LaneConnex: An Integrated Biomedical Digital Library Interface

Debra S. Ketchell, Ryan Max Steinberg, Charles Yates, and Heidi A. Heilemann

This paper describes one approach to creating a search application that unlocks heterogeneous content stores and incorporates integrative functionality of Web search engines. LaneConnex is a search interface that identifies journals, books, databases, calculators, bioinformatics tools, help information, and search hits from more than three hundred full-text heterogeneous clinical and bioresearch sources. The user interface is a simple query box. Results are ranked by relevance with options for filtering by content type or expanding to the next most likely set. The system is built using component-oriented programming design. The underlying architecture is built on Apache Cocoon, Java Servlets, XML/XSLT, SQL, and JavaScript. The system has proven reliable in production, reduced user time spent finding information on the site, and maximized the institutional investment in licensed resources.

Most biomedical libraries separate searching for resources held locally from external database searching, requiring clinicians and researchers to know which interface to use to find a specific type of information. Google, Amazon, and other Web search engines have shaped user behavior and expectations.[1] Users expect a simple query box with results returned from a broad array of content ranked or categorized appropriately with direct links to content, whether it is an HTML page, a PDF document, a streaming video, or an image. Biomedical libraries have transitioned to digital journals and reference sources, adopted OpenURL link resolvers, and created institutional repositories. However, students, clinicians, and researchers are hindered from maximizing this content because of proprietary and heterogeneous systems. A strategic challenge for biomedical libraries is to create a unified search for a broad spectrum of licensed, open-access, and institutional content.

Background

Studies show that students and researchers will use the search path of least cognitive resistance.[2] Ease and speed are the most important factors for using a particular search engine. A University of California report found that academic users want one search tool to cover a wide information universe, multiple formats, full-text availability to move seamlessly to the item itself, intelligent assistance and spelling correction, results sorted in order of relevance, help navigating large retrievals by logical subsetting and customization, and seamless access anytime, anywhere.[3] Studies of clinicians in the patient-care environment have documented that effort is the most important factor in whether a patient-care question is pursued.[4] For researchers, finding and using the best bioinformatics tool is an elusive problem.[5]

In 2005, the Lane Medical Library and Knowledge Management Center (Lane) at the Stanford University Medical Center provided access to an expansive array of licensed, institutional, and open-access digital content in support of research, patient care, and education. Like most of its peers, Lane users were required to use scores of different interfaces to search external databases and find digital resources. We created a local metasearch application for clinical reference content, but it did not integrate result sets from disparate resources. A review of federated-search software in the marketplace found that products were either slow or they limited retrieval when faced with a broad spectrum of biomedical content. We decided to build on our existing application architecture to create a fast and unified interface.

A detailed analysis of Lane website-usage logs was conducted before embarking on the creation of the new search application. Key points of user failure in the existing search options were spelling errors that could easily be corrected to avoid zero results; lack of sufficient intuitive options to move forward from a zero-results search or change topics without backtracking; lack of use of existing genre or role searches; confusion about when to use the resource, OpenURL resolver, or PubMed search to find a known item; and results that were cognitively difficult to navigate. Studies of the Web search engine and the PubMed search log concurred with our usage-log analysis: A single term search is the most common, with three words maximum entered by typical users.[6] A PubMed study found that 22 percent of user queries were for known items rather than for a general subject, confirming our own log analysis findings that the majority of searches were for a particular source item.[7] Search-term analysis revealed that many of our users were entering partial article citations (e.g., author, date) in any query

Debra S. Ketchell (debra.ketchell@gmail.com) is the former Associate Dean for Knowledge Management and Library Director; **Ryan Max Steinberg** (ryan.max.steinberg@stanford.edu) is the Knowledge Integration Programmer/Architect; **Charles Yates** (charles.yates@stanford.edu) is the Systems Software Developer; and **Heidi A. Heilemann** (heidi.heilemann@stanford.edu) is the former Director for Research & Instruction and current Associate Dean for Knowledge Management and Library Director at the Lane Medical Library & Knowledge Management Center, Information Resources & Technology, Stanford University School of Medicine, Stanford, California.

box expecting that article databases would be searched concurrently with the resource database. Our displayed results were sorted alphabetically, and each version of an item was displayed separately. For the user, this meant a cluttered list with redundant title information that increased their cognitive effort to find meaningful items. Overall, users were confronted with too many choices upfront and too few options after retrieving results. Focus groups of faculty and students were conducted in 2005. Attendees wanted local information integrated into the proposed single search. Local information included content such as how-to information, expertise, seminars, grand rounds, core lab resources, drug formulary, patient handouts, and clinical calculators. Most of this content is restricted to the Stanford user population. Users consistently described their need for a simple search interface that was fast and customized to the Stanford environment.

In late 2005, we embarked on a project to design a search application that would address both existing points of failure in the current system and meet the expressed need for a comprehensive discovery-and-finding tool as described in focus groups. The result is an application called LaneConnex.

Design objectives

The overall goal of LaneConnex is to create a simple, fast search across multiple licensed, open-access, and special-object local knowledge sources that depackages and reaggregates information on the basis of Stanford institutional roles. The content of Lane's digital collection includes forty-five hundred journal titles and forty-two thousand other digital resources, including video lectures, executable software, patient handouts, bioinformatics tools, and a significant store of digitized historical materials as a result of the Google Books program. Media types include HTML pages, PDF documents, JPEG images, MP3 audio files, MPEG4 videos, and executable applications.

More than three hundred reference titles have been licensed specifically for clinicians at the point of care (e.g., UpToDate, eMedicine, STAT-Ref, and Micromedex Clinical Evidence). Clinicians wanted their results to reflect subcomponents of a package (e.g., results from the Micromedex patient handouts). Other clinical content is institutionally managed (e.g., institutional formulary, lab test database, or patient handouts). More than 175 biomedical research tools have been licensed or selected from open-access content. The needs of biomedical researchers include molecular biology tools and software, biomedical literature databases, citation analysis, chemical and engineering databases, expertise-finding tools, laboratory tools and supplies, institutional-research resources, and upcoming seminars.

The specific objectives of the search application are the following:

- The user interface should be fast, simple, and intuitive, with embedded suggestions for improving search results (e.g., Did you mean? Didn't find it? Have you tried?).
- Search results from disparate local and external systems should be integrated into a single display based on popular search-engine models familiar to the target population.
- The query-retrieval and results display should be separated and reusable to allow customization by role or domain and future expansion into other institutional tools.
- Resource results should be ranked by relevance and filtered by genre.
- Metasearch results should be hit counts and filtered by category for speed and breadth. Results should be reusable for specific views by role.
- Finding a known article or journal should be streamlined and directly link to the item or "get item" option.
- The most popular search options (PubMed, Google, and Lane journals) should be ubiquitous.
- Alternative pathways should be dynamic and interactive at the point of need to avoid backtracking and dead ends.
- User behavior should be tracked by search term, resource used, and user location to help the library make informed decisions about licensing, metadata, and missing content.
- Off-the-shelf software should be used when available or appropriate with development focused on search integration.
- The application should be built upon existing metadata-creation systems and trusted Web-development technologies.

Based on these objectives, we designed an application that is an extension of existing systems and technologies. Resources are acquired and metadata are provided using the Voyager integrated library system (ILS). The SFX OpenURL link resolver provides full-text article access and expands the title search beyond biomedicine to all online journals at Stanford. EZproxy provides seamless off-campus access. WebTrends provides usage tracking. Movable Type is used to create FAQ and help information. A locally developed metasearch application provides a cross search with hit results from more than three hundred external and internal full-text sources. The technologies used to build LaneConnex and integrate all of these systems include Extensible Stylesheet Language

Transformations (XSLT), Java, JavaScript, the Apache Cocoon project, and Oracle.

Systems Description

Architecture

LaneConnex is built on a principle of separation of concerns. The Lane content owner can directly change the inclusion of search results, how they are displayed, and additional path-finding information. Application programmers use Java, JavaScript, XSLT, and Structured Query Language (SQL) to create components that generate and modify the search results. The merger of content design and search results occurs "just in time" in the user's browser.

We use component-oriented programming design whereby services provided within the application are defined by simple contracts. In LaneConnex, these components (called "transformers") consume XML information and, after transforming it in some way, pass it on to some other component. A particular contract can be fulfilled in different ways for different purposes. This component architecture allows for easy extension of the underlying Apache Cocoon application. If LaneConnex needs to transform some XML data that is not possible with built-in Cocoon transformers, it is a simple matter to create a software component that does what is needed and fulfills the transformer contract.

Apache Cocoon is the underlying architecture for LaneConnex, as illustrated in figure 1. This Java Servlet is an XML–publishing engine that is built upon a component framework and uses a pipeline-processing model. A declarative language uses pattern matching to associate sets of processing components with particular request URLs. Content can come from a variety of sources. We use content from the local file system, network file system, HTTP, and a relational database. The XSLT language is used extensively in the pipelines and gives fine control of individual parts of the documents being processed. The end of processing is usually an XHTML document but can be any common MIME type. We use Cocoon to separate areas of concern so things like content, look and feel, and processing can all be managed as separate entities by different groups of people with little effect on another area. This separation of concerns is manifested by template documents that contain most of the HTML content common to all pages and are then combined with content documents within a processing pipeline. The declarative nature of the sitemap language and XSLT facilitate rapid development with no need to redeploy the entire application to make changes in its behavior.

The LaneConnex search is composed of several components integrated into a query-and-results interface: Oracle resource metadata, full-text metasearch application, Movable Type blogging software, "Did you mean?" spell checker, EZproxy remote access, and WebTrends tracking.

Full-text Metasearch

Integration of results from Lane's metasearch application illustrates Cocoon's many strengths. When a user searches LaneConnex, Cocoon sends his or her query to the metasearch application, which then dispatches the request to multiple external, full-text search engines and content stores. Some examples of these external resources are UpToDate, Access Medicine, Micromedex, PubMed, and MD Consult. The metasearch application interacts with these external resources through Jakarta Commons HTTP clients. Responses from external resources are turned into W3C Document Object Model (DOM) objects, and XPath expressions are used to resolve hit counts from the DOM objects. As result counts are returned, they are added to an XML–based result list and returned to Cocoon.

The power of Cocoon becomes evident as the XML–based metasearch result list is combined with a separate display template. This template-based approach affords content curators the ability to directly add, group, and describe metasearch resources using the language and look that is most meaningful to their specific user communities. For example, there are currently eight metasearch templates curated by an informationist in partnership with a target community. Curating these templates requires little to no assistance from programmers.

In Lane's 2005 interface, a user's request was sent to the metasearch application, and the application waited five seconds before responding to give external resources a chance to return a result. Hit counts in the user interface included a link to refresh and retrieve more results from external resources that had not yet responded. Usability studies showed this to be a significant user barrier, since the refresh link was rarely clicked. The initial five second delay also gave users the impression that the site was slow. The LaneConnex application makes heavy use of JavaScript to solve this problem. After a user makes her initial request, JavaScript is used to poll the metasearch application (through Cocoon) on the user's behalf, popping in result counts as external resources respond. This adds a level of interactivity previously unavailable and makes the metasearch piece of LaneConnex much more successful than its previous version.

Resource metadata

LaneConnex replaces the catalog as the primary discovery interface. Metadata describing locally owned and

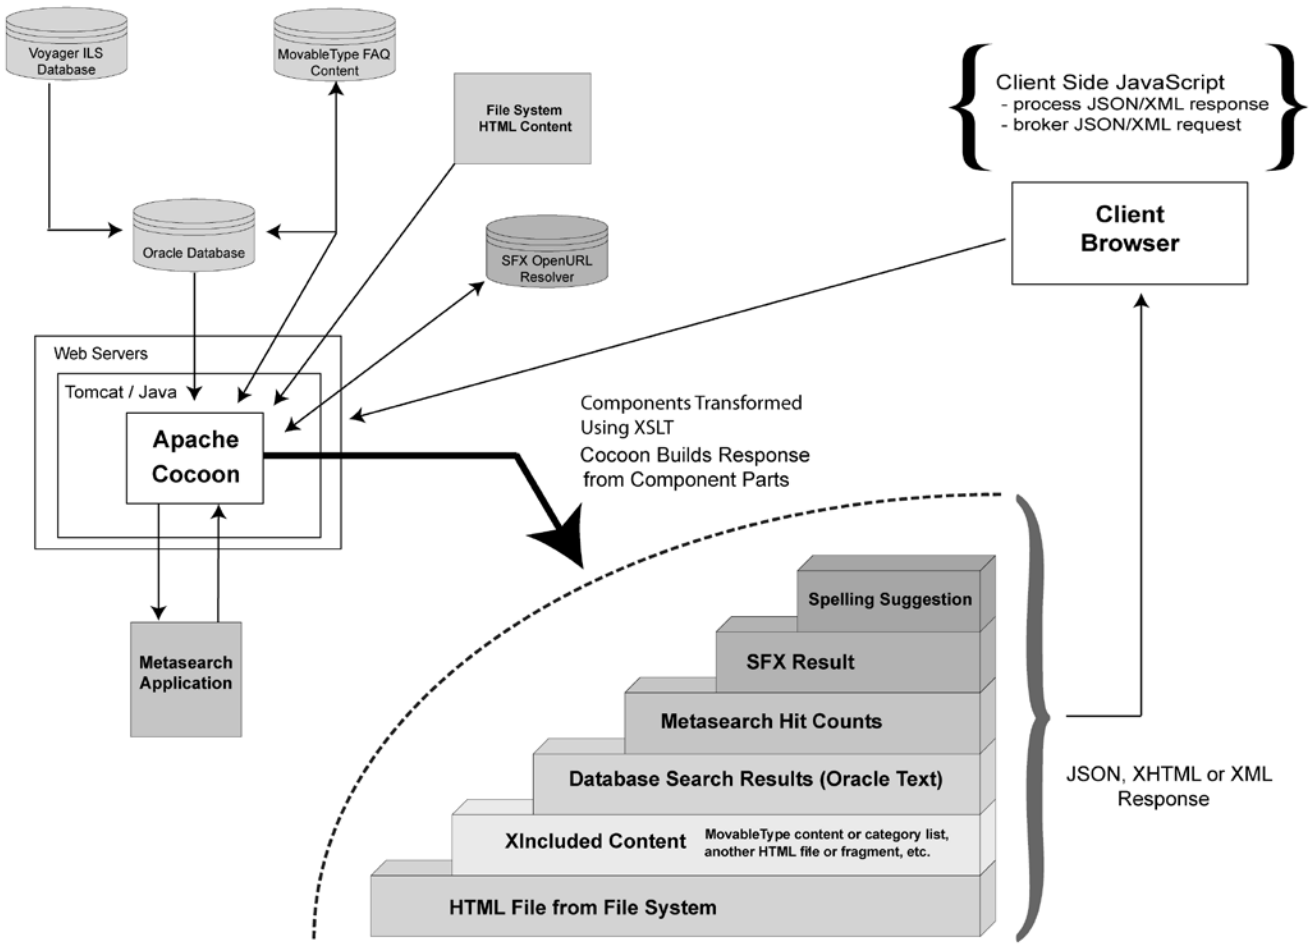

Figure 1. LaneConnex Architecture.

licensed resources (journals, databases, books, videos, images, calculators, and software applications) are stored in the library's current system of record, an instance of the Voyager ILS. LaneConnex makes no attempt to replace Voyager's strengths as an application for the selection, acquisition, description, and management of access to library resources. It does, however, replace Voyager's discovery interface. To this end, metadata for about eight thousand digital resources is extracted from Voyager's Oracle database, converted into MARCXML, processed with XSLT, and stored in a simple relational database (six tables and twenty-nine attributes) to support fast retrieval speed and tight control over search syntax. This extraction process occurs nightly, with incremental updates every five minutes. The Oracle Text search engine provides functionality anticipated by our Internet-minded users. Key features are speed and relevance-ranked results. A highly refined results ranking insures that the logical title appears in the first few results. A user's query is parsed for wildcard, Boolean, proximity, and phrase operators, and then translated into an SQL query. Results are then transformed into a display version.

Related services

LaneConnex compares a user's query terms against a dictionary. Each query is sent to a Cocoon spell-checking component that returns suggestions where appropriate. This component currently uses the Simple Object

Access Protocol (SOAP)–based spelling service from Google. Google was chosen over the National Center for Biotechnology Information (NCBI) spelling service because of the breadth of terms entered by users; however, Cocoon's component-oriented architecture would make it trivial to change spell checkers in the future.

Each query is also compared against Stanford's OpenURL link resolver (FindIt@Stanford). Client-side JavaScript makes a Cocoon-mediated query of FindIt@Stanford. Using XSLT, FindIt@Stanford responses are turned into JavaScript Object Notation (JSON) objects and popped into the interface as appropriate. Although the vast majority of LaneConnex searches result in zero FindIt@Stanford results, the convenience of searching all of Lane's systems in a single, unified interface far outweighs the effort of implementation.

A commercial analytics tool called WebTrends is used to collect Web statistics for making data-centric decisions about interface changes. WebTrends uses client-side JavaScript to track specific user click events. Libraries need to track both on-site clicks (e.g., the user clicked on "Clinical Portal" from the home page) and off-site clicks (e.g., the user clicked on "Yamada's Gastroenterology" after doing a search for "IBS"). To facilitate off-site click capture, WebTrends requires every external link to include a snippet of JavaScript. Requiring content creators to input this code by hand would be error prone and tedious. LaneConnex automatically supplies this code for every class of link (search or static). This specialized WebTrends method provides Lane with data to inform both interface design and licensing decisions.

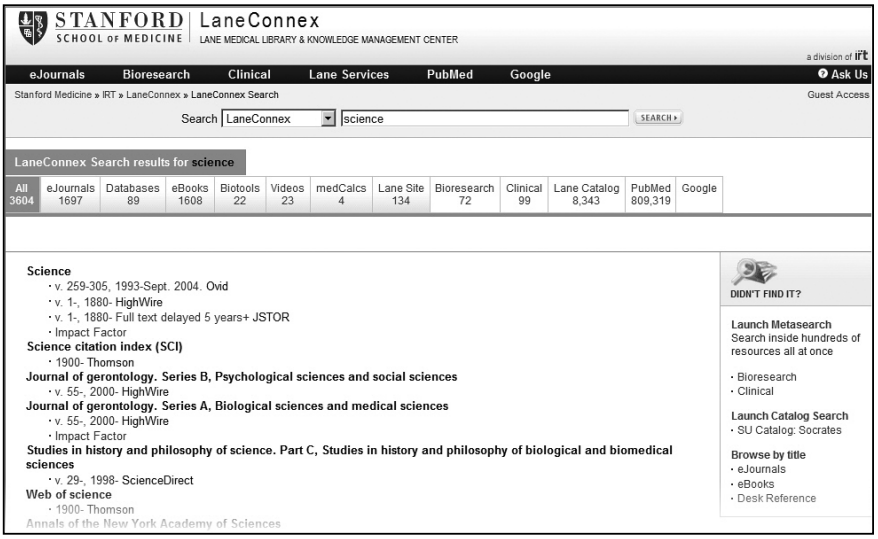

Figure 2. LaneConnex Resource Search Results. Resource results are ranked by relevance. Single word titles are given a higher weight in the ranking algorithm to insure they are displayed in the first five results. Uniform titles are used to co-locate versions (e.g., the three instances of *Science* from different producers). Journals titles are linked to their respective impact factor page in the ISI Web of Knowledge. Digital formats that require special players or restrictions are indicated. The metadata searched for eJournals, Databases, eBooks, Biotools, Video, and medCalcs are Lane's digital resources extracted from the integrated library system into a searchable Oracle database. The first "All" tab is the combined results of these genres and the Lane Site help and information.

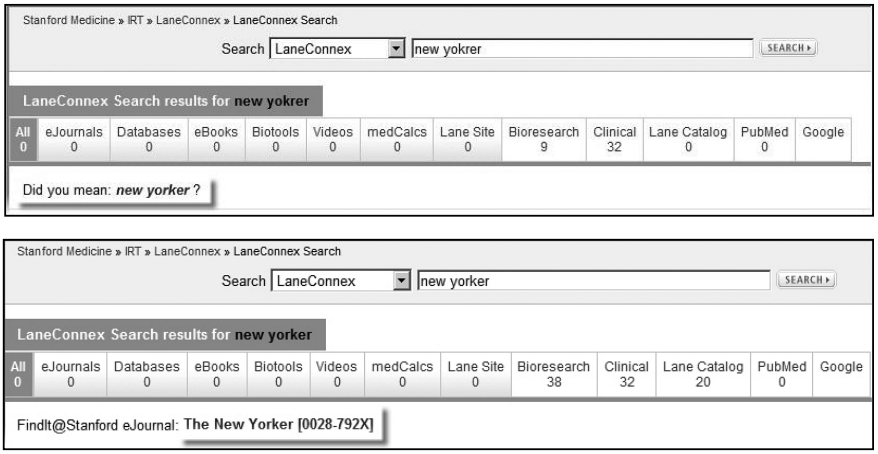

Figure 3. LaneConnex Related Services Search Enhancements. LaneConnex includes a spell checker to avoid a common failure in user searches. AJAX services allow the inclusion of search results from other sources for common zero results failures. For example, the Stanford link resolver database is simultaneously searched to insure online journals outside the scope of biomedicine are presented as a linked result for the user.

Results

LaneConnex version 1.0 was released to the Stanford biomedical community in July 2006. The current application can be experienced at http://lane.stanford.edu. The production version has proven reliable over two years. Incremental user focus groups have been employed to improve the interface as issues arose. A series of vignettes will be used to illustrate how the current version of

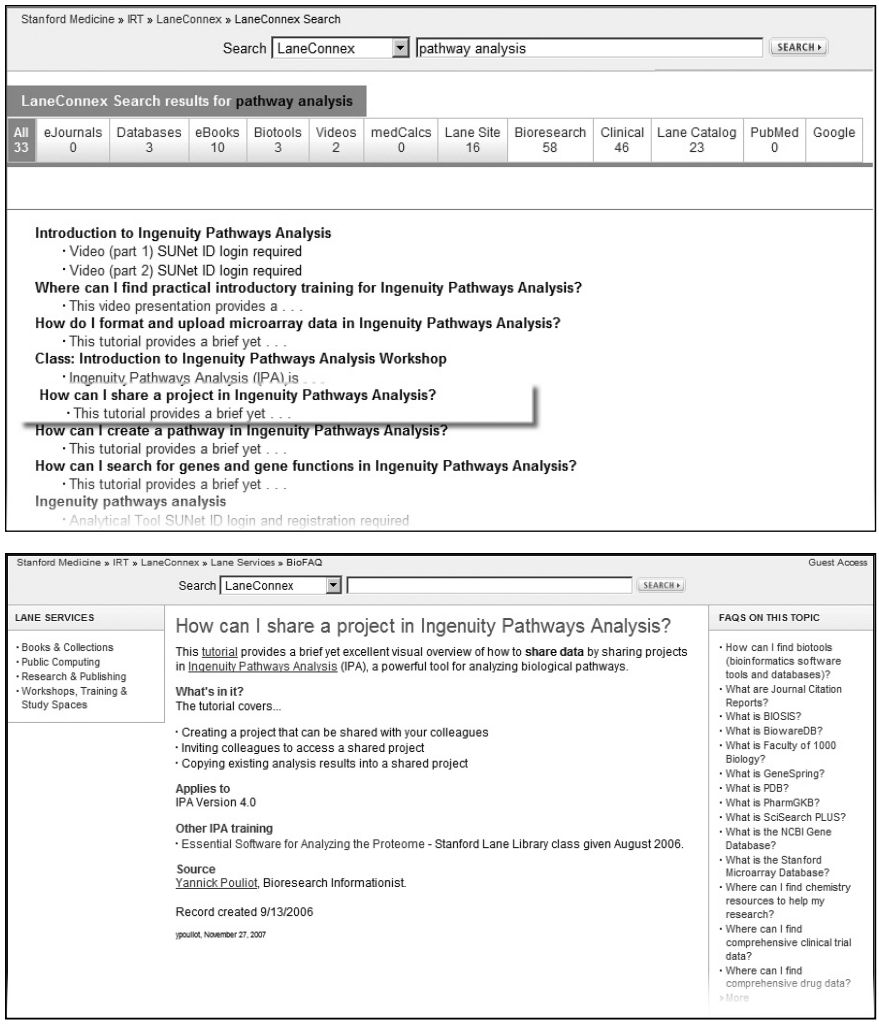

Figure 4. Example of Integration of Local Content Stores. Help information is managed in Moveable Type and integrated into LaneConnex search results.

LaneConnex meets the design objectives from the user's perspective.

- User query: "science." A graduate student is looking for the journal *Science*. The LaneConnex results are listed in relevance order (see figure 2). Single-word titles are given a higher weight in the ranking algorithm to insure they are displayed in the first five results. Results from local metadata are displayed by uniform title. For example, Lane has three instances of the journal *Science*, and each version is linked to the appropriate external store. Brief notes provide critical information for particular resources. For example, restricted local patient education documents and video seminars note that the "SUNetID login" is required.
- User query: "new yokrer." A faculty member is looking for an article in the *New Yorker* for a class reading assignment. He makes a typing error, which invokes the "Did you mean?" function (see figure 3). He clicks on the correct spelling. No results are found in the resource search, but a simultaneous search of the link-resolver database finds an instance of this title licensed for the campus and displays a clickable link for the user.
- User query: "pathway analysis." A post–doc is looking for information on how to share an Ingenuity pathway. Figure 4 illustrates the integration of the locally created Lane FAQs. FAQs comprise a broad spectrum of help and how-to information as described by our focus groups. Help text is created in the Movable Type blog software, and made searchable through the LaneConnex application. The Movable Type interface lowers the barrier to HTML content creation by any staff member. More complex answers include embedded images and videos to enable the user to see exactly how to do a particular procedure. Cocoon allows for the syndication of subsets of this FAQ content back into static HTML pages where it can be displayed as both category-specific lists or as the text for scroll-over help for a link. Having a single store of help information insures the content is updated once for all instances.
- User query: "uterine cancer kapp." A resident is looking for a known article. LaneConnex simultaneously searches PubMed to increase the likelihood of user success (see figure 5). Clicking on the PubMed tab retrieves the results in the native interface; however, the user sees the PubMed@Stanford version, which includes embedded links to the article based on our OpenURL link resolver. The ability to retrieve results from bibliographic databases that includes article resolution insures that our biomedical community is always using the correct URL to insure maximum full-text article access. User testing in 2007 found that adding the three most frequently used sources (PubMed, Google, and Lane Catalog) into our one-box LaneConnex search was a significant time saver. It addresses

the expectation on the part of our users that they could search for an article or a journal title in a single search box without first selecting a database.

- User query: "serotonin pulmonary hypertension." A medical student is looking for the correlation of two topics. Clicking on the "Clinical" tab, the student sees the results of the clinical metasearch in figure 6. Metasearch results are deep searches of sources within licensed packages (e.g., textbooks in MD Consult or a specific database in Micromedex), local content (e.g., Stanford's lab-test database), and open-access content (e.g., NCBI databases). PubMed results are tailored strategies tiered by evidence. For example, the evidence-summaries strategy retrieves results from twelve clinical-evidence resources (e.g., BUJ, Clinical Evidence, and Cochrane Systematic Reviews) that link to the full-text licensed by Stanford. An example of the bioresearch metasearch is shown in figure 7. Content selected for this audience includes literature databases, funding sources, patents, structures, clinical trials, protocols, and Stanford expertise integrated with gene, protein, and phenotype tools.

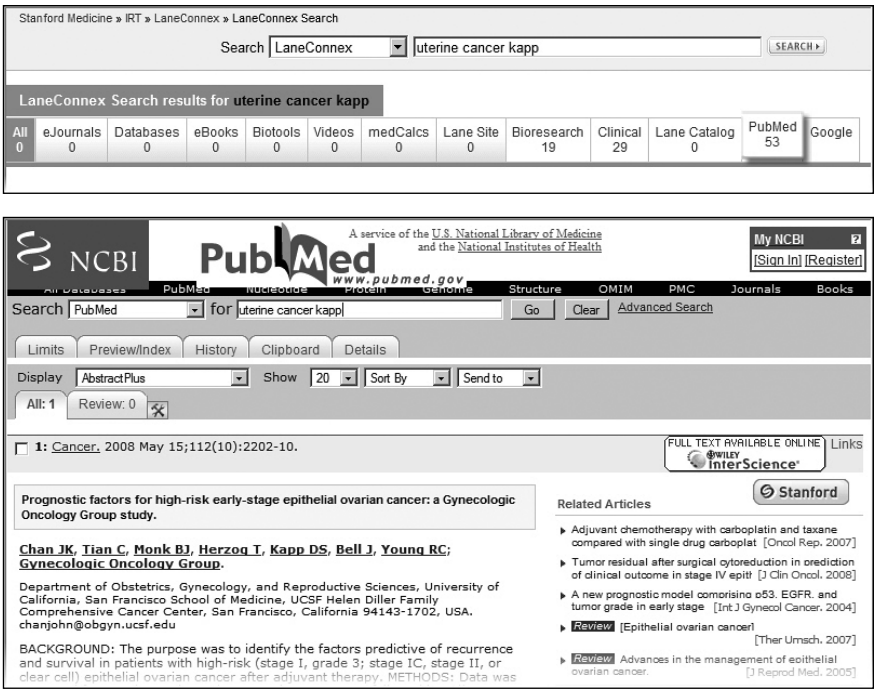

Figure 5. Example of Integration of Popular Search Engines into LaneConnex Results. Three of the most popular searches based on usage analysis are included at the top level. PubMed and Google are mapped to Lane's link resolver to retrieve the full article.

User testing revealed that many users did not click on the "Clinical" tab. The clinical metasearch was originally developed for the Clinical portal page and focused on clinicians in practice; however, the results needed to be exposed more directly as part of the LaneConnex search. Figure 8 illustrates the "Have you tried?" feature that displays a few relevant clinical-content sources without requiring the user to select the "Clinical" tab. This feature is managed by the SmartSearch component of the LaneConnex system. SmartSearch sends the user's query terms to PubMed, extracts a subset of articles associated with those terms, extracts the MeSH headings for those articles, and computes the frequency of headings in the articles to determine the most likely MeSH terms associated with the user's query terms. These MeSH terms are mapped to MeSH terms associated with each metasearch resource. Preliminary evaluation indicates that the clinical content is now being discovered by more users.

Creating or editing metasearch templates is a curator-driven task. Programming is only required to add new sources to the metasearch engine. A curator may choose from more than three hundred sources to create a discipline-based layout using general templates. Names, categories, and other description information are all at the curator's discretion. While developing new subspecialty templates, we discovered that clinicians were confused by the difference in layout of their specialty portal and their metasearch results (e.g., the Cardiology portal used the generic clinical metasearch). To address this issue, we devised an approach that merges a portal and metasearch into a single entity as illustrated in figure 9. A combination of the component-oriented architecture of LaneConnex and JavaScript makes the integration of metasearch results into a new template patterned after a portal easy to implement. This strategy will enable the creation of templates contextually appropriate to knowledge requests originating from electronic medical-record systems in the future.

Direct user feedback and usage statistics confirm that search is now the dominant mode of navigation. The amount of time each user spends on the website has dropped since the release of version 1.0. We speculate that the integrated search helps our users find relevant

Figure 6. Integration of metasearch results into LaneConnex. Results from two general, role-based metasearches (Bioresearch and Clinical) are included in the LaneConnex interface. The first image shows a clinician searching LaneConnex for *serotonin pulmonary hypertension*. Selecting the Clinical tab presents the clinical content metasearch display (second image), and is placed deep inside the source by selecting a title (third image).

information more efficiently. Focus groups with students are uniformly positive. Graduate students like the ability to find digital articles using a single search box. Medical students like the clinical metasearch as an easy way to look up new topics in texts and customized PubMed searches. Bioengineering students like the ability to easily look up patient care–related topics. Pediatrics residents and attendings have championed the development of their portal and metasearch focused on their patient population. Medical educators have commented on their ability to focus on the best information sources.

Discussion

A review of websites in 2007 found that most biomedical libraries had separate search interfaces for their digital resources, library catalog, and external databases. Biomedical libraries are implementing metasearch software to cross search proprietary databases. The University of California, Davis is using the MetaLib software to federate searching multiple bibliographic databases.[8] The University of South California and Florida State University are using WebFeat software to search clinical textbooks.[9] The Health Sciences Library System at the University of Pittsburgh is using Vivisimo to search clinical textbooks and bioresearch tools.[10] Academic libraries are introducing new "resource shopping" applications, such as the Endeca project at North Carolina State University, the Summa project at the University of Aarhus, and the VuFind project at Villanova University.[11] These systems offer a single query box, faceted results, spell checking, recommendations based on user input, and Asynchronous JavaScript and XML (AJAX) for live status information.

We believe our approach is a practical integration for our biomedical community that bridges finding a resource and finding a specific item through

a metasearch of multiple databases. The LaneConnex application searches across digital resources and external data stores simultaneously and presents results in a unified display. The limitation to our approach is that the metasearch returns only hit counts rather than previews of the specific content. Standardization of results from external systems, particularly receipt of XML results, remains a challenge. Federated search engines do integrate at this level, but are usually slow or limit the number of results. True integration awaits Health Level Seven (HL7) Clinical Decision Support standards and National Information Standards Organization (NISO) MetaSearch initiative for query and retrieval of specific content.[12]

One of the primary objectives of LaneConnex is speed and ease of use. Ranking and categorization of results has been very successful in the eyes of the user community. The integration of metasearch results has been particularly successful with our pediatric specialty portal and search. However, general user understanding of how the clinical and biomedical tabs related to the genre tabs in LaneConnex has been problematic. We reviewed Web engines and found a similar challenge in presenting disparate format results (e.g., video or image search results) or lists of hits from different systems (e.g., NCBI's Entrez search results).[13] We are continuing to develop our new specialty portal-and-search model and our SmartSearch term-mapping component to further integrate results.

Conclusion

LaneConnex is an effective and open-ended search infrastructure for integrating local resource metadata and full-text content used by clinicians and biomedical researchers. Its effectiveness comes from the recognition that users prefer a single query box with relevance or categorically organized results that lead them to the most likely

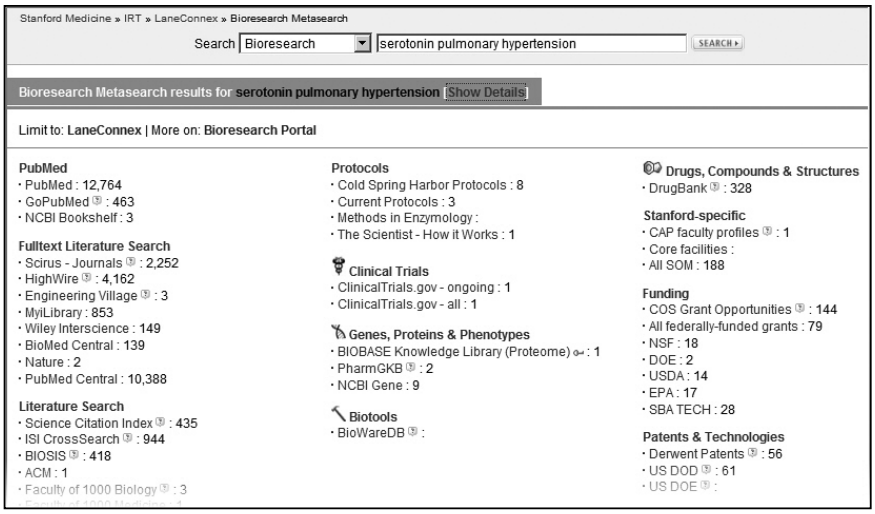

Figure 7. Example of a Bioresearch Metasearch.

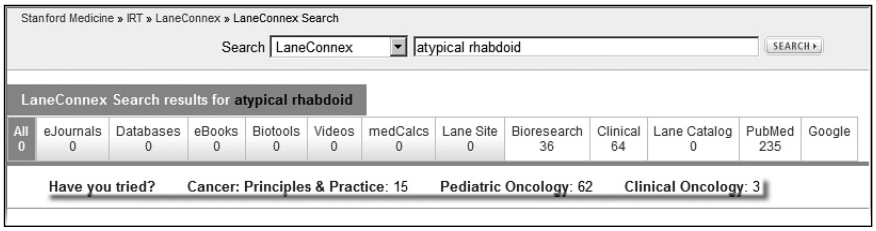

Figure 8. The SmartSearch component embeds a set of the metasearch results into the LaneConnex interface as "Have you tried?" clickable links. These links are the equivalent of selecting the title from a clinical metasearch result. The example search for atypical malignant rhabdoid tumor (a rare childhood cancer) invokes oncology and pediatric textbook results. These texts and PubMed provide quick access for a medical student or resident on the pediatric ward.

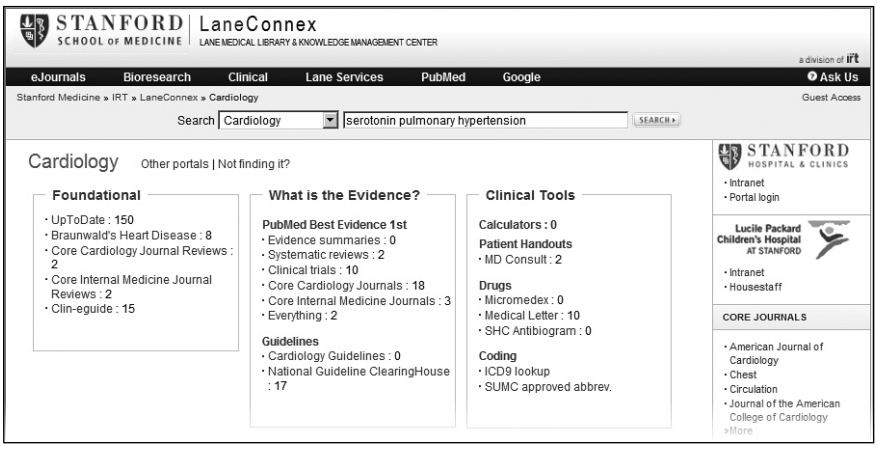

Figure 9. Example of a Clinical Specialty Portal with Integrated Metasearch. Clinical portal pages are organized so metasearch hit counts can display next to content links if a user executes a search. This approach removes the dissonance clinicians felt existed between separate portal page and metasearch results in version 1.0.

answer to a question or prospects in their exploration. The application is based on separation of concerns and is easily extensible. New resources are constantly emerging, and it is important that libraries take full advantage of existing and forthcoming content that is tailored to their user population regardless of the source. The next major step in the ongoing development of LaneConnex is becoming an invisible backend application to bring content directly into the user's workflow.

Acknowledgements

The authors would like to acknowledge the contributions of the entire LaneConnex technical team, in particular Pam Murnane, Olya Gary, Dick Miller, Rick Zwies, and Rikke Ogawa for their design contributions, Philip Constantinou for his architecture contribution, and Alain Boussard for his systems development contributions.

References

1. Denise T. Covey, "The Need to Improve Remote Access to Online Library Resources: Filling the Gap between Commercial Vendor and Academic User Practice," *Portal Libraries and the Academy* 3 no.4 (2003): 577–99; Nobert Lossau, "Search Engine Technology and Digital Libraries," *D-Lib Magazine* 10 no. 6 (2004), www.dlib.org/dlib/june04/lossau/06lossau.html (accessed Mar. 1, 2008); OCLC, "College Students' Perception of Libraries and Information Resource," www.oclc.org/reports/perceptionscollege.htm (accessed Mar 1, 2008); and Jim Henderson, "Google Scholar: A Source for Clinicians," *Canadian Medical Association Journal* 12 no. 172 (2005).

2. Covey, "The Need to Improve Remote Access to Online Library Resources"; Lossau, "Search Engine Technology and Digital Libraries"; OCLC, "College Students' Perception of Libraries and Information Resource."

3. Jane Lee, "UC Health Sciences Metasearch Exploration. Part 1: Graduate Student Gocus Group Findings," UC Health Sciences Metasearch Team, www.cdlib.org/inside/assess/evaluation_activities/docs/2006/draft_gradReport_march2006.pdf (accessed Mar. 1, 2008).

4. Karen K. Grandage, David C. Slawson, and Allen F. Shaughnessy, "When Less is More: a Practical Approach to Searching for Evidence-Based Answers," *Journal of the Medical Library Association* 90 no. 3 (2002): 298–304.

5. Nicola Cannata, Emanuela Merelli, and Russ B. Altman, "Time to Organize the Bioinformatics Resourceome," *PLos Computational Biology* 1 no. 7 (2005): e76.

6. Craig Silverstein et al., "Analysis of a Very Large Web Search Engine Query Log," www.cs.ucsb.edu/~almeroth/classes/tech-soc/2005-Winter/papers/analysis.pdf (accessed Mar. 1, 2008); Anne Aula, "Query Formulation in Web Information Search," www.cs.uta.fi/~aula/questionnaire.pdf (accessed Mar. 1, 2008); Jorge R. Herskovic, Len Y. Tanaka, William Hersh, and Elmer V. Bernstam, "A Day in the Life of PubMed: Analysis of a Typical Day's Query Log," *Journal of the American Medical Informatics Association* 14 no. 2 (2007): 212–20.

7. Herskovic, "A Day in the Life of PubMed."

8. Davis Libraries University of California, "QuickSearch," http://mysearchspace.lib.ucdavis.edu/ (accessed Mar. 1, 2008).

9. Eileen Eandi, "Health Sciences Multi-eBook Search," *Norris Medical Library Newsletter* (Spring 2006), Norris Medical Library, University of Southern California, www.usc.edu/hsc/nml/lib-information/newsletters.html (accessed Mar. 1, 2008); Maguire Medical Library, Florida State University, "WebFeat Clinical Book Search," http://med.fsu.edu/library/tutorials/webfeat2_viewlet_swf.html (accessed Mar. 1, 2008).

10. Jill E. Foust, Philip Bergen, Gretchen L. Maxeiner, and Peter N. Pawlowski, "Improving E-Book Access via a Library-Developed Full-Text Search Tool," *Journal of the Medical Library Association* 95 no. 1 (2007): 40–45.

11. North Carolina State University Libraries, "Endeca at the NCSU Libraries," www.lib.ncsu.edu/endeca (accessed Mar. 1, 2008); Hans Lund, Hans Lauridsen, and Jens Hofman Hansen, "Summa—Integrated Search," www.statsbiblioteket.dk/publ/summaenglish.pdf (accessed Mar. 1, 2008); Falvey Memorial Library, Villanova University, "VuFind," www.vufind.org (accessed Mar. 1, 2008).

12. See the Health Level Seven (HL7) Clinical Decision Support working committee activities, in particular the Infobutton Standard Proposal at www.hl7.org/Special/committees/dss/index.cfm and the NISO Metasearch Initiative documentation at www.niso.org/workrooms/mi (accessed Mar 1, 2008).

13. National Center for Biotechnology Information (NCBI) Entrez cross-database search, www.ncbi.nlm.nih.gov/Entrez (accessed Mar. 1, 2008).

Index to Advertisers

ACRL	5	LITA	cover 2, cover 3
ALCTS	15	Jaunter	cover 4