Did We Get What We Paid For? An Institutional Assessment of User Behavior and System Performance of Web-Scale Discovery Tools

Doralyn Rossmann, Brian Rossmann, Kelly Meadow, Scott Young, and Greg Notess
Montana State University, USA

Abstract

The purpose of this study is to assess an implemented web-scale discovery (WSD) tool with regard to user behavior, system performance, and collection coverage. An academic library that implemented Serials Solutions' Summon WSD in July 2010 serves as the source of data for this study. The assessment consists of four points of analysis. First, a quantitative design is used to assess link success from search results using a set of researcher-generated search queries. These results are categorized into full-text and non-full-text links, and are reviewed for their success in reaching the targeted item, including how many clicks it took to reach the item. The second portion of the study uses Summon transaction log analysis over a two-year period and addresses the hypotheses that search query quality is low and that quality of searches improved over the two-year period. The third section examines Google Analytics data for Summon for query types and overall Summon usage in comparison to the main library website. Finally, the study evaluates the coverage of library holdings in Summon and explores the implications of these findings.

This research design and subsequent findings provide other libraries with tools and benchmarks for conducting similar studies of WSD tools prior to selecting a product and after having implemented one. Also, in addition to assessing these tools, the findings will have implications for better understanding: (1) user behavior, (2) system performance, and (3) collection coverage and implications relative to WSD tools. These studies will inform how libraries can work with vendors and users to bridge gaps between these three areas. Originally, the study was to include a fourth area, user expectations, but that analysis was not complete at the time of this writing.

Introduction

Montana State University (MSU) Library acquired a web-scale discovery (WSD) tool, Serials Solutions' Summon, in July 2010 with a three-year contract. A search box for Summon, locally branded "CatSearch," was placed front-and-center on the main web page for the Library. With the contract at its midway point, the Dean of the MSU Library formed a Summon Assessment Group to assess the Summon product. It is intended that the results of this study will serve several purposes: help inform MSU Library's decisions when the contract for Summon comes up for renewal, inform MSU Library about how its patrons use Summon, suggest areas it might pursue to improve Summon functionality, and determine the location and promotion of whatever WSD tool is on the MSU Library's website. The study intentionally uses a variety of information sources and data points for the assessment to provide a more complete picture of Summon usage and functionality. This research design and subsequent findings provide other libraries with tools and benchmarks for conducting similar studies of WSD tools prior to selecting a product and after having implemented one.

Review of the Literature:

Because of the relatively new nature of WSD tools, the literature is just now beginning to offer a wide variety of approaches to assessing such products. As of the last Library Assessment Conference in October 2010, five major WSD tools existed and three of those had been released just that year: OCLC's WorldCat Local (released November 2007), Serials Solutions' Summon (released July 2009), EBSCO's EDS (released January 2010), Innovative Interfaces' Encore Synergy (released May 2010), and Ex Libris' Primo Central (released June 2010).[1] When the Summon Assessment Group first convened in January 2012, it considered issues like the "Google factor" and how users might interact with a search box located front-and-center on the

library's web site.[2] Other studies, such as those at Dartmouth, highlight the need to understand what is indexed in Summon relative to library holdings and subscriptions.[3] At the beginning of the study, much of the information available was anecdotal from MSU librarians using Summon.[4] The Summon Assessment Group decided to expand its body of knowledge to include more substantive data such as linking success rates from Summon to information resources. While the value of having qualitative data from user experience testing was identified, time and staffing constraints did not allow for exploration into that area in this phase of the project.[5]

Areas of Analysis:

Summon Link Analysis

Purpose:

While users report anecdotal issues with broken links to search results or having to click multiple times to get to the desired item, it was decided that studying click-through success rates over time would provide a clearer picture of system function. The following hypotheses were established:

- The majority of successful full-text links take users three or fewer clicks to reach the full-text item.
- Successful linking to full-text resources improved during the first two years of implementation.

Methods:

A quantitative design is used to assess link success from search results using a set of researcher-selected search queries. These twenty-six topics come from actual queries found in the Summon usage logs. They were chosen by a member of the research team who identified these as searches for a subject rather than known items, such as specific books or journal articles. These queries were selected as they reflected a snapshot of different subject areas in Summon. These results are categorized into full-text and non-full-text links and are reviewed for their success in reaching the targeted item, including how many clicks it took to reach the item. These queries were conducted three times: in fall 2010, fall 2011, and summer 2012. The first twenty-five links for each query were included in the study, for a total of 650 items analyzed each time.

Results:

Table 1 shows that while failed links to full text dropped from the first year of analysis to the second, they increased slightly from year two to three. Still, the overall failure rate between years one and three declined significantly. Of the failed links in 2012, 34 percent were from the Lexis Nexis Academic database, 13 percent were not clear in their source as they went directly to a publisher site, and 13 percent were from Gale Opposing Viewpoints in Context. The remaining errors were less than ten percent per database spread over 25 databases. Removing the highest failed link source, LexisNexis Academic, from the results still results in a 20 percent failure rate. The study's first hypothesis, that successful linking to full-text resources improved during the first two years of implementation, is supported by these findings, but it is worth noting that the success rate was higher mid-study than in the latest year of analysis.

Table 1: Failed links to full text

Semester	Percentage failure rate
Fall 2010	45
Fall 2011	23
Summer 2012	27

Table 2 shows that in the 2012 portion of the study, 97 percent of successful links required three or

fewer clicks to reach the item. This supports the study's hypothesis that the majority of successful full text links take users three or fewer clicks to reach the full text item.

Table 2: Full text: number of clicks to reach full text,

Failed link	158
1 click	82
2 clicks	194
3 clicks	136
4 clicks	6
5 clicks	4
6 clicks	2
Non-full-text link	68
Total	650

Discussion:
Looking at the figures from 2010, 2011, and 2012, there are some known changes that have taken place during the time of the study, and there are other factors that may account for the presence of errors. Since its product launch, Serials Solutions has made improvements to its system with fixes and enhancements launched every two and, later, every three weeks. Also, some vendors have improved the quality of their linking. For example, a vendor had been putting a hyphen in its metadata for journals which included two issues published together, which caused a broken link from Summon. This vendor now follows the proper OpenURL metadata standard for this type of citation, which results in properly functioning links. Locally, MSU Library began regularly reporting errors to Serials Solutions or the source vendor when encountering problems. Also, MSU Library started going through its list of resources in Serials Solutions 360 Resource Manager (its electronic resource management system) to reduce erroneously activated titles to which it does not

have access. Five main factors can be identified as possible sources of the remaining errors:
- The Summon system with its indexing and linking technologies.
- The content provider itself with metadata or linking technologies.
- The in-between step of the OpenURL resolver used by MSU Library, Serials Solutions' 360 Link, could be improperly resolving links.
- The link within the 360 Resource Manager to a resource or the items listed as contained in that source could be inaccurate.
- MSU Library may have erroneously selected a source as a part of its subscription base when it, in fact, is not.

Conclusion:
It is difficult to isolate some of the sources of problems with broken links encountered in any WSD tool. MSU Library hopes to get a better sense of the role of the Summon product and the 360 Link product by having colleagues at other libraries run the same searches against systems

that use different WSD and OpenURL Linking products. Also, it has since turned off access to LexisNexis Academic in the Summon search results since it is such a highly problematic source, and it corrected a problem in the URL used for Business Insights, which has resolved issues with that source. It continues to report problems to Serials Solutions and content providers. Looking beyond MSU Library, this type of study can help any library get away from anecdotal reports as a means of assessing a WSD and can help identify areas that can be improved.

Summon Transaction Log Analysis
Purpose:
The queries typed into a search box can be a window into a better understanding of user behavior within an information retrieval system. When Summon went live at MSU Library in mid-2010, the Summon Administration Console was not made available until later that year. The Summon Administration Console provides statistics on volume of use as well as the queries entered into the Summon search box, which can help with understanding the quality of search queries performed in Summon and how these queries looked over multiple semesters. It was a startling discovery to learn that *facebook.com* was the most common query entered into the search box when looking at the Summon Administration Console in late 2010. The frequency of *facebook.com* was problematic due to the fact that Summon does not support web addresses the way search engines do. In April of 2011, a member of the Summon Assessment Group performed her own study to better understand the *facebook.com* phenomenon.[6] The following hypotheses were established:
- Queries performed within Summon are of a low quality.
- Query quality improved during the first two semesters of implementation.

Samples of 100 queries per month from August 2010 through April 2011 were coded by query type, resulting in a dataset of 900 queries divided into query types. Surprisingly, both of these hypotheses were proven false as will be shown in more detail in the methods and results section of this report. As a part of the formal Summon assessment, a second iteration of this study was conducted with some minor changes. Here, what follows is a comparison of the data from both studies.

Methods:
The total number of queries from each month from August 2010 through April 2011 and August 2011 through April 2012 were downloaded from the Summon Administration Console. All query processing and statistical analyses were performed in the R data analysis software application.[7] The dataset was constructed from a stratified random sample of 100 queries for each month. Since Serials Solutions lists the query and its frequency, queries were multiplied by how many times they were performed before the random sample was extracted. The month of origin for each query was retained for analysis of results. Thus, a total of 900 queries were used for the first year of study and 1,000 queries were used for the second year of study.

The majority of queries listed in the Serials Solutions administration console were blanks, meaning nothing was entered into the search box. A blank search box could mean that users are skipping the search box and navigating to other areas in the interface, such as the advanced search; thus, blank queries were removed from the data set. In the first year's study, in order to input the data into R, some special characters, such as *, {}, "", and ~, were also removed. The removal of these characters from the dataset eliminated some potentially useful information, especially since Summon supports the use of * and "" symbols. A workaround for the removal of special characters was established for the second year's study; thus, special characters were included in the samples for the second year.

Each of the original 900 queries was assigned to one of the following seven query types: URL, invalid, natural language, database/journal, subject, known-item, and Boolean operator. The 1,000 queries in the second year of the study were assigned one of eight query types. It was decided that determining which of the invalid searches could have been considered a site search was useful information; thus, a site-search category was added for the second year. The subjectivity of query quality and type was addressed by creating a set of rules to determine the query types.[8]

After the queries were coded, the query types were grouped into high and low quality. The concept of quality for this study was dependent upon whether

474

or not the system would support the methods implemented within the query. For example, a URL is an effective way to locate information within a search engine; however, URLs are not supported by Summon and thus they were labeled as a low-quality query. The low-quality grouping consisted of the URL, invalid, site search, and natural language query types. The high-quality grouping consisted of the subject, Boolean operator, known-item and database/journal query types.

In the first year's study, two-sample t-tests were used to compare the proportions of high-versus low-quality queries for the entire academic year, as well as for both semesters individually, and the proportions of high-quality searches by semester were also compared to detect a change in search quality between the semesters. Simple linear regression was performed on each search type, as well as for the combined high-quality queries,

through the academic year to test for a change in any of the categories over time.

Once the study was performed again in the second year, two-sample t-tests were performed on each search type to detect significant changes between each year. The site-search type was included in the invalid type for this comparison.

Results:
Figure 1 shows the percentages of each query type by month from August 2010 through April 2011, with high-quality categories as cool colors and low-quality as warm colors. Subjects were clearly the most common type of query, followed by known items. It is apparent that low-quality queries are a small proportion of the overall queries coded in the study. This also shows that the quality of queries did not appear to change over time.

Figure 1: Frequency of query types by month for August 2010–April 2011. Low-quality queries are displayed in warm colors. High-quality queries are displayed in cool colors.
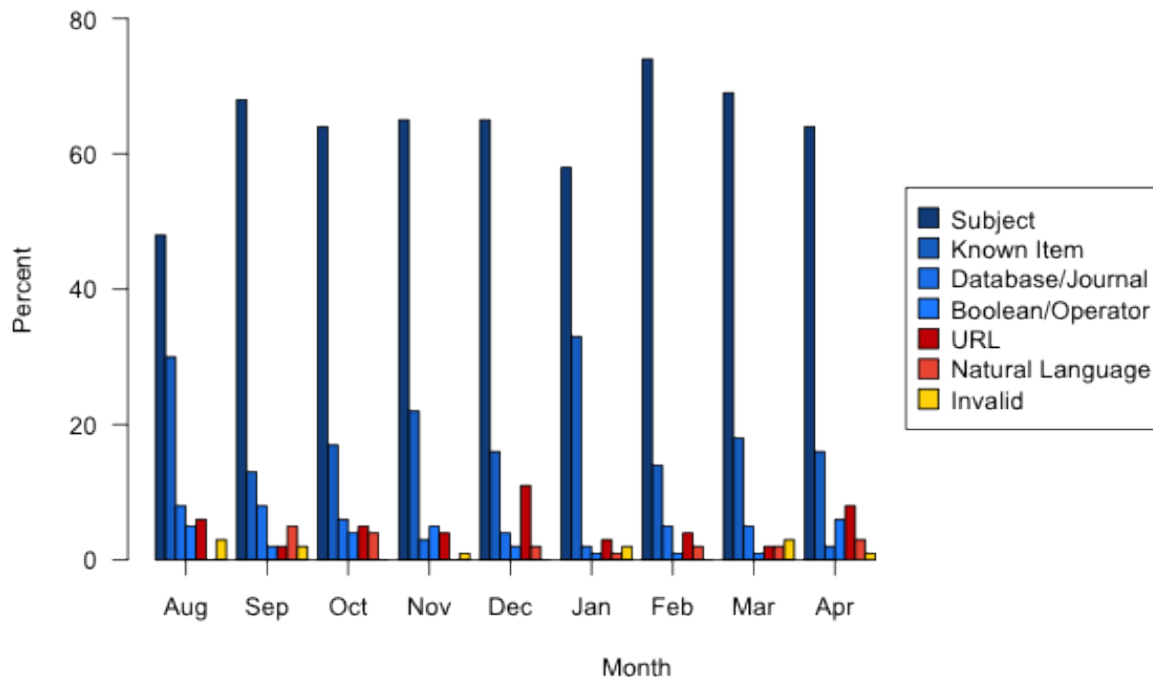
The high-quality search percentage is similar between semesters with no significant change (p value = 0.536, from a two-sample t-test). The mean for high-quality queries in semester 1 is 91%, and the mean for high-quality queries in semester 2 is

92.25%. The standard deviation for semesters 1 and 2 is 2.83 and 2.87, respectively.

There was a significant difference between high-

and low-quality queries over the entire academic year (p < 0.001, from a two-sample t-test). The mean of the high-quality queries was 91.56%, and mean of the low-quality queries was 8.44%.

The results put the hypotheses of the original study into perspective. The first hypothesis of this study was that the quality of queries within Summon is low. This hypothesis is rejected given the large and consistent differences between high- and low-quality queries. The second hypothesis of this study was that the quality of queries in

Summon improved during the first two semesters of implementation. This hypothesis is also rejected since there was no significant difference in high-quality searches between the two semesters.

Figure 2 shows the percentage of each query type by month from August 2011 through May 2012, with high-quality categories as cool colors and low-quality as warm colors. The results are similar to the previous year with subjects as the most common query type followed by known items.

Figure 2: Frequency of query types by month for August 2011 – May 2012. Low-quality queries are displayed in warm colors. High-quality queries are displayed in cool colors.
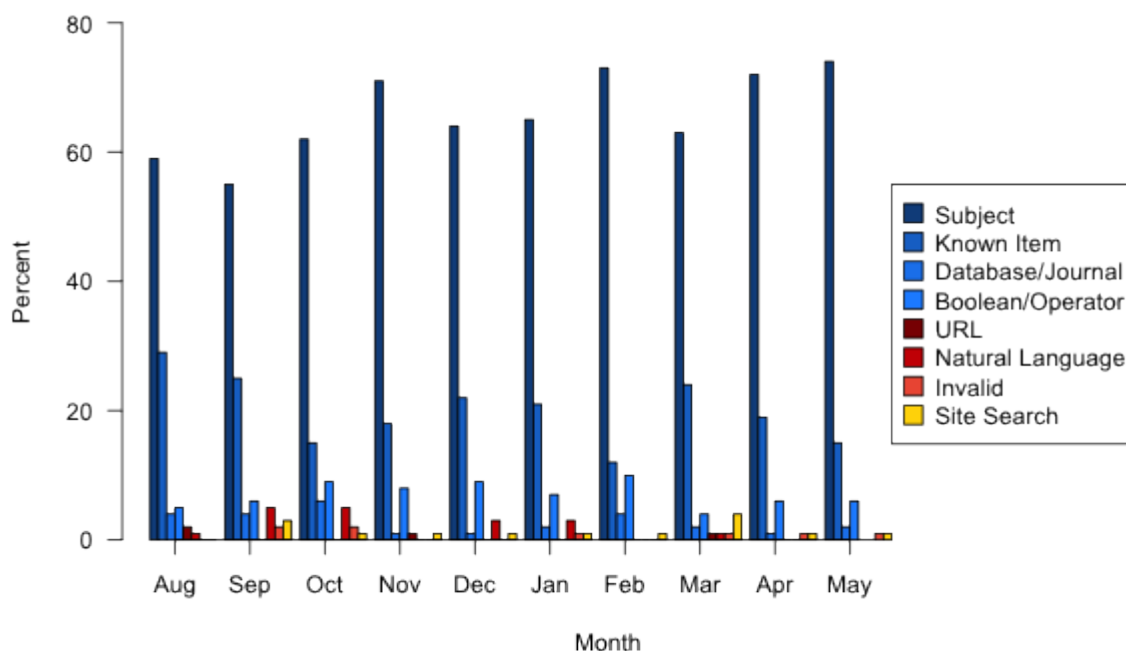

Figure 3 shows the percentage of each query type for each year of the study and the percent difference of each query type by year. There were two query types with a statistically significant difference from the first year to the second year. Boolean operator queries increased a significant amount (P<0.001) in year two, although this is

caused by the inclusion of special characters in the second year of the study. The difference is not significant if the queries utilizing special characters are removed from the dataset. URL queries decreased a significant amount (P = 0.0015) in year two.

Figure 3: Percentage of each query type for each year of the study and the percent difference of each query type by year
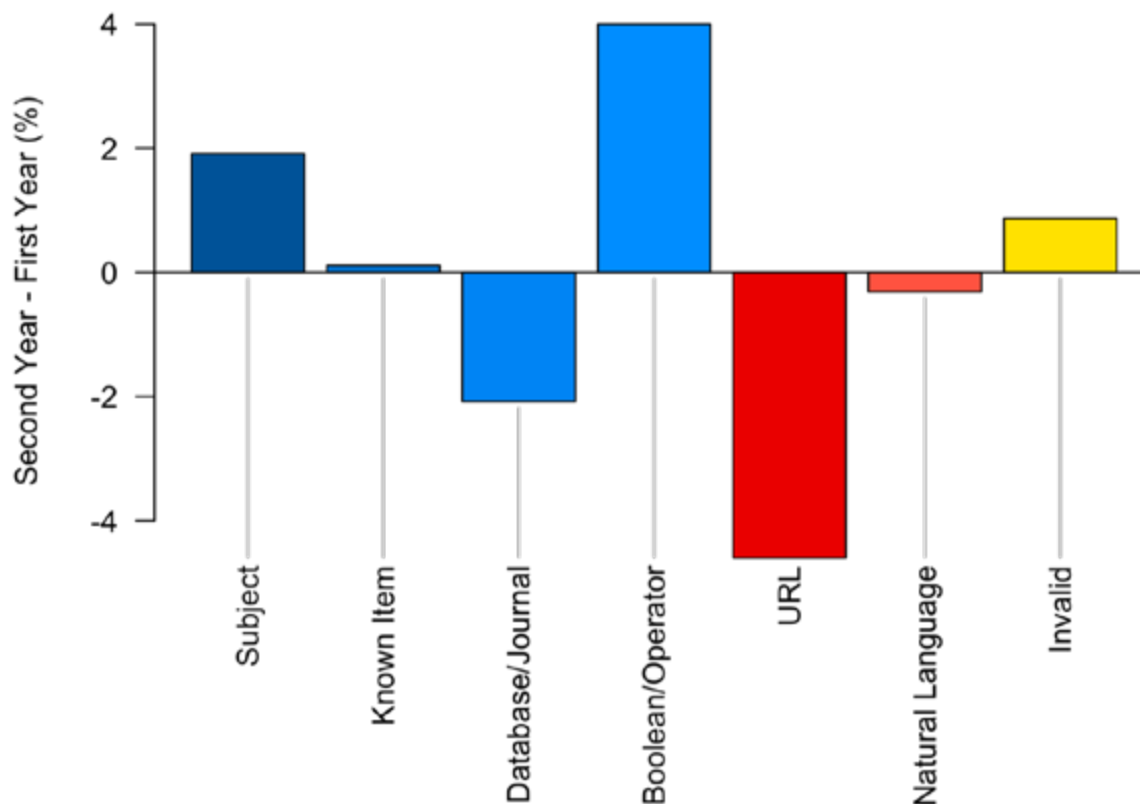
Discussion:

The Concept of Quality

These results could signify that the single search box model of discovery via Summon is sufficient for most students. However, there are some who do not understand the most effective ways to search, and others who do not even understand the meaning of searching library resources. The fact that most search queries are legitimate could mean that library resources are becoming easier to access through Summon.

There are some implications for further instruction when only the high quality queries are considered. Table 3 shows the percentage of query types for the high quality queries performed in each year. Since the use of Boolean operators and other operators yields more effective results, a higher percentage of this query type would be an indication of effective searching behavior. Even with special characters added in year two, Boolean/operator queries are still a small percentage of the total queries.

Table 3: Percentage of query types for high-quality queries performed each year

Query Type	First Year (%)	Second Year (%)	Overall (%)
Subject	69.8	68.9	69.3
Known Item	21.7	20.9	21.3
Database/Journal	5.2	2.8	3.9
Boolean/Operator	3.3	7.3	5.4

Unfortunately, within the limitations of this study, it is impossible to determine whether or not the users considered their search sessions to be successful. Overall, the concept of quality is

relative, and the ability to form an effective search query appears to be out of reach for many users of the system.

Facebook.com

The impetus for this study was that *facebook. com* was the most common query, after a blank, entered into Summon. According to the Summon Administration Console, the last time this search was performed was November 2011. Thirty-five searches for *facebook* without the *.com* appear between December 2011 and April 2012, but users may be doing research about Facebook or looking for the site itself. Other URL searches continue to show up in the data, but the numbers have dropped significantly since the first year of the study. This is an encouraging sign that users better understand what the Summon search box does not do.

Conclusion:

This study is merely a window into a better understanding of how Summon is being used at MSU Library. The data from the administration console shows that Summon is being heavily used. Although it is unlikely that the majority of queries entered into Summon are of high quality from the perspective of skilled information seekers, it is apparent that the majority of queries are valid. Further instruction on how to form an effective query may be necessary in order to distill the information library users need from the volume of resources available via discovery layers. The various query types that emerged from the study may leave room for a response from an interface design perspective as well. For example, the amount of site search queries performed within Summon may necessitate a site search feature in the Summon interface. This study merely scratches the surface of what can be accomplished in understanding Summon and how users are interacting with it. By better understanding user behavior, this study comprises a piece of the puzzle in evaluating the effectiveness of Summon and informs future discussions of Summon's role at MSU Library.

Google Analytics Transaction Log Analysis

Purpose:

In addition to the data provided through the Summon Administration Console, data provided from external web analytics tools can provide different insights into the system. Web analytics has emerged in recent years as a valuable assessment tool in understanding online user behavior. By providing insight into how users navigate a library website, the study of web analytics can lead to more effective delivery of web content and web services.

In February of 2012, MSU Library activated Google Analytics (GA) user tracking within Summon. With GA in place, valuable user and site performance data became available to us. A full account of GA capabilities and a robust analysis of GA data are beyond the scope of this study, but a focused tour and examination of data provided by GA will prove instructive in demonstrating the value of web analytics in understanding user behavior within a WSD.

Methods:

Among the full suite of metrics offered through GA, two particular metrics that are valuable for understanding Summon usage were identified: landing pages and total page views. The landing pages metric provides data showing which page within a website a user begins navigation. To aid understanding of this metric, an independent GA account was established for the Summon web server, which allowed control for a user's navigation within Summon as distinct from MSU Library's primary website. This distinction is important for allowing one to see which percentage of library home page visitors initiated a Summon search and then landed into the Summon web server.

Results and Discussion:

Comparing total page views of MSU Library website's homepage with the total landing page visits into the Summon server provides an indication of the overall number of Summon searches that originated from the MSU Library home page. Total page views for the MSU Library home page for the sixth-month period February 13, 2012–August 13, 2012, numbered 198,447 (Figure 4). Total landing page visits for MSU Library Summon searches for the same period numbered 21,838 (Figure 5). Of these 21,838 landing page visits, GA referral data shows us that 20,905 of those visits originated from the MSU Library home page (Figure 6). This data indicates that 9.5% of users visiting the MSU Library's home page have

initiated a search using Summon.

Figure 4: Total page views for MSU Library Summon searches in six months, 2012.
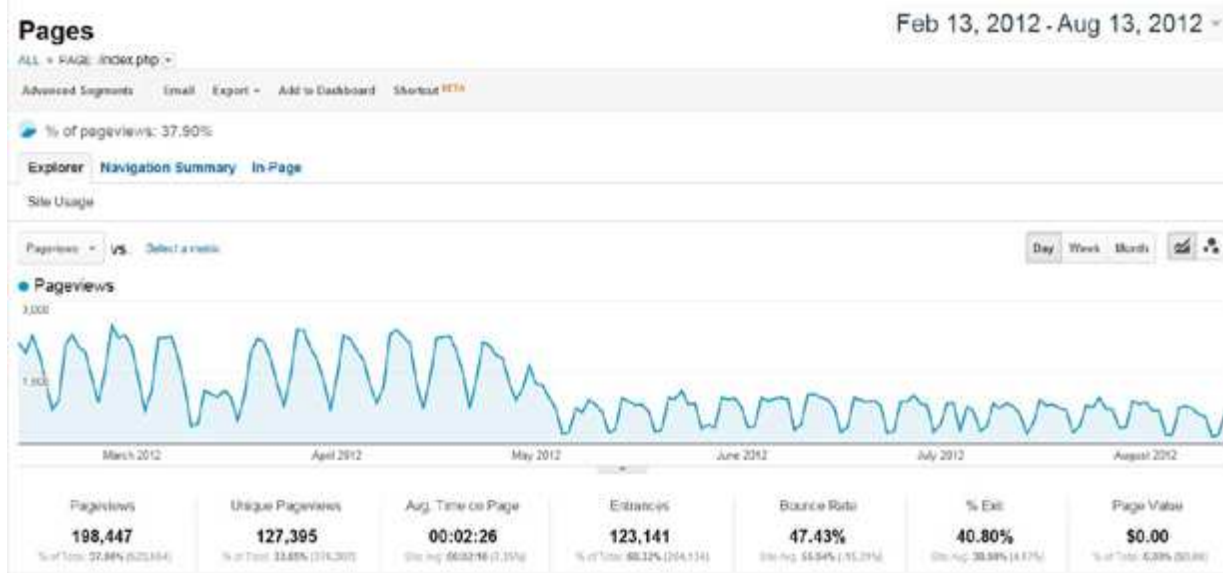

Figure 5: Total landing page visits for MSU Library landing page in six months, 2012.
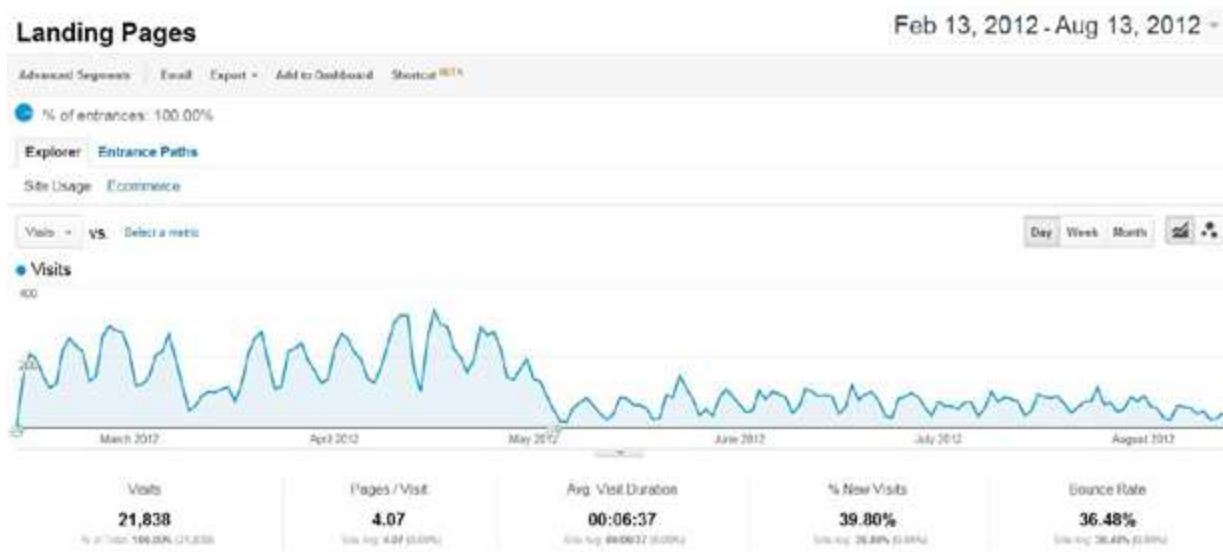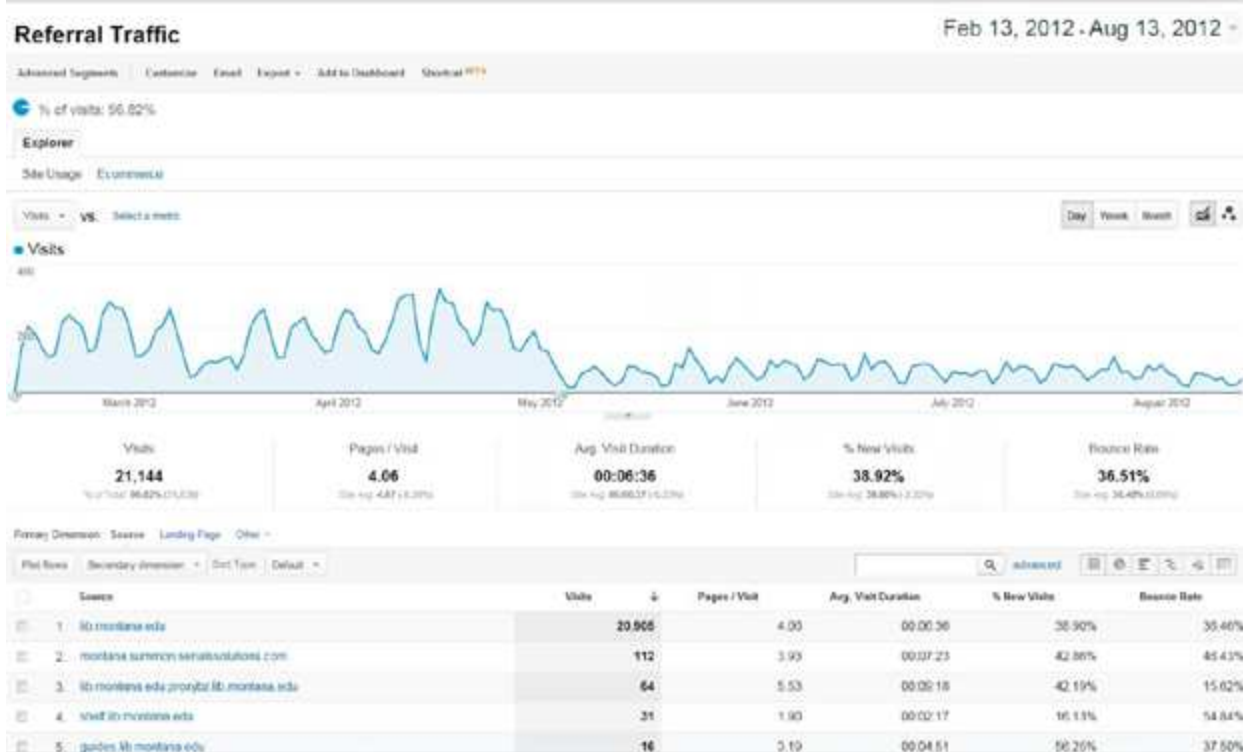
Figure 6: Total referral pages for MSU Library landing page in six months, 2012.

Providing this figure with additional context allows one to more fully understand it. One can look at additional data with GA to help us interpret not only how many users are engaging Summon, but also how they are using Summon. Looking at new versus returning visitors allows one to see how many users are coming to Summon for the first time and how many are returning to the tool after having used it before (Figure 7). The evaluation over time of this metric combined with landing page visits may indicate user satisfaction with Summon.

Figure 7: New vs. returning visitors to Summon
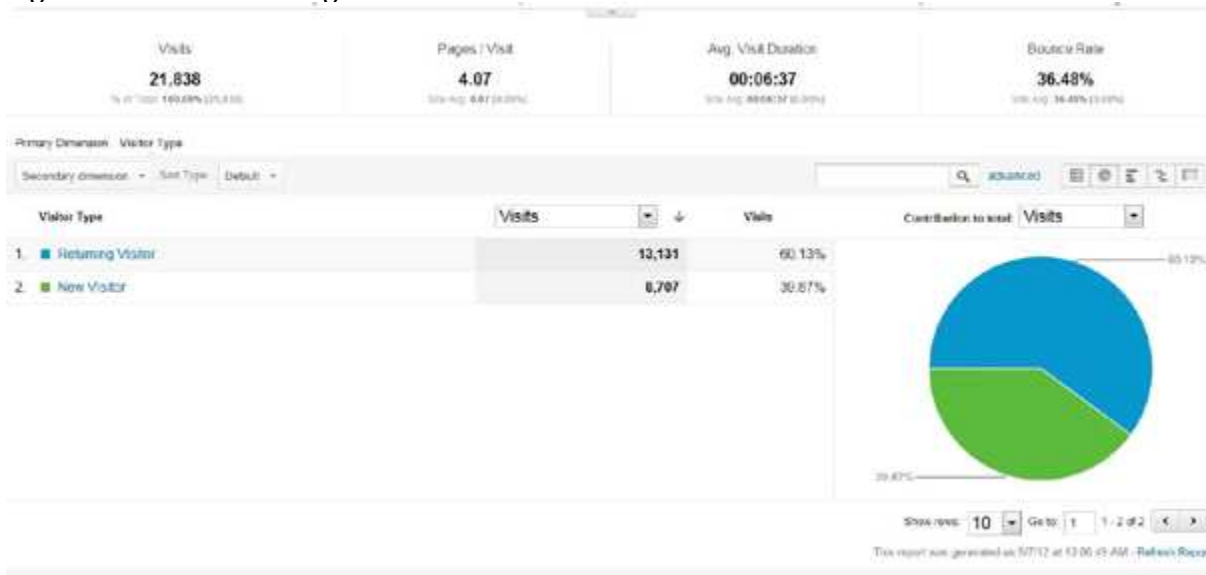
Further insight into user behavior within Summon is available through a second important metric, total page views (Figure 8). This metric allows one to see the total number of searches, search queries, and the number and type of facets that users selected in conducting searches. For the sixth-month period February 13, 2012–August 13, 2012, GA recorded 72,787 searches. This large dataset can be mined from within GA by using the included search tool. The methodology is still being developed for analyzing these results, but already shows promise in charting user capability in conducting searches. For example, if one identifies faceting as an important factor in successful Summon searches, one can examine, over time, the relative use of faceting by users to determine if the functions of Summon are being employed to their full advantage. Likewise, one can examine the full range of search queries recorded through GA to understand the information-seeking behavior of users and consequently tailor and refine the Summon tool.

Figure 8: Total page views for Summon at MSU Library for six month period
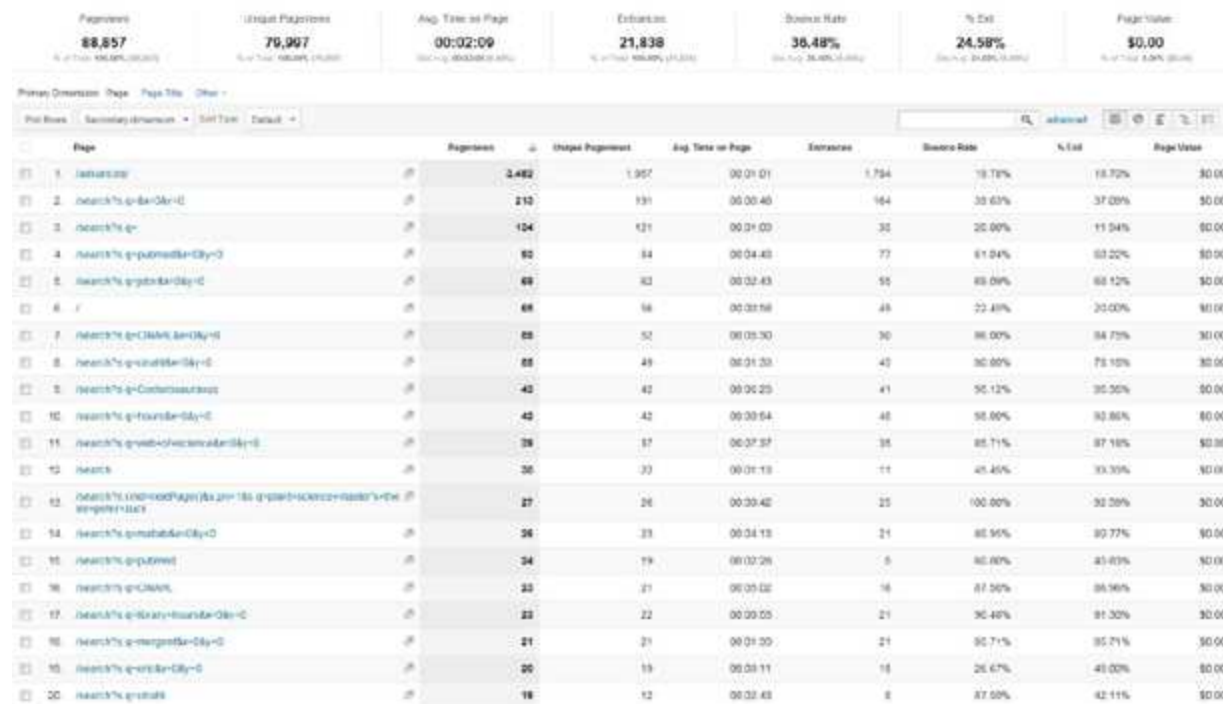
It is crucial to state at this point the importance of establishing comparative benchmarks for the interpretation of GA data. Benchmarks can be set over time and used to compare similar ranges of data as changes are made in web design and service delivery. In August 2012, for instance, MSU Library launched a redesigned and restructured home page that more prominently features the Summon search box. After collecting GA data for another six-month period following the launch of the new design, a comparison of that data to the previous six-month period may provide user behavior insight into the connection between library web design and the use of a web-scale discovery tool such as Summon.

Equally important for the evaluation of GA data is establishing goals, which are necessary for providing contextual information. A library should identify, for example, an acceptable ratio of home page visits to Summon visits. If GA indicates that 9.5% of home page visits result in Summon searches, the challenge for libraries is in determining the value of that figure. A new home page design may initially increase that number, but over time stagnating or decreasing Summon usage may indicate that users are unhappy with the tool.

Conclusion:
While GA can provide extensive user behavior data, the challenge of interpreting that data reveals the limits of web analytics. Tools like GA

481

are successful in providing vast data related to where users go within a website and how they move around. Web analytics, in essence, records the choices a user makes, but it cannot tell us why a user makes those choices. Acknowledging this limitation is an important factor in fully understanding not only Summon user behavior but also library website user behavior. When appropriately contextualized, data gathered from web analytics can be a valuable component in the overall assessment of user behavior and WSD tools.

Holdings and Indexing Comparison
Purpose:
As part of the study, the coverage of resources was evaluated in the WSD in comparison to library holdings. This information informs several decision points. First, if something is not indexed in Summon, should MSU Library cancel that subscription or find an alternative source, or encourage Serials Solutions and the providing vendor to get that item indexed in Summon? Also, if there are a large percentage of items not indexed in Summon, should MSU Library explore other WSD tools to see if there is better coverage?

Methods:
In August 2012, MSU Library requested that Serials Solutions provide a title-level analysis of its holdings against the Summon index. This is a free service offered upon request. MSU Library provided a list of ISSNs to Serials Solutions, which was generated from the resources that it had set as active in its instance of 360 Resource Manager. Serials Solutions then took the unique ISSNs from this list and de-duplicated titles with multiple ISSNs to provide the information requested. In addition, the MSU Library took its list of subscribed databases and compared them against Serials Solutions' list of indexes with full-text content that are indexed in Summon.[9] Of the items in the database list, ones that are abstracting and indexing services only were determined by consulting the Ulrichsweb database.[10]

Results:
The MSU Library subscribes to a total of 139 databases. Of these, 59 are abstracting and indexing databases, so they are not included in the Summon indexing. Of the remaining 80 indexes for consideration, 65 are indexed in Summon. Three more databases that are not in Summon's partner

list for indexing have MARC records for the items that are included in the Library's catalog and are, therefore, indexed in MSU Library's Summon instance. In total, 85 percent of MSU Library's full-text subscription databases are indexed in Summon. Fifteen percent, or 12 titles, are not indexed in Summon.

As was suggested by a representative at Serials Solutions, a more complete picture of Summon's coverage is gained by understanding how many individual serial titles are indexed in Summon. The file provided to Serials Solutions included 79,757 entries. After Serials Solutions de-duplicated titles from this list, there were 42,464 unique titles to analyze. Of the titles with active ISSNs, only 2,679 items are not yet covered in Summon. And of these titles not yet in Summon, only 709 titles are peer-reviewed sources. According to the report from Serials Solutions "we can say that we are already in active negotiations with some, if not most, of the content sources on this list." So, 6.3 percent of the library's titles are not indexed in Summon, and only 1.6 percent of peer-reviewed titles are not indexed.

Discussion:
Returning to the questions posed prior to analyzing the amount of indexed content in Summon, if something is not indexed in Summon, should MSU Library cancel that subscription or find an alternative source, or encourage Serials Solutions and the providing vendor to get that item indexed in Summon? Also, if there are a large percentage of items not indexed in Summon, should other WSD tools be explored to see if there is better coverage? The study's findings show that the MSU Library has a relatively small set of items that fall into the not-indexed category, with just 12 databases and several hundred journals. This finding suggests that it is a manageable task to encourage the vendors of those databases and journals to include their content in Summon and to look at alternate sources for this information, when possible.

Conclusion:
Libraries can conduct similar studies before committing to a WSD product to make sure that indexing rates of local holdings are acceptable to a library. Also, this information helps inform collection development decisions if indexing in a WSD is important to the library for discoverability.

Libraries may want to encourage publishers and WSD providers to partner whenever an item is identified as not indexed in the WSD. If these groups know that there could be financial implications for not cooperating, it can be a motivating force to getting the information included.

Study-wide conclusions:

Reviewing the information gleaned in this study, libraries have a great deal of control in shaping user experiences in their interactions with WSD tools. Libraries can make sure their collections have coverage in their WSD by serving as an advocate with WSD vendors and content providers. They can work with these groups to encourage cooperation in having resources indexed in WSDs and in making sure that OpenURL standards are followed. Likewise, information about the purpose of WSDs and their placement and usage on a library website can be presented whenever possible. Library website designers can use tools like Google Analytics to understand user behavior and design website presentation of such tools accordingly. Likewise, librarians interacting directly with users can shape their instruction skills and technologies to inform usage of these tools most productively. While a drop in failed link rates from 43 percent to 27 percent over several years is an improvement, a quarter of searches still failing would likely be considered by most to be unacceptable or highly intolerable. The more information libraries can provide to the WSD producers and the content providers about what is workable within the library arena, the more powerful libraries can be in shaping the tools that are now so important to the users being served. The more informed libraries are in identifying areas of improvement and education with WSD, the better libraries can be in making WSDs tools for actual discovery.

Notes:

1. Jason Vaughn, "Investigations into Library Web Scale Discovery Services," *Information Technology and Libraries* Preprint (January 2011), http://digitalscholarship.unlv.edu /lib_articles/44.

2. Helen Timpson and Gemma Sansom, "A Student Perspective on e-Resource Discovery: Has the Google Factor Changed Publisher Platform Searching Forever?" *The Serials Librarian* (August 2011): 253–256; and Cory Lown, Tito Sierra, and Josh Boyer, "How Users Search the Library from a Single Search Box," *College and Research Libraries* Preprint (December 2011), http://crl.acrl.org/content /early/2012/01/09/crl-321.full.pdf+html.

3. Dartmouth College Library, "An Evaluation of Serials Solutions Summon: As a Discovery Service for the Dartmouth College Library" (November 2009), http://www.dartmouth .edu/~library/admin/docs/Summon_Report .pdf.

4. David Howard and Constance Wiebrands, *Culture Shock: Librarians' Response to Web Scale Search* (February 2011), http://ro.ecu .edu. au/ecuworks/6206/; and Stefanie Buck and Margaret Mellinger, "The Impact of Serials Solutions Summon on Information Literacy Instruction: Librarian Perceptions," *Internet Reference Services Quarterly* (December 2011): 159–181, http://dx.doi.org/10.1080/10875301.20 11.621864.

5. Melissa Becher and Kari Schmidt, "Taking Discovery Systems for a Test Drive," *Journal of Web Librarianship* (2011): 199–219; and Julia Gross and Lutie Sheridan, "Web Scale Discovery: The User Experience," *New Library World* (2011): 236–247.

6. Kelly Meadow and James Meadow, "Search Query Quality and Web-Scale Discovery: A Qualitative and Quantitative Analysis," *College & Undergraduate Libraries* (August 2012): 163–175.

7. R Development Core Team, *R: A language and environment for statistical computing* (2010), http://www.R-project.org.

8. For detailed information about the coding process see Meadow and Meadow, "Search Query Quality," 163–175.

9. Serials Solutions, Summon Key Databases and *Packages—Full-Text* (July 2012), http://www .serialssolutions.com/assets/resources

/Summon_Databases_Full_Text.pdf.

10. Serials Solutons, *Ulrichsweb* (2012), http://www
.ulrichsweb.serialssolutions.com.