLAC_20140912.pdf.
- Keane, Mark T., Maeve O'Brien, and Barry Smyth. "Are People Biased in Their Use of Search Engines?" *Communications of the ACM* 51, no. 2 (February 1, 2008): 49–52. doi:10.1145/1314215.1314224.
- Khoo, Michael, and Catherine Hall. "What Would 'Google' Do? Users' Mental Models of a Digital Library Search Engine." In *Theory and Practice of Digital Libraries*. Edited by Panayiotis Zaphiris et al. (Berlin: Springer Berlin Heidelberg, 2012), 1–12.
- Krejcie, Robert V., and Daryle W. Morgan. "Determining Sample Size For Research Activities." *Educational and Psychological Measurement* 3 (1970): 607–610.
- Lau, Eng Pwey and Dion Hoe-Lian Goh. "In Search of Query Patterns: A Case Study of a University OPAC." *Information Processing & Management* 42, no. 5 (2006): 1316–29. doi:http://dx.doi.org/10.1016/j.ipm.2006.02.003.
- Lundrigan, Courtney, Kevin Manuel, and May Yan. "'Feels Like You've Hit the Lottery': Assessing the Implementation of a Discovery Layer Tool at Ryerson University." *Library and Staff Publications* 19 (2012): 1–25. http://digitalcommons.ryerson.ca/library_pubs/19
- McKay, Dana, and George Buchanan. "One of These Things Is Not Like the Others: How Users Search Different Information Resources." In *Research and Advanced Technology for Digital Libraries*, edited by Stefan Gradmann, Francesca Borri, Carlo Meghini, and Heiko Schuldt, 260–271. Berlin: Lecture Notes in Computer Science, 2011. doi:10.1007/978-3-642-24469-8_28.
- . "Boxing Clever: How Searchers Use and Adapt to a One-Box Library Search." In *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration* (New York: ACM, 2013), 497–506. doi:10.1145/2541016.2541031.
- Maglaughlin, Kelly L., and Diane H. Sonnenwald. "User Perspectives on Relevance Criteria: A Comparison Among Relevant, Partially Relevant, and Not-Relevant Judgments." *Journal of the American Society for Information Science and Technology* 53 (2002): 327–342. doi:10.1002/asi.10049.
- Matthews, Brian. "Database vs. Database vs. Web-Scale Discovery Service: Further Thoughts on Search Failure (Or: More Clicks Than Necessary?) (Or: Info-Pushers vs. Pedagogical Partners)." *Ubiquitous Librarian* (blog). The Chronicle of Higher Education. August 21, 2013. <http://chronicle.com/blognetwork/theubiquitouslibrarian>.
- Meadow, Kelly and James Meadow. "Search Query Quality And Web-Scale Discovery: A Qualitative and Quantitative Analysis." *College & Undergraduate Libraries* 19 (2012): 163–75. doi:10.1080/10691316.2012.693434.
- Mills, Keren. "What Do Students Want from Discovery Tools?" Slides presented at the Internet Librarian International Conference. London, UK. October 23, 2014. <http://www.slideshare.net/mirya/ili14-library-futures-what-do-students-want-from-discovery-tools-slideshare>.
- Namei, Elizabeth and Christal A. Young. "Letting User Search Behavior Lead Us: An Analysis of Search Queries & Relevancy in USC's Web-Scale Discovery Tool." *California Association of Research Libraries Conference Proceedings*, San Jose, CA, April 5, 2014, <http://carl-conference.org/sites/carl-conference.org/files/slides/nameiyoung.pdf>.
- Mussell, Jessica, and Rosie Croft. "Discovery Layers and the Distance Student: Online Search Habits of Students." *Journal of Library & Information Services in Distance Learning* 7, no. 1–2 (January 2013): 18–39. doi:10.1080/1533290X.2012.705561.
- Niu, Xi, Tao Zhang, and Hsin-liang Chen. "Study of User Search Activities With Two Discovery Tools at an Academic Library." *International Journal of Human-Computer Interaction* 30, no. 5 (May 4, 2014): 422–33. doi:10.1080/10447318.2013.873281.
- Pan, Bing, Helene Hembrooke, Thorsten Joachims, Lori Lorigo, Geri Gay, and Laura Granka. "In Google We Trust: Users' Decisions on Rank, Position, and Relevance." *Journal of Computer-Mediated Communication* 12, no. 3 (April 2007): 801–823. doi:10.1111/j.1083-6101.2007.00351.x.
- ProQuest. "Relevance Ranking in the Summon Service." Accessed February 18, 2015. 1–4. <http://media2.proquest.com/documents/Summon-RelevanceRanking-Datasheet.pdf>
- Rose-Wiles, Lisa M., and Melissa A. Hofmann. "Still Desperately Seeking Citations: Undergraduate Research in the Age of Web-Scale Discovery." *Journal of Library Administration* 53, no. 2–3 (February 2013): 147–66. doi:10.1080/01930826.2013.853493.
- Schonfeld, Roger C. "Does Discovery Still Happen in the Library?" *Ithaca S+R* (2014): 1–12. http://sr.ithaka.org/sites/default/files/files/SR_Briefing_Discovery_20140924_0.pdf
- . "Anticipatory Discovery for Current Awareness of New Publications" (conference presentation). *LITA Top Tech*

- Trends 2014, American Libraries Association Conference.* June 29, 2014. <http://www.ala.org/lita/ttt/2014annual>.
- Silverstein, Craig, Hannes Marais, Monika Henzinger, and Michael Moricz. "Analysis of a Very Large Web Search Engine Query Log." *ACM SIGIR Forum* 33, no. 1 (September 1, 1999): 6–12. doi:10.1145/331403.331405.
- Singley, Emily. "Discovery Systems—Testing Known Item Searching." *Usable Libraries*, (March 18, 2014), <http://emilysingley.net/discovery-systems-testing-known-item-searching/>.
- Slaven, Cathy, Barb Ewers, Jessie Donaghey, and Kurt Vollmerhouse. "From 'I Hate It' to 'It's My New Best Friend!': Making Heads or Tails of Client Feedback to Improve Our New Quick Find Discovery Service." (Australian Library and Information Association Conference & Exhibition, Sydney, N.S.W., February 1, 2011): 1-16. http://conferences.alia.org.au/online2011/papers/paper_2011_A3.pdf.
- Stemler, Steven. E. "A Comparison of Consensus, Consistency, and Measurement Approaches to Estimating Interrater Reliability." *Practical Assessment, Research & Evaluation* 9 no. 4 (2004). <http://pareonline.net/getvn.asp?v=9&n=4>.
- Thomsett-Scott, Beth and Patricia E. Reese. "Academic Libraries and Discovery Tools: A Survey of the Literature." *College & Undergraduate Libraries* 19 no.2-4 (2012): 123-143. doi:10.1080/10691316.2012.697009.
- Timpson, Helen and Gemma Sansom. "A Student Perspective on E-Resource Discovery: Has the Google Factor Changed Publisher Platform Searching Forever?" *The Serials Librarian* 61, no. 2 (August 2011): 253–66. doi:10.1080/0361526X.2011.592115.
- "USC Libraries Online Usage Statistics." Last modified January 14, 2015. <http://libguides.usc.edu/usagestats>.
- Varnum, Ken J. "Library Discovery From Ponds to Streams." Edited by K. J. Varnum. *The Top Technologies Every Librarian Needs to Know: A LITA Guide* (Chicago, IL: ALA Tech-Source, 2014), 57-65. <http://hdl.handle.net/2027.42/107042>.
- Weinberger, David. "A Good, Dumb Way to Learn From Libraries." *The Chronicle of Higher Education Blogs: The Conversation*. October 7, 2014. <http://chronicle.com/blogs/conversation/2014/10/07/a-good-dumb-way-to-learn-from-libraries/>.
- Wolfram, Dietmar. "Search Characteristics in Different Types of Web-Based IR Environments: Are They The Same?" *Information Processing & Management* 44 (2008): 1279-1292. doi:10.1016/j.ipm.2007.07.010.
- Yang, Sharon Q. "Charting Discovery System Improvements: (2010-2013)." *Computers in Libraries* 34 (2014): 10-14.
- Zhang, Tao. "User-Centered Evaluation of a Discovery Layer System with Google Scholar." In *Design, User Experience, and Usability: Web, Mobile, and Product Design*. Edited by Aaron Marcus (Berlin, Heidelberg: Springer Berlin Heidelberg, 2013), 313–22. http://link.springer.com/10.1007/978-3-642-39253-5_34.