Linking OpenCourseWares and Open Education Resources: Creating an Effective Search and Recommendation System

Brett E. Shelton, Joel Duffin, Yuxuan Wang, and Justin Ball

Abstract—With a growing number of digital libraries and other open education repositories being made available throughout the world, effective search and retrieval tools are necessary to access the desired materials that surpass the effectiveness of traditional, all-inclusive search engines. This paper discusses the design and use of *Folksemantic*, a platform that integrates OpenCourseWare search, Open Educational Resource recommendations, and social network functionality into a single open source project. The paper describes how the system was originally envisioned, its goals for users, and data that provides insight into how it is actually being used. Data sources include website click-through data, query logs, web server log files and user account data. Based on a descriptive analysis of its current use, modifications to the platform's design are recommended to better address goals of the system, along with recommendations for additional phases of research.

Keywords—Digital libraries, open education, recommendation system, social networks

I. INTRODUCTION

DUE to the exponential growth of educational-related data available via web sites, it is often tedious for a student or teacher to explore and discover relevant items through standard methods. Well-studied and commercialized search engines like Google will often help users to find what they are seeking. However, if those searching do not know exactly what they are looking for, or they do not know the "proper" words to describe what it is that they want, the searching results returned are often unsatisfactory. Further, most search engines do not take into account personalized information—such as known preferences or proclivities of the individual searcher--and thereby produce the same result for users with different interests. For example, if the learner's interest is different from the mainstream, the search result will be less meaningful to the learner. For these reasons, personalized recommendation systems have drawn the attention of e-business web sites as a way to improve user satisfaction and retention [1][2]. Those searching for specific educational resources, and the owners of repositories of educational resources, face the same challenges as those within the e-business realm. Certainly, the archiving, retrieving and social nature of Internet-based educational resources is of great concern and the focus of study to those within instructional technology realms as well as library studies, cognitive and behavioral studies, and learning sciences [3][4].

Open education is a movement to increase learning opportunities worldwide by using technology to provide free access to learning resources [5]. Issues of importance to the movement include collection development, search, adoption incentives, personalization, licensing, sustainability, localization, learner support, lifelong learning, and accreditation. *Folksemantic* is an open source platform that was developed as part of a National Science Foundation (NSF) project to build tools that could help connect OpenCourseWares (OCWs) with open education resources (OERs) in the National Science Digital Library (NSDL). *Folksemantic* supports open education by providing widgets and APIs that collections can use to easily connect their resources to related resources outside of their collections [6]. *Folksemantic* also aggregates metadata from many open educational resource collections and provides simple interfaces for searching across those collections.

The *Folksemantic* project was motivated by a desire to increase cross-pollination between larger sized OCWs and smaller OERs. The inception and initial growth of the OCW movement, which makes course materials for university courses freely available online, was made largely possible through significant funding by the Hewlett foundation [7]. Through programs such as the NSDL, the NSF has invested in the development of a large number of high quality OERs for K-16 such as simulations, videos, animations, student activities, and teacher resources [8]. The premise of the NSDL *Folksemantic* project was that adding recommendations to OERs and OCWs could help connect learners and educators to additional resources to meet their needs.

Since the inception of the *Folksemantic* project, researchers have aggregated metadata from many OCW and OER repositories, created a search portal, built a content-based recommendation system that learns from usage data and widgets, and created APIs that third parties can use to retrieve recommendations and search results. Additionally, social network functionality was added that allows people to sign up, create a profile, register feeds that they produce such as their blogs and bookmarks, connect to other users, and view

Brett E. Shelton is with Utah State University, Logan, UT 84322 USA (phone: 435-797-2393; fax: 435-797-2693; e-mail: brett.shelton@usu.edu).

Joel Duffin, is with Utah State University, Logan, UT 84322 USA (e-mail: joel.duffin@usu.edu).

Yuxuan Wang is with is with Utah State University, Logan, UT 84322 USA (e-mail: yuxuan.wang@aggiemail.usu.edu).

Justin Ball, is with Utah State University, Logan, UT 84322 USA (e-mail: justinball@gmail.com).

activity in the system. Third party projects including the Open University's OpenLearn and Utah State University's OpenCourseWare project integrated OER search and recommendations into their websites [9]. A popular platform for hosting OCWs called eduCommons added an option that allows a system administrator to easily turn on *Folksemantic* recommendations for an eduCommons site. Ongoing efforts focus on system evaluation, aggregating additional resources, increasing adoptions by collections, and refining recommendation algorithms to personalize them to individual users.

The goals for *Folksemantic*, in parallel with creating a meaningful user experience for search and access to digital library and OCW resources, include:

• Establish the needed connections between OCWs and key digital libraries, such as the NSDL.

• Address the issue of requiring a unique algorithm for searching these resources and presenting meaningful findings.

• Gather pertinent data in which to evaluate the effectiveness of the algorithm to check how well the tool serves the intended audiences.

• Address the idea of adoption—how can the tool be proliferated among resource partners and users?

II. INITIAL EFFORTS AT CUSTOMIZED SEARCH TOOLS

A. OCW Finder

OCW Finder was initially a static light-weight client-side OCW search tool, created early in the OCW movement. It attracted attention because of its simple multi-column tag interface for browsing and searching across multiple OCW collections [10]. *Folksemantic* extended the functionality of OCW Finder by adding the ability for people to register new collections and by putting a database and search engine behind the finder so that as new courses were found they would be immediately available to search. Figure 1 shows the current OCW Finder interface.

B. OER Recommender

OER Recommender is a content-based recommendation system that recommends related resources based on the semantic relatedness of their metadata. Collection metadata is harvested from RSS feeds, OAI endpoints, and other types of data sources using the ROME and other open source libraries. The system uses Lucene to index, search resource metadata and to calculate resource similarities. The recommendation algorithm takes resource title, tags, and descriptions into account, weighting each differently. It uses standard TF/IDF term vector techniques to calculate similarity of resource pairs. The system tracks user clicks and time on page data and uses it to adapt the ordering of recommendations based on this user data [11].
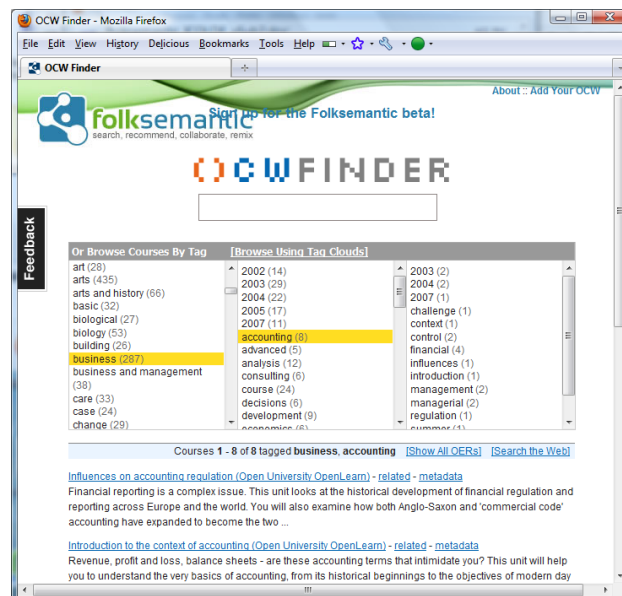

Fig. 1 Current OCW Finder interface

The simplest way to add *Folksemantic* OER recommendations to a web page is to place a small snippet of HTML in the web page. Additionally, recommendations can be retrieved in RSS, OAI, HTML, JSON, and XML formats. Figure 2 shows an example of a web page that has included the OER Recommender HTML snippet.
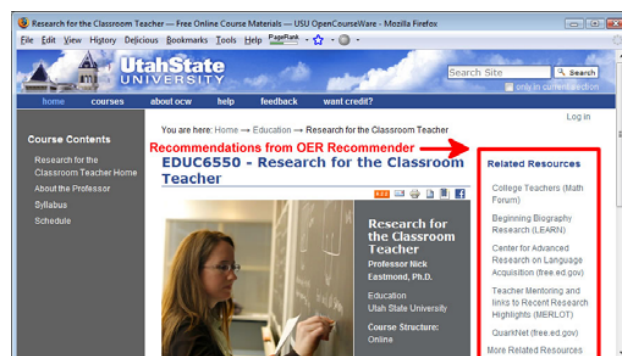

Fig. 2 An example of integrating OER Recommender into a web page

To encourage adoption of the OER Recommender service by third party OER collections, a Greasemonkey [12] script was developed that can be used to easily demonstrate what it would be like to integrate the OER Recommender service into a website. The script includes the ability to retrieve recommendations calculated in real-time for resources not already in the *Folksemantic* index.

C. Learning from User Feedback and Dealing with Ordering Bias

Recommender tracks user clicks and time on page in order to use that implicit feedback to improve recommendations. Because users are naturally inclined to click on results that

appear higher in a list of results, it can create bias in systems that learn from click data. The OER Recommender system attempts to address this potential bias by grouping recommendations into categories and randomly ordering items in those categories. The system begins by identifying the top 20 most semantically related items. It then divides these into two groups called highly recommended and recommended. The highly recommended group contains items with relevance scores at least one standard deviation higher than the average relevance score of the entire 20 items. Before an item has sufficient user feedback data gathered, OER Recommender displays items in the highly recommended category first by randomly ordering them, followed by items in the recommended category which are also randomly ordered. The system uses user feedback data to develop a popular category. The popular category contains items that have at least one standard deviation higher than the average number of clicks. Because the popular category represents user preference, the system displays items in that category before items in the highly recommended category and in strict order of user preference.

III. FOLKSEMANTIC

In previous projects, *Folksemantic* researchers developed Web2.0 tools to increase the impact of open education by supporting human interaction around OERs [3]. In order to facilitate automated and human to human personalization of OERs, *Folksemantic* researchers decided to integrate OCW Finder and OER Recommender and social network components from previous projects into a single platform. Key aspects of the integration effort were to (1) merge the underlying index used by both OCW Finder and OER Recommender into a single index, (2) translate the OCW Finder interface into additional languages, (3) transition OCW Finder to use Lucene for its underlying search and browse functionality, (4) add a user account system to allow people to sign up and create profiles, and (5) make it possible for users to register feeds they produce such as blogs and bookmarks so they can be displayed on user profiles and to use to personalize recommendations.

A. Collection Aggregation

To facilitate adding collections to the *Folksemantic* index, an interface was developed that allowed registered users to submit collections. When system administrators and trusted users submit feeds for collections, the system immediately begins harvesting metadata from those feeds. When other users submit feeds, system administrators are notified and must approve the feed before the system begins harvesting metadata from it. The *Folksemantic* aggregator supports harvesting metadata from OAI endpoints, RSS feeds, and SQI sources.

B. Personalized Recommendation Algorithm

Folksemantic researchers have conducted research on personalized recommendations and begun to add personal recommendation capabilities to the Folksemantic system. An effective personalized recommendation system automatically tracks the behavior of users (in profiles), analyzes those behaviors and based on the analysis results, designates different interests for different users. The engines then recommend relevant items to users based on their inferred interests [13]. To simplify, personalized recommendation systems map items to users based on the analysis of the users' profiles. Different users receive different recommendation lists. To create recommendations, the system calculates recommendation scores for assigned items it recommends to users. The higher the recommendation score, the higher the item appears in the recommendation list. Finally, a method is employed to shrink the recommendation list to the appropriate length, often referred to as "top-N method" or "threshold."

Profiles used for personalized recommendations are usually the ratings that the users give to the items. For instance, Youtube.com employs a "5 star" rating system for videos while The Internet Movie Database (IMDB) uses a "10 star" system for rating movies. During the last decade researchers have given significant attention to efforts at designing efficient personalized recommendation algorithms [14][15]. One element researchers seem to have ignored is the importance of the quantization of those ratings. The common strategy for recommendation algorithms is to equalize the ratings into the same interval. To address this weakness, Folksemantic researchers have explored a method that maps the user ratings to discriminative weights. Experimental results showed that the new method dramatically improved the performance of the personalized recommendation system. To demonstrate the generalizability of the discriminative weights approach and to compare the results in different algorithms, the approach was applied to two personalized recommendation strategies. The first algorithm it was applied to is the widest applied Collaborative Filtering (CF [16]) method and the other one is the recently developed Bipartite Graph Projection method (BG [17]). The new algorithm, when applied to the recommendation of educational resources may impact how communities form around educational interests and applications as well as increase the effective use of educational materials.

The groundwork has been laid to finish implementing and to deploy personalized recommendations in *Folksemantic*, personalized recommendations will be displayed on the users dashboard and emailed to them periodically if they choose that option. The basis for the personalizing recommendations will be user attention metadata. User attention metadata used by the system will include:

• Identity feeds – RSS feeds that users register that represent their interests. For example, their blog or their delicious account.

• Clicks – The articles that the user clicks on.

• Shares – The articles that the user shares to others.

• Comments – Articles that the user comments on.

• Time on page – Amount of time that a user spends on an article before moving on.

• Searches – Searches the user executes.

C. Recommendation Assumptions

Some of the assumptions of the system include:

• Semantic relatedness – The more semantically similar an article is to articles that a user has paid attention to, the more interesting to the user.

• Attention types – Different types of attention should be given different weights. For example, following a link to an article should not give it as much weight as writing the article.

• Attention details – The particulars of a given type of attention might make it more important than another attention of the same type. For example, if a person shares an article with 100 people, it might be reasonable to infer that it is more important than an article that they share with one person.

• Entry recency – The more recently an article has been added to the system, the more interesting to the user (they probably have not seen it before).

• Attention recency – The more recently a user has showed attention to an article, the more weight that should be given to it.

• Attention frequency – The more frequently a user has showed attention to an article, the more weight that should be given to it.

The system takes into account the difference between relevance and certainty by weighting different types of attention metadata differently. For example, while an item that a user clicks on may be more relevant than an blog article they have written, it is harder to be certain of that and so the system gives the click less weight than the article.

D. Recommendation Score

Articles are scored by considering the factors of:

• relevance
• attention type
• attention details
• attention recency
• article recency

For all articles that a user has paid attention to, the system creates a list of 20 "related articles" using this algorithm. It ranks the scores and caches the top 20 (that the user has not already clicked on) to recommend to the user. This approach will serve as a starting place for exploring personalized recommendations in the context of the Folksemantic system.

IV. Data Sources and analysis

To facilitate analysis researchers implemented consistent logging and analysis methods across the *Folksemantic*, OER Recommender, and OCW Finder web sites. All three web sites share the same application database and Lucene index. In order to gather data from which to analyze the Folksemantic system, researchers installed Google Analytics, implemented custom query logging, and archived standard web server log files. In addition, researchers used the web application database as a data source for analysis. Application database data analyzed included click tracking and time on page data

that the application uses to improve recommendations. The web application click tracking approach omits multiple clicks on the same recommendation in a given user session, an approach people might use to try to "game the system".

To understand how people are using *Folksemantic*, researchers analyzed data web server log files, custom logging, Google Analytics, and the application database. Researchers gathered Google Analytics data by accessing the reports it provides. Web server log files were analyzed by writing custom scripts that matched and counted specified patterns. The scripts were then applied to archived log files to mine the desired data. Scripts identified and omitted requests from web crawlers such as the Googlebot crawler. The web application database was analyzed by writing custom SQL queries and generate reports. The earliest month that statistics are given for is August 2009 because that is when the OER Recommender and OCW Finder websites were modified to use the same database and index and the *Folksemantic* website was launched (see Table 1).

V. Conclusion

Researchers have begun adding personal recommendation capabilities to the recommender systems. Previously, the system recommended "related resources" based on semantic similarities, and adapted rankings based on click- and time-on-page data, but did not personalize recommendations. New efforts focus on personalizing recommendations to individuals based on metadata gathered about users' attentions. The new system allows a user to register, create a profile, list their bookmark, blog, and comment and share resources with other learners. With data from these sources, in addition to click-and-search metadata, the system generates personalized recommendations, thereby creating the potential for increasing the effective locating and use of educational resources. This paper assists in efforts to increase the impact of open education by providing informational technology perspectives and insight into search infrastructure for online educational resources.

To gain greater understanding of how *Folksemantic* is being used and how it can be improved, researchers plan to conduct user tests, interviews, focus groups, and surveys. Researchers also plan to add more collections, invite registered collections to integrate *Folksemantic* recommendation and search results into their websites, complete and deploy personalized recommendation algorithms, implement OER rank and incorporate it into search and recommendation ranking algorithms, and conduct additional research on recommendation algorithms.

To provide good recommendations, an index needs to have enough high quality resources to cover the topics that interest users. The *Folksemantic* index currently includes 7,160 courses from 46 OpenCourseWare repositories and a total of 87,941 resources from 628 OER repositories. Researchers believe that they can increase the number of courses it indexes dramatically by focusing on including additional existing

OCW repositories. Even while the *Folksemantic* interface has been translated into 8 languages and the algorithms are implemented for language specific index, search, and recommendation, the majority of indexed resources are in English. Again, broadening to include different languages is an area where focused effort will increase the number of available resources, resulting in greater utility for different language users and repository holders. Researchers are part of the GLOBE worldwide learning objects organization that recently reported that its indexes contain over 320,000 records that *Folksemantic* researchers are evaluating for addition into the *Folksemantic* index.

Additional research will include the analysis of user videos using *Folksemantic* captured through a collection strategy of UserFly videos. UserFly videos will be analyed by coding randomly sampled videos. Videos will note unique activities observed and issues recognized. Each video will be classified according to the type of activities that are observed. Items that could be improved in the system will be noted.

Researchers will continue their efforts to complete and deploy personalized recommendation algorithms in the *Folksemantic* system. Personalized recommendation algorithms will incorporate advancements from research on algorithms that combine discriminative weights with a novel sparse matrix clustering method, and modeling users interests using multiple term vectors (one for each interest) by extracting vectors from closely related (clustered documents) based on user attention.

Another planned improvement is to implement an "OER Rank" measure of resource quality that will be taken into account when ranking recommendations. Similar to PageRank [18] OER will take into account factors such as overall inbound links, size, and availability of the resource. In addition, relevance scores will be analyzed and cutoff level will be established to avoid making poorly scored recommendations just because better recommendations are not available.

REFERENCES

[1] Brusilovsky, P., Kobsa, A., & Nejdl, W, (Eds.). (2007). *The adaptive web methods and strategies of web personalization*. Berlin: Springer.

[2] Chu, W., & Park, S. T. (2009). Personalized recommendation on dynamic content using predictive bilinear models. In *Proceedings of the 18th international World Wide Web Conference.*

[3] Johnson-Henson, S. (2008). The *folksemantic* web: Tools for a human-centered approach to learning. *International Conference of the Learning Sciences (ICLS) 2008*, Utrecht, Netherlands.

[4] Seow, P. (2008). Towards a framework for seamless learning environments. In *Proceedings of the International Conference of the Learning Sciences (ICLS) 2008*. Utrecht, Netherlands.

[5] Cape Town Open Education Declaration (2008) Read the declaration — Retrieved November 13, 2008, from www.capetowndeclaration.org/read-the-declaration

[6] Duffin, J. & Muramatsu, B. (2008). OER Recommender: Linking NSDL Pathways and OpenCourseWare Repositories. 2008 Joint Conference on Digital Libraries, Pittsburgh, PA.

[7] Tom Caswell, Shelley Henson, Marion Jensen, & David Wiley (February – 2008). Open Educational Resources: Enabling universal education. Retrieved March 8, 2010 from http://www.irrodl.org/index.php/irrodl/article/view/469/1001

[8] Zia, Lee L. "The NSF National Science, Mathematics, Engineering, and Technology Education Digital Library (NSDL) Program: A Progress Report." D-Lib Magazine, October 2000. Available at http://www.dlib.org/dlib/october00/zia/10zia.html

[9] Muramatsu, B., Ball, J., Duffin, J., Wiley, D., Gray, J., Godwin, A., Reedstrom, R., Cooper, C. & Whipkey, C. (June, 2008). Open Educational Resources Interoperability Sprint Outcomes: Linking OER Providers. Paper presented at OpeniWorld: Europe 2008, Lyon, France.

[10] Wiley, D. (2005, July 13). OCW Finder. iterating towards openness. Blog. Retrieved March 31, 2010 from http://opencontent.org/blog/archives/175

[11] Duffin, J., Muramatsu, B. & Henson Johnson, S. L. (2007). OER Recommender: A recommendation system for open educational resources and the National Science Digital Library. White paper funded by the Andrew W. Mellon Foundation for the Folksemantic.org project. (Updated March 7, 2008).

[12] Boodman, A., Greasemonkey (2007). https://addons.mozilla.org/en-US/firefox/addon/748

[13] Mobasher, B., & Anand, S. S., (2005). Intelligent techniques for web personalization. In *Intelligent Techniques for Web Personalization*. Bamshad Mobasher and Sarabjot Singh Anand (Eds.), Lecture Notes in Artificial Intelligence (3169) , (pp. 1-37). Springer.

[14] Belkin, N. J. (2000). Helping people find what they don't know. *Communications of the ACM,* 43(8), 58-61.

[15] Montaner, M., López, B., & De la Rosa, J. L. (2003). A taxonomy of recommender agents on the Internet. *Artificial Intelligence Review*. 19(4), 285-330.

[16] Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). *GroupLens*: An open architecture for collaborative filtering of NetNews. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*. ACM Press, New York, 175–186.

[17] Zhou, T., Ren, J., Medo, M., & Zhang, Y. C. (2007). Bipartite network projection and personal recommendation. *Physics Review E.,* 76.

[18] Brin, S. & Page, L. (1998). The Anatomy of a Large-Scale Hypertextual Web Search Engine. Proceedings of the 7th international conference on World Wide Web (WWW). Brisbane, Australia. pp. 107–117. http://dbpubs.stanford.edu:8090/pub/1998-8.

TABLE I
FOLKSEMANTIC RESOURCES, ADOPTERS, AND USERS

Measure	Period	Total, Daily Average, Daily Std Dev
OpenCourseWare Collections †	Mar 2010	46
Total Collections †	Mar 2010	628
Indexed Courses †	Mar 2010	7160
Indexed Resources (All OERs) †	Mar 2010	87941
Languages †	Mar 2010	8
Searches £	Mar 2010	**Aggregate:** 10,377; 346, 100 *Folksemantic:* 4,258, 142, 83 **Recommender:** NA **Finder:** 6,119, 204, 75
Websites Using *Folksemantic* Widgets ‡	Mar 2010	18
Visitors (Unique IPs) *††	Aug 2009 Mar 2010	1893 8140
Requests (hits) ‡	Mar 2010	**Aggregate**: 874,613; 29,154; 9,097 *Folksemantic*: 316,329; 10,544; 6,226 **Recommender**: 501,270; 16,709; 4,269 **Finder**: 57,014; 1,900; 452
Registered Users †	Mar 2010	1763
New Users †	Jul 2009 Aug 2009 Sep 2009 Oct 2009 Nov 2009 Dec 2009 Jan 2010 Feb 2010 Mar 2010	908 226 131 105 66 57 94 98 73
Logged in Users ‡	Mar 2010	**Aggregate**: 142, 4.7, 2.6 *Folksemantic:* 108, 3.6, 2.2 **Recommender:** 34, 1.1, 1.2 **Finder:** NA
Recommendation Lists Displayed ‡	Mar 2010	**Aggregate:** 437,890; 14,596; 3588 *Folksemantic:* 33,666; 1,122; 1,077 **Recommender:** 404,224; 13,474; 476 **Finder:** NA
Clicked Recommendations †	Mar 2010	47974 (11%), 3198, 838
Average Time on Page †	Mar 2010	46 seconds

For each measure, four sets of statistics are presented: (a summary of all 3 *Folksemantic* websites, followed by statistics for the *Folksemantic*, OER Recommender, and OCW Finder websites, listed in that order).
* Data was gathered from Google analytics. Note that this does not include visits to websites that display recommendations using *Folksemantic* widgets.
† Data was gathered from the application database.
‡ Data was gathered from the web server log files.
£ Data was gathered from custom query log files.
††The maximum number of visitors in a single day was on July 8, 2009 when the sites had 600 unique visitors in response to being highlighted on a nationally syndicated technology radio show.