

A multi-campus usability testing study of the new Primo interface

W. Jacobs, Mike DeMars & J. M. Kimmitt

To cite this article: W. Jacobs, Mike DeMars & J. M. Kimmitt (2020) A multi-campus usability testing study of the new Primo interface, *College & Undergraduate Libraries*, 27:1, 1-16, DOI: [10.1080/10691316.2019.1695161](https://doi.org/10.1080/10691316.2019.1695161)

To link to this article: <https://doi.org/10.1080/10691316.2019.1695161>



Published online: 09 Dec 2019.



Submit your article to this journal [↗](#)



Article views: 382



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)



A multi-campus usability testing study of the new Primo interface

W. Jacobs^a, Mike DeMars^b, and J. M. Kimmitt^c

^aJ. Paul Leonard Library, San Francisco State University, San Francisco, California, USA; ^bPollak Library, California State University Fullerton, Fullerton, California, USA; ^cUniversity Library, California State University Dominguez Hills, Carson, California, USA

ABSTRACT

When the 23-member California State University Libraries merged integrated library systems in 2017, a unique opportunity arose to conduct multi-tenant usability testing on user experience in the Primo discovery service, collectively branded OneSearch. In 2018, librarians at five campuses conducted a round of task-based individual user testing on their separate instances, seeking to review individual campuses' decisions about customizations designed to provide optimal performance of the system and provide insight into group decisions about our shared customization choices. Along with the findings of the test results and recommendations for enhancements to existing modifications, we discuss the challenges and opportunities that are present in conducting multi-campus usability testing and share recommendations for future consortium-based testing efforts.

ARTICLE HISTORY

Received 15 October 2019
Revised 14 November 2019
Accepted 15 November 2019

KEYWORDS

Consortium; CSU; discovery system; multi-campus; Primo; usability testing

Introduction

The California State University (CSU) Libraries decided to migrate from a series of disparate library systems to a Unified Library Management System. A request for proposals was drafted in 2014 and the contract was awarded to Ex Libris and their product, Alma. Along with the move to Alma, the CSU libraries elected to use Ex Libris' Primo as the discovery system for the 23 campus libraries. The decision to use Primo was driven in part by the large amount of customization options that Primo provides. Primo allows administrators to customize much of the look, feel, and functionality of the system, including relevancy rankings.

The CSU Chancellor's Office (CO) served as the main organizer of this transition. The CSU Council of Library Deans (COLD) formed several working groups to help with the migration project, one of which focused exclusively on discovery. The Discovery Migration and Implementation

CONTACT Joanna Messer Kimmitt  jkimmitt@csudh.edu  University Library, California State University Dominguez Hills, 1000 E. Victoria St. Carson, CA 90737, USA.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/wcul.

© 2019 The Author(s). Published with license by Taylor & Francis Group, LLC

Group did some user testing prior to implementation, work which laid the foundation for the Discovery Working Group formed after going live in June 2017. The Discovery group formed task forces to look into relevancy ranking, third party integration, user outreach, and user experience (UX). The UX group was tasked with making modifications to the user interface of stock Primo to improve and enhance the user experience. The group conducted usability studies and environmental scans to develop a list of best practices, with the goal of creating a unified interface design and discovery experience across the 23 campus CSU system based on the most effective customizations (California State University Libraries 2017).

There is extensive documentation in the literature of usability testing for many types of library discovery interfaces, mostly focused upon implementation of specific features and customization within a single-campus implementation. Multiple studies also include sentiment analysis of user perceptions of the discovery layer. Among the best practices established in user experience testing are studies conducted within the last ten years on Primo, the discovery layer associated with Ex Libris's Alma library management system. Early adopters of the Primo discovery system, such as the group at Flinders University, combined a survey instrument with a lengthy set of tasks using traditional usability testing methods (Jarrett 2012). These small testing cohorts evaluated system settings for the user interface (UI) in multiple iterations of system settings prior to their official go-live, allowing the designer to fine-tune the implementation before wide release. At the University of Kansas, researchers conducted two rounds of usability testing on specific tasks in 2012 and 2013, as well as mining Google Analytics and search logs for aggregated user behavior over the course of three academic semesters to support their qualitative observations (Hanrath and Kottman 2015). They found that most users tended to access facets for search limiting when conducting open-ended research, and were less likely to use those limiters when searching for a specific or "known" item. In his foundational work, Comeaux (2012) studied the first Primo UI post-implementation with five scenarios designed to glean information about where a typical user would encounter roadblocks to making best use of the interface's capabilities in search and information sharing, like citation, email, and placing holds and requests for items. A team at the University of Vermont (Nichols et al. 2014) conducted usability testing on their implementation of the first Primo user interface with similar research questions in mind, and with a more traditional usability testing setting: a thirty-minute session, also with five discrete tasks that included several subdivisions. Each of these studies determined that users had difficulty with accessing hold and request services, location and facet/limiter labeling, and the use of facets for focusing searches. Even when users had issues with accessing information or

completing more complex tasks as assigned, they found it easy to use, rating high satisfaction overall in findings from multiple studies (Nichols et al. 2014; Comeaux 2012; Jarrett 2012; Hanrath and Kottman 2015).

Other groups tested the new Primo UI after going live with the first version of Primo, examining specific customizations and measuring the adjustments in information-seeking behavior that users made from the first discovery interface to the second implementation. [Galbreath, Johnson, and Hvizdak \(2018\)](#) conducted usability testing with eight users, covering a large task list. They found that users were able to easily complete tasks that were similar in workflow between the two interfaces, but where there were significant design changes between the classic and the new UI, they tended to have more difficulty.

Other studies surfaced the problems users face with selecting options like tabs and facets that available in a discovery layer. Some of the issues may stem from a lack of understanding of library jargon and the definition of “scholarly” or “peer-reviewed.” A team at the University of Houston conducting testing of Primo came to this conclusion, among others, when observing undergraduate students attempting to search for peer-reviewed articles, and instead clicking on the “Reviews” facet, limiting to reviews of books rather than journal articles (Brett, Lierman, and Turner 2016). Valentine and West (2016) sought feedback from student assistants via usability testing and simplified language in their Primo instance to reflect their needs. They further recommend that teaching and learning librarians be involved in usability testing to gain understanding about how to provide research and instruction support to students and faculty who are using the system for resource discovery, stressing that usability testing is an iterative and localized process, even as more large organizations are implementing discovery systems in cohorts.

In recent years, larger consortia have implemented Alma and Primo as a group effort, making use of the resource sharing capabilities of the system and the opportunity to build user groups and, since August 2016, to share centralized Primo customization packages within these existing governance structures. As outlined by Moore and Mealey (2016), the Orbis Cascade Alliance created a consortial working group to investigate implementation of a “central package,” particularly to benefit member institutions that have less technical capacity and resources available for development, and ultimately recommended constructing a central package for their membership to choose to implement as they wished. This created a similar structure for searching experience across their consortium, with customizations floated on top to reflect campus web standards and individual libraries’ interface choices.

Following publication of the Orbis Cascade study, the CSU decided to implement a central package for similar reasons. In 2016, the CSU’s

Discovery Working Group conducted initial usability studies prior to implementation, as discussed below (DeMars 2017).

The CSU system went live with Primo and Alma in June of 2017. Prior to go-live, Ex Libris provided the 23 libraries in the system with a test environment in June of 2016, giving the CSUs a full calendar year to adapt to the new systems and make requisite changes to workflows in various library departments. It was during this time that the Discovery Migration and Implementation Team made several decisions on how Primo would be implemented on the 23 campuses. Based upon recommendations from this team, COLD mandated that a unified look and feel should be applied to the interface. The CSU rebranded Primo as OneSearch, adopting a consistent logo that could be customized by each campus to match their school colors. The customizability of Primo meant that the CSU libraries could still apply local changes, but overall the look-and-feel of the interface would remain fairly consistent across the system.

In addition to branding and stylistic recommendations, the Discovery Migration and Implementation team performed two rounds of user testing on the new Primo interface during the test environment phase. Four CSUs participated in this study, Fullerton, Los Angeles, Fresno, and Northridge, each of which had applied their own modifications to Primo. Each campus selected testers to be in line with the user population, primarily undergraduates with some graduates and faculty. All of the campuses made an effort to ensure that the groups studied would be similar across the four campuses. Each tester was given five tasks to complete, ranging from finding a known item, to finding a peer-reviewed article on a given topic. The testers evaluated the user's ability to complete each task. Time to completion was also recorded. Once the data was gathered and analyzed, the most successful aspects of the various layouts and configurations were applied at the Fullerton campus instance. Once the changes were applied, a second round of testing was performed on the Fullerton instance. The changes to the interface and search scopes resulted in greater user satisfaction and lead to improvements in both discoverability and time to completion.

Based upon the results from these initial studies, the Discovery Team developed a set of implementation guidelines as well as a best practices document that was shared across the CSU consortium and was posted on the CSU Unified Library Management System Wiki page. Additionally, these changes were later incorporated into the Primo Central Package, a collection of javascript and configuration settings that can be applied to Primo to manipulate the user interface. To inherit changes from the Central Package, individual CSUs only needed to opt in, which the majority of campuses electing to do so. This ensured that the look and feel across many of the campuses is similar; however, local administrators can choose

to override some changes from the central package. Customizations applied locally were often implemented to streamline the interface, reduce visual clutter, and enhance discoverability. Default search scopes also varied from campus to campus. These small modifications meant that upon go-live there were still many variations from campus to campus.

Methods

After going live with Primo, the Implementation Team was disbanded, making way for the Discovery Working Group. This working group formed a number of small task forces, each charged with investigating or enhancing the performance of Primo. **One of these task forces, the UX group, was tasked with building off the user studies that were mentioned above.**

The UX Group first decided that it would be necessary to establish just how many modifications to the central package were being applied to local instances. The first step in the process was to send out a survey to the 23 Primo administrators. The short 9 question survey asked administrators to detail the CSS and Javascript changes they had made to their system, as well as questions about customizing labels and facet changes. Additionally, members of the group conducted an environmental scan to see what changes were made to search scopes and default search values. This data, which was also published on the Chancellor's Office Wiki, showed that there was still a good deal of variation among the campuses, despite the attempt to standardize the experience through the central package. In most cases, these customizations were done on top of those provided by the CO. Given these variations, the UX Task Force decided more user testing was in order to determine which layouts best enhanced the discovery experience.

The UX Task Force identified five campuses in the system that would be ideal for user testing, as they supplied the greatest variances in the user interface. The Fullerton, Dominguez Hills, San Marcos, San Francisco, and Sonoma campuses all agreed to participate in the study and a **test administrator was identified at each location.** For this round of testing, the group decided to rethink the UX testing process, as the first round of user testing proved quite instructive and showed some issues with our testing procedures. For the first round of testing, the five main tasks assigned were identical, however each campus was afforded the opportunity to come up with additional questions more germane to their campus. **This meant that tests at some locations had 12 tasks while others had between 5 and 10.** This severely impacted the the time it took to complete the test. At some campuses the test took 15 minutes for users to complete, at others it took users two to three times as long. The additional optional questions made it

difficult to tell if extended time to completion was due to layout issues, or simply the length of the test. Additionally, each campus recorded the results differently; some using paper and others using Google Forms. This made compiling the data unnecessarily difficult. The Discovery Working Group determined that the next round of testing would be more streamlined, efficient, and identically applied across all of the participating campuses.

The Discovery Working Group identified key areas of the interface that they wanted the UX Task Force to investigate. Many of the questions, or tasks, that were devised were similar to those on the first and second round of testing, including tasks such as finding a known item and identifying a peer-reviewed article. Additionally, the CSU had recently established a CSU-first resource sharing platform called CSU+, and the group wanted to test the users' ability to navigate this system. Other questions related to layout and wording changes were also included. For this round, campuses were instructed to stick with the agreed upon 10 tasks and three demographic questions, which were included to insure that we had a test group that matched the general makeup of typical library users. No additional questions would be included as this would impact time to completion.

The participating campuses also agreed on using a common recording method for all test results. The group created a single Google Form for all of the campuses to use, which would simplify data analysis. This differed greatly from the first round in which some responses were recorded on paper and others on a separate Google Form. Despite this single online input form, some campuses still elected to use paper to record the results. Those campuses then added the information they recorded to the Google Form post-test.

Test administrators used different methods to record what steps the test subjects took to complete each task. Many of the questions on the UX test were yes or no answers, i.e., "Can you find a copy of the book 'Animal farm?'" The researchers were interested not only in users' abilities to complete each task, but also in what steps users took to complete each task and at what points they were having difficulties. One of the test administrators used the screen recording software Camtasia to capture the entirety of the user's testing experience. This allowed the administrator to leave the user on their own and exit the room in an effort to reduce the subtle pressure that having the test administrator in the room might generate. These videos were reviewed later during the data analysis phase of the project. The other four campuses in the study took a more traditional approach and elected to take notes during the test which they then entered on the Google Form. This data was also reviewed in the analysis phase.

Which of the following best describes you?

39 responses

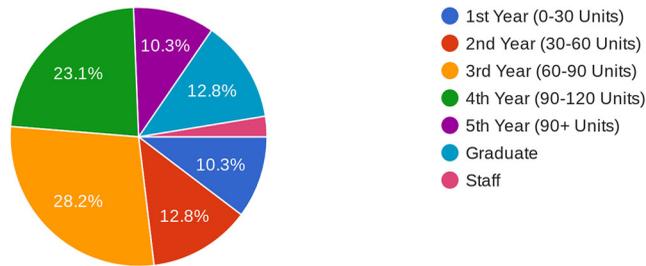


Figure 1. User Identification.

Results

The test administrators at each campus made an effort to create a testing group that reflected our user base. This included a majority of students that were in their 3rd year or greater (Figure 1). Two of the campuses also elected to include Faculty members in the test group. The test also asked users to indicate how often they use the library in an effort to identify highly active and potentially more savvy users. Of the 40 users tested, 15% of respondents indicated that they used the library daily, and 62% indicated that they use the library website weekly, monthly, or “a few of times a semester.” Roughly 22% indicated that they rarely or never used the library website.

The first task users were asked to complete involved a known item search. Users were asked to find a copy of the book *Animal Farm*, a title which each campus owns. This task was easy for most to complete as the result was usually first or second on the results list. Over 92% of users were able to locate the book. It is difficult to determine why three users were not able to locate the book. One campus had two people who were unable to find the item, perhaps because it was listed second after the film version. In the other case it is difficult to understand why the user reported they could not find the item, because at that campus the book was the first result.

We also asked users to find a specific journal, *Criminal Justice Ethics*. Again, there was a fairly good success rate here, with over 82% of testers across the five campuses able to identify the title (Figure 2). There was less success at the San Francisco campus, despite the fact that the journal does come up first on the results list. However, we found that many users limited their search to articles, and searched futilely for “Criminal Justice Ethics” after having removed the journal from the results list. Results where the students found an article with the words “criminal justice ethics” in the

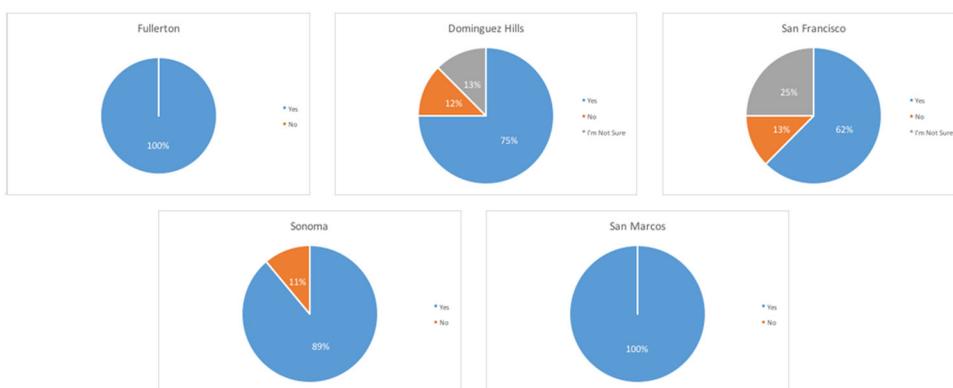


Figure 2. Can you find the journal titled Criminal Justice Ethics?

title that happened to be in that journal were recorded as “I’m Not Sure” as the students were unclear on how they had arrived on the journal’s home page. Confusion about the difference between the terms “article” and “journal,” which is not uncommon, appears to be the reason that users employed the articles limiter.

In one task, users were asked to find a peer-reviewed article on the Pacific Ocean Garbage patch. This question was intended to test user’s ability to recognize the default peer-reviewed icon created by Ex Libris. In addition, the test administrators were curious to see if users employed the Peer-Review facet. Results across the five campuses were very similar, with 75% of users in aggregate able to correctly identify a peer-reviewed article (Figure 3). This task was more difficult than it may appear, because there is a trove of articles from other print sources that are not subject to peer-review. In the case of San Francisco, the first peer-reviewed article did not appear until the 18th record, and at Dominguez Hills it is the 12th record. Both of these campuses had a 10 result limit, meaning that users had to click Load More Results to find an item meeting this criteria. This is impressive and shows the resilience of determined users. Employing the Peer-Review facet would have added these users. At San Francisco, several users looked for peer review under format. **This indicates that the usage of the word “format” to solely indicate the physical form of an item (often theoretical for digitized items) is not clear to users.** The high success at San Marcos is due to the fact that they display the full Primo Central Index (PCI) by default. When that setting is applied, the very first article is peer reviewed. At the time, the other campuses in the study displayed only the collections that the library holds, not the full PCI. Representatives from other campuses indicated that they made this choice for fear of overwhelming their interlibrary loan departments with requests.

This question also exposes a critical need in information literacy instruction. There are still a significant number of users who struggle with

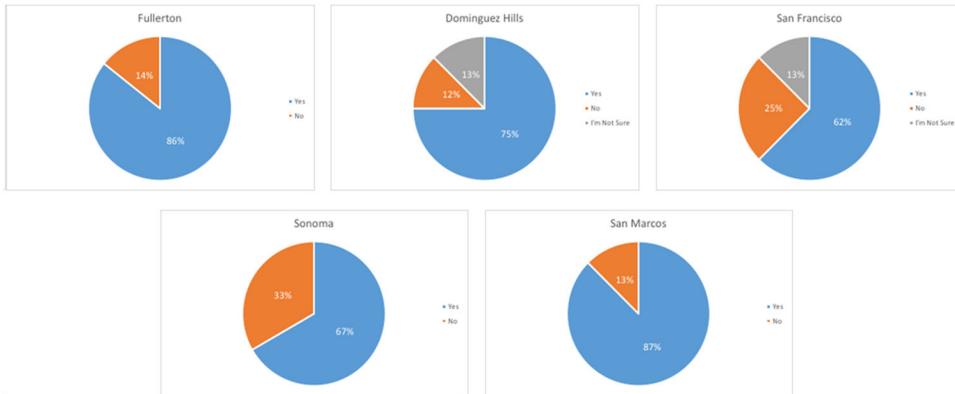


Figure 3. Can you find a peer reviewed journal article on the Pacific Ocean garbage patch?

identifying peer-reviewed material. This illustrates the need to engage with users about what the term “peer review” encompasses, and how not all articles tagged with the “peer reviewed” icon are actually research articles. In Primo, “peer reviewed” results may also include book reviews, poster session abstracts, and other PCI data feeding from metadata sources describing peer-reviewed journal contents.

Users were also asked to email an article to themselves. This was an attempt to determine if placement or the number of options in the “Send To” section impacts usage. This question tests the “Send To” location for the email icon on the full results page. The success rate for this task was 87% when averaged out over the five campuses. Several campuses have moved the “Send To” location to areas of the page other than the out-of-the-box position for these tool links. No matter where it was positioned within the page, users were generally able to locate the link and email it to themselves. Fullerton renamed this section “Tools” and that had no impact either.

While wording and placement did not affect user’s ability to send items to their email, it had a profound effect on another task. Users were asked to find a known item, the article “Stress During Residency Training.” This article was selected because none of the five campuses own this item. For users to access it they have to use the “Expand My Results” feature at the top of the facet panel. There was a great deal of variation across the campuses in this regard. At the time of testing, only San Marcos displayed the full PCI collection, so the article appears at the top of the list, which undoubtedly was the reason for the high rate of success at that location. Fullerton took the opposite approach, electing to move the option to the bottom of the facet panel. This meant that users had to scroll down the page to access that setting. Additionally, the default language “Expand My Results” was not altered. Dominguez Hills had the greatest success rate

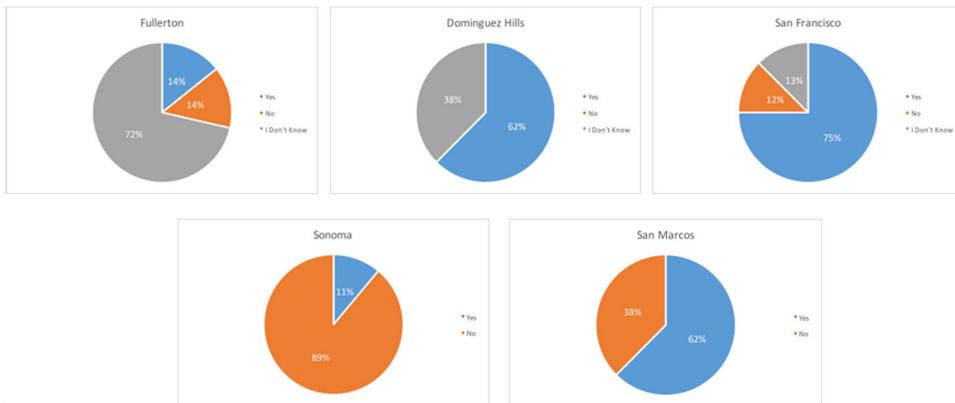


Figure 4. If you found a copy of the article, can you request it from another library?

among the non-PCI group. At that campus, they changed the language in an attempt to make it more comprehensible to users. They chose to change it to read “Include results outside of library databases,” which was more readily understood by users. Clearly positioning of the link for checking other libraries’ holdings, and the language used does matter.

The wide variance in success rate could also be due to differences in testing methodology. At San Francisco, the test administrator showed users how to find a copy of the article “Stress During Residency Training” to ensure this question was independent from the one asking them to locate the article only available at another campus. Not all testers took that step. Those that didn’t, either because the tester wasn’t in the room and recording the session, or because the tester did not demonstrate how to find the article (i.e., Dominguez Hills), would have an automatic failure or not-sure on this question from the users who failed at the previous one.

We asked users who found the article if they could request it from another library. This was to test their understanding of interlibrary loan, and their ability to access that feature. Each campus used different language to indicate that the item was not available in the home library, but could be requested from another. **This is another case where wording makes a big difference in the user experience.** At Dominguez Hills the message displays as, “we don’t have a physical copy in the library, but you can still get it,” whereas at Fullerton it read “To request to have this resource delivered to you (ILLiad) please sign in.” As evidenced in [Figure 4](#), the language used at Fullerton was not comprehensible to the few users that were able to locate the item.

There was a great deal of interest in the CSU consortium on the usability of the new CSU + resource sharing platform. CSU + is a CSU-first resource sharing system that makes the entire physical collection of the CSU visible to users. To test a user’s ability to find resources at other institutions, the

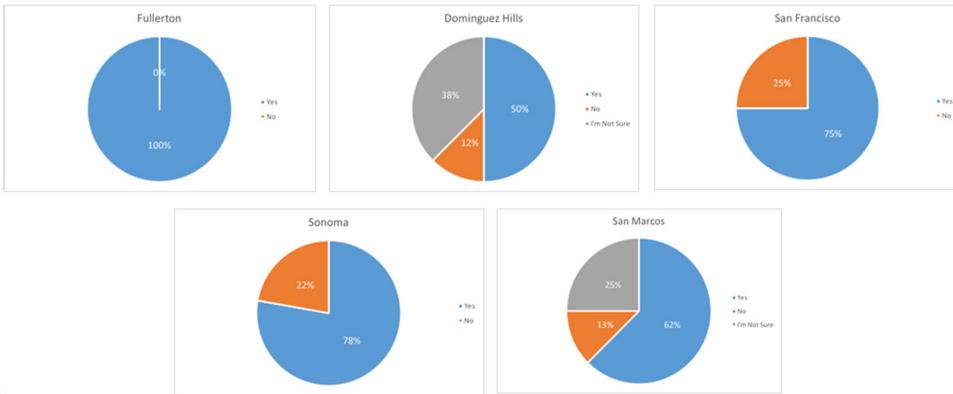


Figure 5. The library does not own the book “The Hat Party,” however another CSU does. Can you find which campus owns a copy of “The Hat Party”?

test designers had to locate a book that was not held by any of the five institutions participating in the test. The book “The Hat Party” was chosen as it was only held by the Chico campus. In order for users to succeed at this task they must find the record for the book, then click on the brief record to open the full record, and finally, locate the link to check other campuses’ holdings. The first step is complicated by different sorting options set at each campus. Additionally, results were affected by whether the user includes the word “The” in their search. Notably, the book “The Hat Party” is not present on the first page of results at Dominguez Hills and Sonoma when the “The” is omitted from the search, but is for the other campuses.

As with other aspects of the Primo configurations, there was variation among the five campuses in how they provided access to CSU+. Primo allows for the creation of separate search scopes, with a default search, which in the case of the CSU is called Everything, that usually consists of the print collection as well as the content from PCI. The CSU Discovery Working Group recommended that all campuses include all of the CSU + collection in their default search scope. Most campuses elected to include CSU+, content from the Primo Central Index and local collections to their default search scope. However, roughly $\frac{1}{3}$ of the campuses elected to only show local content in their default search. Testers on campuses that did not include CSU + content by default had difficulty finding “The Hat Party,” while users at Fullerton, which provides the full catalog had greater success finding the item, as seen in Figure 5.

A follow-up question asked users to identify which campus held a copy of “The Hat Party.” Users’ ability to complete this task would be predicated on their ability to complete the prior question. In this instance, the Fullerton campus displays the name of the library that holds the item available by default. On other campuses, such as Sonoma, the user would need

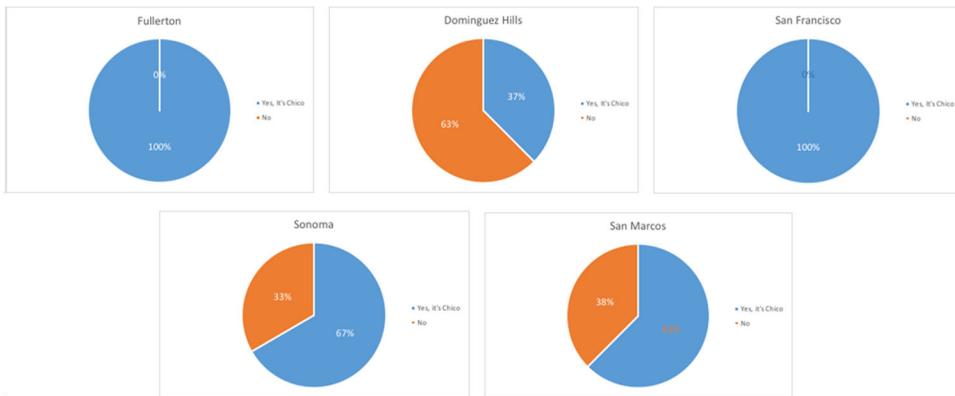


Figure 6. Were you able to tell which CSU owns a copy of “The Hat Party”?



Figure 7. If you found “The Hat Party,” are you able to request it from another library?

to sign in to Interlibrary loan to Primo which campus held the title, which created a barrier for the users. Similarly, at the San Marcos campus, the option to check holdings at other campuses is not visible in the record without signing in. At Dominguez Hills users were required to click on an icon that said “Show Libraries” in order to see the name of the campus holding the item, which is another barrier to access (Figure 6).

In an additional test of resource sharing usability, we asked users if they were able to request this item from the library that holds it (Figure 7). Again, this question is predicated by the ability of the users to find the item in the first place. Testing revealed that there is a general lack of confidence among most users about how and where to find information to request items from either CSU+ or Illiad. Through their individual General Electronic Services setup, Dominguez Hills has display language designed explicitly to indicate how to request a copy through CSU+ (“We don’t have a physical copy at CSUDH, but you can still get it. Sign-in to request it from another library”). Dominguez Hills also exposes the Interlibrary

Loan option earlier in the results display, before users log into Primo. Even if users didn't feel confident in finding "The Hat Party," they understood how to place a request for the title. Other libraries have no indication that requesting the item via CSU + or interlibrary loan is an option until after the user has logged in. Likewise, other institutions do not provide a clear explanation of what ILLiad, the interlibrary loan service, or CSU + is.

Discussion

Due to discrepancies in testing across campuses, uncertainty was introduced into result sets.

Further standardizing user experience testing across campuses and agreeing on a protocol for test administration up front, beyond the questions themselves, is a goal for future testing. A potential barrier to participation may come in the ability of different organizations to agree to committing similar time and resources to a group UX testing effort.

We identified areas where users are consistently able to perform certain tasks: sending results links to themselves and logging into a library account. The placement for these tasks within the UI is similar across multiple discovery systems and databases, and users may have had external familiarity with accomplishing these goals.

Users had more difficulty finding known items located outside of the library's collections, understanding how to initiate a lending request from either the consortial borrowing or Interlibrary Loan platform, and viewing which borrowing partners have it available. The CSU has invested significant resources in the CSU + consortial partnership; thus further testing and educational efforts should focus on effective use of our borrowing and lending systems.

Additionally, this study confirms that labels for certain tasks and facets are difficult for users to understand. Thanks to having a multi-tenant platform, we could test labeling customizations across multiple instances. We were then able to make important adjustments based on the demonstrated successes of existing customizations from partner libraries. For example, Dominguez Hills's language for the PCI facet produced favorable results in testing; Fullerton has now implemented this language for their PCI facet and experienced much greater usage of this functionality. Users may also benefit from further customizations to expose answers to frequently asked questions, like how to find a book on the shelf or initiating an interlibrary loan request.

While this study was somewhat hampered by the difficulty of multi-campus user testing, we were able to determine **several key recommendations for Primo configuration**. Many of our recommendations are informed by

users failing at the tasks given them. As a result, testing flaws will inform the best practices we will use in our future testing and suggest for others undertaking similar projects.

In terms of configuration suggestions, it became clear that adding extra information to guide users toward resources and services available was beneficial. Results strongly indicate the value of making books and articles not held locally visible by default. However, libraries should consult with their interlibrary loan departments before taking this step, as our experience, upon implementation of this, indicates that their workload will be substantially increased. Further research is necessary to determine how best to display items not available locally to minimize user confusion.

Libraries that choose not to include expanded search results by default should use clear language, visible high on the search results page, to explain that option. Consortial shared catalogs is a new concept for many users, so they will often need some details to understand it. Even the sharing of items between libraries is unfamiliar to some users, so explanations of interlibrary loan services should be made visible on catalog records even to users not logged in. Otherwise, they will never know that they have options for obtaining the item they found.

We recommend adding journal titles to searches limited to articles. Users don't consider the intermediate step of finding a journal and will take the option of searching for articles when it is presented to them even if they are search of a journal. In addition, when the search is limited to articles, users are often puzzled by the common following step of limiting to scholarly peer reviewed articles. A brief explanation on the results page of peer review and how one can limit to just those articles would be of great help to them. Underclassmen in particular often search for scholarly articles without understanding the concept.

Multi-campus user testing requires proper sharing of information as well. We found Google Docs to be an effective way to record and share results. It was convenient to use on a tablet or laptop while observing the user or while transcribing from hand-written notes. It automatically consolidated the data and made it simple to collate for analysis. However, we found the most useful information to be in the mode of failure when users were unable to complete the test tasks. We neglected to include spaces in Google Docs for recording the process, and then had to reconstruct events from notes taken. It would be much better to record at least some key details in Google Docs at the time.

Though we standardized test questions, we did not do enough to standardize the testing methodology. It is important to indicate whether the tester should be present and level of interaction they should have with the test subject between questions in general. For each question, the starting page and conditions need to be specified, particularly to note if they are

meant to be dependent on performance on the previous question or independent. Failure to do this adds additional variables to questions and makes results hard to interpret.

Finally, we recommend sharing preliminary results early in the process to identify these sorts of uncontrolled variables and miscommunications. Even if the first one or two tests at each location will need to be discarded, a revised testing procedure will ensure more reliable data. Our testing offers useful insights and can be a model for future iterations of multi-tenant testing for consortia.

Notes on contributors

William Jacobs is the Associate Librarian for STEM at California State University San Francisco's J. Paul Leonard Library. His professional interests include systematic reviews, user interface design, and online forum moderation. He earned his B.S. from the University of Delaware, and Masters from San Diego State University, Boston University, and Pratt Institute.

J. Michael DeMars is the Interim Associate Dean for Public Services at California State University Fullerton's Pollak Library, where he has also served as Systems and Reference & Instruction Librarian. His research interests include point-of-need information services, e-learning, mobile technologies, and web metrics. He earned his B.A. from the University of California, Santa Cruz, and his MLIS from San Jose State University.

Joanna Messer Kimmitt is the User Services Coordinator at California State University, Dominguez Hills. Her professional interests focus on user experience design, student and staff professional development, and access services workflows. She earned her BM from Lawrence University (WI), her MM from Carnegie Mellon University, and her MLIS from the University of Wisconsin-Milwaukee. She has worked in academic, public, and special libraries, as well as library consortial offices.

References

- Brett, K., A. Lierman, and C. Turner. 2016. "Lessons Learned: A Primo Usability Study." *Information Technology and Libraries* 35 (1):7–25. doi: [10.6017/ital.v35i1.8965](https://doi.org/10.6017/ital.v35i1.8965).
- California State University Libraries. 2017. "Discovery Migration and Implementation." *Unified Library Management System Wiki*. <https://calstate.atlassian.net/wiki/spaces/ULMSD/pages/88178689/Discovery±Migration±and±Implementation>.
- Comeaux, D. J. 2012. "Usability Testing of a Web-Scale Discovery System at an Academic Library." *College & Undergraduate Libraries* 19 (2-4):189–206. doi: [10.1080/10691316.2012.695671](https://doi.org/10.1080/10691316.2012.695671).
- DeMars, J. M. 2017. "Discovering Our Users: A Multi-Campus Usability Study of Primo." *IFLA Library*. <http://library.ifla.org/1810/1/S10-2017-demars-en.pdf>.
- Galbreath, B., C. Johnson, and E. Hvizdak. 2018. "Primo New User Interface: Usability Testing and Local Customizations Implemented in Response." *Information Technology and Libraries* 37 (2):10. doi: [10.6017/ital.v37i2.10191](https://doi.org/10.6017/ital.v37i2.10191).
- Hanrath, S., and M. Kottman. 2015. "Use and Usability of a Discovery Tool in an Academic Library." *Journal of Web Librarianship* 9 (1):1–21. doi: [10.1080/19322909.2014.983259](https://doi.org/10.1080/19322909.2014.983259).

- Jarrett, K. 2012. "Findit@Flinders: User Experiences of the Primo Discovery Search Solution." *Australian Academic & Research Libraries* 43 (4):278–99. doi: [10.1080/00048623.2012.10722288](https://doi.org/10.1080/00048623.2012.10722288).
- Moore, D., and N. Mealey. 2016. "Consortial-Based Customizations for New Primo UI." *Code4Lib Journal* (34). <http://journal.code4lib.org/articles/11948>.
- Nichols, A., A. Billey, P. Spitzform, A. Stokes, and C. Tran. 2014. "Kicking the Tires: A Usability Study of the Primo Discovery Tool." *Journal of Web Librarianship* 8 (2): 172–95. doi: [10.1080/19322909.2014.903133](https://doi.org/10.1080/19322909.2014.903133).
- Valentine, B., and B. West. 2016. "Improving Primo Usability and Teachability with Help from the Users." *Journal of Web Librarianship* 10 (3):176–96. doi: [10.1080/19322909.2016.1190678](https://doi.org/10.1080/19322909.2016.1190678).