"Readers who borrowed this also borrowed…":
Recommender Systems in UK libraries

1. Recommender Systems

The emergence of the internet as an unparalleled information source has led to a number of significant and well documented problems constituting significant areas of research in fields as diverse as economics, sociology, computer and information science, and psychology. One relatively recent development, and one exercising academics in all of these fields, has been the recognition of so called 'information overload,' the difficulty users face navigating and processing the information available online. It has been noted that humans' capacity to find information advances more slowly than the pace at which new information is made available (Cosley et al., 2003), and the current exponential growth of the world wide web, which offers an increasingly vast and heterogeneous repository of information, presents significant challenges to users in many types of online environments. Recommender systems, applications that attempt to suggest useful content to users, have emerged as an important means of addressing these challenges, and now constitute a significant area of research in the field of Information Science.

Many of the issues that proved a catalyst to the development and implementation of recommender systems on the wider web can be seen to affect the library community. The huge growth in available resources, particularly through the increased provision of electronic journals and e-books, has led to a wealth of accessible information with the potential to inform scholarship and facilitate teaching and research, while the provision of that information over the web has dramatically increased the range and scope of services available to users (Webster et al., 2004). Despite these undoubted benefits, however, it has become increasingly clear that traditional models of information retrieval frequently fail to best connect library users with potentially relevant material. As Baez et al. note, the sheer volume of available resources often demands that searches are narrowed to identify core resources rather than broadened to take in potentially useful but perhaps ill-defined items (Baez et al., 2010). This is exacerbated by the increasing preference of users, born out of practice and familiarity, for self guided web-based searches, bypassing the traditional role of the librarian or information professional (Webster

et al., 2004). Furthermore, in an increasingly competitive information market place, institutions are becoming ever more conscious of the need for their services to match the expectations of users in both appearance and functionality. The commercial application of recommender systems in e-commerce environments offered a means of presenting specifically those products deemed potentially attractive to a user, leading to improved browser-to-buyer rates, extended cross selling opportunities and increased customer loyalty, while also reducing the time and effort spent by consumers on searching (Schafer et al. 1999 , Hervas-Drane, 2007). It is not difficult to see how these benefits might be paralleled in library settings.

This paper examines the current state of recommender systems in the context of UK library OPACs. A brief overview of recommender system technology is given, and two key issues in systems development are discussed. Three existing systems offering recommendations within the OPAC are introduced, and the results of a recent comparison of UK library systems are presented. Finally this paper introduces a collaboration between the University of Sheffield and OCLC that aims to develop a recommender system for WorldCat.org - the aggregated catalogue of OCLC member libraries (http://www.worldcat.org).

2. Types of Recommender Systems

Recommender systems "analyze user profiles, content items, and the connections between them, and try to predict future user behaviour" (Prekopcsák, 2007, 8). This process results in a presentation of suggested content to the user. Such systems differ from more traditional forms of information retrieval (for example search engines) in the sophistication with which an item's potential utility is calculated, and the extent to which the system explicitly attempts to add value though a personalised approach (Burke, 2002). While the earliest development of recommender systems occurred in non-commercial environments, and the roots of the theories underpinning their implementation can be found in areas such as cognitive science, approximation theory and information retrieval (Adomavicius & Tuzhilin, 2005), it was the emergence in the late 1990s of e-commerce as a key growth area that drove many of the most significant advances in recommender systems.

While recommender system uses and applications are diverse, the recommender problem essentially can be reduced to the question of how best to forecast a rating value for an item that a user has not seen (Adomavicius & Tuzhilin, 2005). Two main models of recommender system design have emerged to address this problem. Content-based (CB) systems can be linked closely with traditional information

retrieval research, and approach recommendation as a user-specific classification problem (Burke, 2007). CB models create profiles of users inferred from their previous interaction with the system, and link these profiles to static data about items. The effectiveness of CB systems relies on the richness of available data about content, and the amount of information they can obtain about the users (O'Donovan & Smyth, 2005). Since recommendations in a CB system are based solely on past user preferences, CB systems are most likely to recommend things similar to those preferred by users in the past, rather than suggest new or unusual items. The major alternative to the CB system is Collaborative Filtering (CF). CF Systems also establish some form of user profile, but generate recommendations based on the preferences of other users identified as having similar preferences. CF systems have a number of advantages for users; they offer a rich means of discovering new items, obtaining advice about a selected item, and connecting with other like-minded users. Research also indicates that CF systems tend to offer greater recommendation diversity than other models, making serendipitous discovery more likely and presenting users with unfamiliar but potentially valuable niche material (Burke, 2007).

Three other models also have been suggested. Knowledge-Based (KB) systems incorporate a third layer of data representing domain knowledge, allowing the system to infer functional links between the user's needs and items that might fulfil them (Resnick & Varian 1997). Utility-Based (UB) systems make recommendations based on a computation of utility. While in some case this might equate to the most practically useful item for the user, the utility to other parties also can be incorporated (e.g. the profitability of items might influence recommendations on an e-commerce site) (Shani et al., 2005). Finally Demographic recommender systems attempt to categorise users according to their personal attributes, making recommendations according to extrapolated demographic classes (Resnick & Varian 1997).

As might be expected, an understanding of the unique and complementary benefits of different types of systems has led to the development of Hybrid applications – systems which utilise at least two of the models outlined above in order to generate recommendations. Burke (2007) characterises 53 flavours of hybrid systems, and evaluates 41 of these in a laboratory setting, identifying significant disparities in recommendation type between hybrid models and concluding that system designers must carefully tailor hybrid type to system function. It has been suggested that any attempt to develop the recommender systems field will certainly utilise hybrid systems (Yager, 2003).

3. Data Acquisition and Recommender Systems

A key consideration for recommender system developers is what data to use as a basis for recommendations, and how the system collects that data. Methods for acquiring data can be broadly categorised as 'implicit' or 'explicit' (or sometimes 'extensional' or 'intentional') (Yager, 2003). Implicit methods gather data from standard engagement with the system, for example purchasing a product or withdrawing a book, and are not noticed by the user. Explicit methods are those requiring the active participation of the user, for example rating an item (Prekopcsák, 2007). Implicit methods have the advantage of making no demands on the user, but the relationship of the data gathered with optimal ratings may be imprecise. One area of research has examined the role of temporal information, particularly the time spent viewing a page or item, as a means of inferring ratings. Parsons et al. (2004) note that a correlation between viewing time and preference has been broadly established in other fields, and such a method has proven effective as a means of augmenting e-commerce recommender systems (Lee et al., 2008). Other web usage data (for example click- though rates and query log analysis) also has been successfully applied in this context (Cho et al., 2002).

Explicit methods are potentially much more valuable to recommender systems, since they usually represent an unambiguous statement of preference (Chen & Pu, 2007). The most common type of explicit data collection comes in the form of user ratings. These can be either unary (e.g. "like"), binary ("good" or "bad"), or integer based (a Likert numeric scale) (Schafer et al., 2007). Since users are generally assumed to avoid or seek to reduce cognitive effort, collecting such data was assumed to be problematic, because the user must be persuaded to actively participate in the process (Gretzel & Fesenmaier, 2005). This has led some to observe that the drive for increased recommendation accuracy must be balanced with the levels of user effort required (Rashid et al., 2008). Recent research, however, has indicated that users are frequently willing to engage with ratings processes, motivated by the opportunity to improve their profile (and, therefore, recommendations they receive), express themselves, and help or influence others (Herlocker et al., 2004). Researchers also have examined the ratings process itself, establishing that users prefer more detailed or "fine-grained" scales (Cosley et al., 2003), while Gretzel & Fesenmaier (2005) suggest that the structure, content and layout of a ratings process can substantially influence a user's perception of subsequent recommendations. Sinha & Swearingen (2002) posit that systems should seek to elicit ratings at particularly opportune moments, suggesting a 'conversational and collaborative' model by which additional rating requests can be

triggered at moments when the user is assumed to have particular motivation to do so (for example when the user is surprised by a particularly high or low rating). Chen & Pu (2007) expand this conversational element to create an ongoing dialogue with the user to refine and develop ratings and ratings ranking. The development of unobtrusive yet comprehensive interfaces for the gathering of ratings has been identified as a crucial area for future research in recommender systems (Perugini et al., 2004).

4. Explanation and Recommender Systems

The early recommender system has been characterised as a 'black box,' offering the user no information about how suggestions were computed (Schafer et al., 2007). The result is a system more akin to a search engine, and far removed from the type of word-of-mouth recommendations that users are accustomed to receiving (Bonhard & Sasse, 2006). Recent research has indicated that users interact far more effectively with systems that clearly explain the relationship between ratings and recommendations. This was demonstrated most clearly by Sinha & Swearingen's (2002) findings that users wanted an explanation even for items they were recommended and already liked. Herlocker et al. (2000) also have suggested four key benefits to incorporating an explanation function into a recommender system: justification (helping users understand why a recommendation has been made), user involvement (making users feel more involved in the process), education (helping users better understand the scope of the system) and acceptance (giving the users greater confidence in recommendations, thereby mitigating trust issues discussed above). They also identify that explanations offer a means of handling recommendation errors, be they process or data based.

While explanations are beneficial in principle, it also has been shown that the form of the explanation is also important (Tintarev & Masthoff, 2007). Some explanation methods trialed by McSherry (2005) were found to have a negative effect - particularly those of a technical nature. By contrast the second most effective form of explanation consisted solely of a statement of past performance (e.g. 'the system has predicted correctly for you 80% of the time'). Other research has experimented with a confidence display, essentially a graphic representation of the system's faith in the recommendation (McNee et al., 2003). It seems clear that the extent to which users trust recommendations is heavily influenced by the extent to which they understand how those recommendations have been generated.

5. Recommender Systems and the OPAC

Moves towards more interactive services on the wider web have led libraries to add an increasing range of functionality to their OPACS. Features such as tagging, reviews and recommendations offer libraries a means not only of engaging with their users but also supplementing traditional means of metadata generation, thereby further enriching their catalogue. While the authors were able to find no study focusing solely on recommendation functions in library OPACs, a number of researchers have attempted to analyze the prevalence of "next generation" OPAC features. Yang and Hofmann (2011) sampled 260 of the approximately 2,500 academic libraries in North America, evaluating each institution's OPAC against a list of features including "recommendations / related material." They found that 34% of OPACs offered some form of "recommendation language," although their definition of such language was broad, including functions such as "Find more by this author" and "Nearby items on shelf." Taking a systems rather than institutional approach, Yang and Wagner (2010) reviewed twenty OPAC discovery tools, finding that thirteen provide a recommendation function – although it should be noted that no precise definition of what was deemed to constitute a recommender system was given. Mercun and Zumer (2008) compared the functionality of six selected library catalogues (five of which had been recently modernized) to that of the Amazon website. Two features were included in the study relating to recommendations: i) "New items, most popular, recently returned items and recommendations lists" and ii) "Personalised recommendations." While none of the library OPACs studied were found to incorporate the latter feature, three were judged as "very good" in the former category, with one rated "good with some limitations." However the broad scope of feature i) makes it difficult to determine exactly what those libraries provide. Tam et al. (2009) analyzed the prevalence of ten next-generation features in 153 UK University OPACs, including "Borrowing Suggestions." They found only three examples of this feature in the library catalogues.

Comment [SW1]: Paragraph moved from the "Comparison" section

A number of different recommender systems models have seen practical implementation within library OPACs , three of which will be discussed here. The University of Huddersfield has developed an in-house system (see Figure 1) that mines circulation logs to present users viewing an item on the catalogue with suggested alternatives (Pattern, 2008). Their approach matches the item being viewed in the OPAC with a list of all the library users who have withdrawn that item. The full circulation records of those users are then collated, and the most commonly withdrawn items are presented back to users. Users are initially presented the three recommendations, with an option to view more suggestions, or the entire list of

relevant items. Even with this relatively simple method, circulation has increased significantly since the system's introduction (Pattern, 2009), although it should be noted that further analysis is required before the precise impact of the system can be gauged.

Figure 1: Borrowing Recommendations at the University of Huddersfield

A more sophisticated approach is taken with the BibTip project (Monnich & Spiering, 2005 & 2006). They suggest that the use of circulation data alone is insufficient for a comprehensive recommender system, since reference items cannot be properly integrated, and the availability of items to be loaned is not considered. BibTip therefore uses three software agents to build recommendations based on implicitly inferred data. An *Observation Agent* identifies all the titles selected within each discrete user session, passing the data to an *Aggregation Agent* which collates co-occurrences and builds links between items. Finally a *Recommendation Agent* mines these links to provide recommendations to users in the form of hypertext links in the browser (see Figure 2). At each stage statistical evaluations of the data are employed to refine connections between users and titles viewed. Users of the prototype system at the

University of Karlsruhe rated the quality of the recommendations as 4.21 on a 1-5 Likert scale, and the system is now employed by a number of German Universities.

Figure 3: Example of BibTip Recommendations (text translates to "What is also interesting, perhaps...")

A third model for providing OPAC recommendations can be found with *Library Thing For Libraries* (LTFL). *Library Thing* is an online service allowing members to catalogue their book collections, and supplement this catalogue with ratings, reviews and tags. With more than a million members, this represents a significant amount of explicit feedback for collaborative filtering recommendations. LTFL is sold to libraries as an OPAC overlay, using a small piece of Javascript to query *Library Thing*'s database for the ISBN of the item being viewed in the OPAC. Relevant tags, reviews and recommendations are then exported to the OPAC interface and can be viewed by the user (see Figure 3). Since the overlap of the Library Thing aggregated catalogue with University OPACS is around 50%, this offers a readymade means of adding significant value to a large part of the corpus (Westcott et al., 2009). While no research

yet has focused specifically on the utility of the LTFL recommendations, Mendes et al. compared the resources discovered by users through the LTFL tag cloud with those found using traditional Library of Congress Subject Headings. They determined that "for every new book a user discovers using LCSH headings they will discover 4 books using LTFL" (Mendes et al., 2009, 10).

Figure 3: Library Thing for Libraries Recommendations

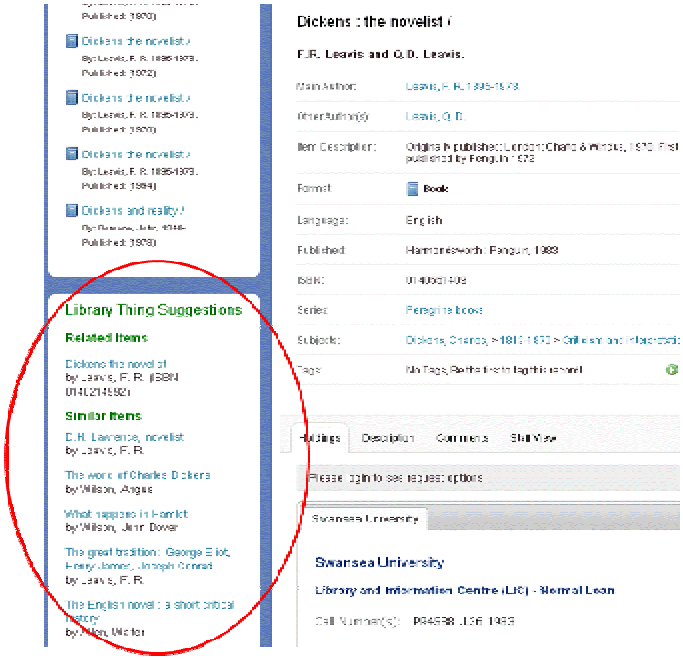

6. A Comparison of UK OPACs

6.1 Method

To determine the extent to which recommender systems have permeated the UK library environment, OPACs from the 118 UK universities ranked in the 2011 Guardian University League Table were analysed, along with those of the 211 local councils appearing on the UK Public Libraries Website, with the aim of identifying whether any form of recommendation was offered to users. Each OPAC was accessed and keyword searches run using the terms "accountancy" and "Dickens." The item pages for the top five search results were then viewed. For the purposes of this study, a recommendation was characterized as the presentation to the user of one or more specific resources not directly related to the ranking of the search results, and not requiring the use of a faceted browsing function (e.g. "find more by this author"). The analysis also noted the type of recommendation (collaborative filtering or content based), when the recommendations were presented to users, and the language used to introduce and explain the recommendation. In addition, three other widely discussed next-generation features were assessed; user-generated reviews, user-generated tags, and user ratings. It was hoped that this additional data might provide some context for the recommendation results. Finally the LMS software used by each institution was noted, along with any discovery overlay.

Two minor limitations of the comparison should be mentioned. All the OPACs were accessed as a guest – i.e. without any form of log-in – and it is therefore conceivable that certain features of some OPACs were not visible. It should also be noted that a number of libraries in both the public and academic sector are in the process of migrating to or testing new OPACs or discovery tools, meaning that visitors are given the option of choosing which catalogue to access. Where this was found to be the case, the newest OPAC was viewed. It is however possible that in accessing catalogues through publically accessible websites, some newer implementations may have been missed.

6.2 Results

Table 1: Results of UK OPAC Comparison

Feature	Public Library OPACs (of 211)	University Library OPACs (of 118)
Recommendations	4 (2%)	13 (11%)
Tags	4 (2%)	19 (16%)
Reviews	53 (25%)	12 (10%)
Ratings	12 (6%)	4 (3%)

Table 1 offers a broad overview of the results of the analysis, showing that the number of OPACs offering recommendations to library users is extremely low. Of the thirteen university OPACs found to feature recommendations, eight were examples of the Encore discovery tool, which offers a list of items new to the library that are related to the user's search terms (see Figure 4). This constitutes a purely content-based approach, since recommendations are based on pre-existing metadata (a combination of subject areas and keywords, and acquisitions data). These recommendations are presented at the search results page, rather than at item level, and there was some variation across the eight OPACs in how the recommendations were introduced, with text varying from "Recently Added" to "Recently Added: people looking for this may also want…" Only two libraries offer recommendations based on some form of collaborative filtering – the University of Huddersfield, which has developed its own recommendation system based on circulation data (as described above), and the University of Swansea, which has incorporated the Library Thing for Libraries overlay.

In the public library domain, three organisations offer collaboratively filtered recommendations through their Vubis OPAC software (see Figure 5), with recommendations presented using language similar to commercial recommendations applications: "Readers who borrowed this also borrowed…", "Top ten suggestions based on borrower choices", "Click to see what people reading this also read." In all, only 5 of the 329 OPACs viewed incorporate recommendations based on a collaborative filtering model (see

Figure 4: Encore "Recently Added" Recommendations

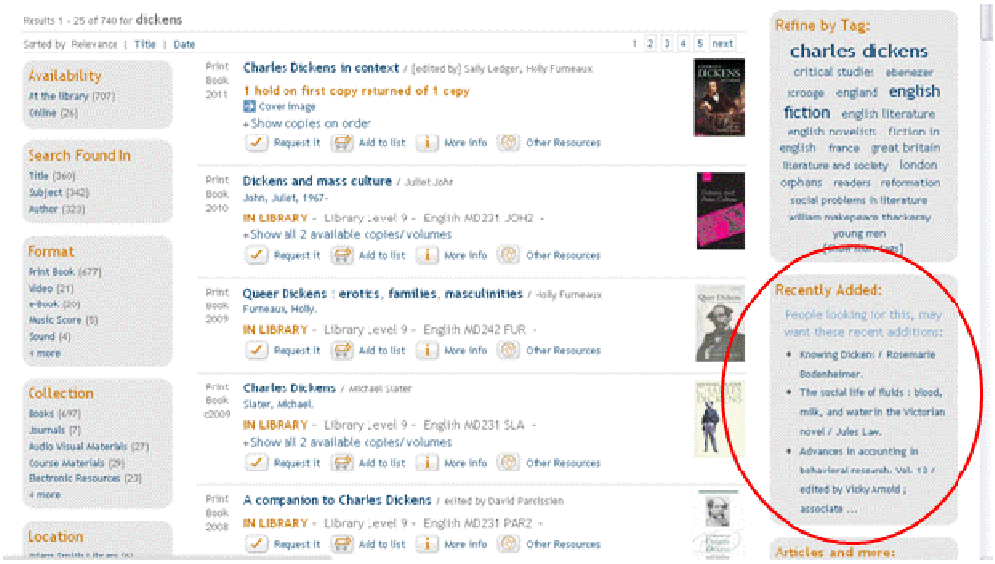

Figure 5: Vubis Recommendations

Table 2), although in one case it proved impossible to accurately determine the method used. No OPAC was found to offer any form of detailed explanation about how recommendations were generated.

Table 2: Types of Recommendation

	Type of Recommendation			
	Content Based	Collaborative Filtering	Unknown	Total
Public Libraries	1 (25%)	3 (75%)	0 (0%)	4 (100%)
University Libraries	10 (77%)	2 (15%)	1 (8%)	13 (100%)

Given the relatively small numbers of features found, it is not surprising that a majority of both public and university OPACs demonstrate none of the four functions examined (see table 3). In the university library domain, the OPACs featuring recommendations averaged a total of 2.15 features, with 10 out of the 13 institutions in question having at least one other function available to users. For public libraries this average drops to 1.5, with three of the four OPACs that show recommendations having no other features. Only two OPACs – one public and one university – offer both recommendations and user generated ratings.

Table 3: Number of features found

Number of Features	Public Libraries (of 211)	University Libraries (of 118)
0	149 (71%)	93 (79%)
1	52 (25%)	7 (6%)
2	9 (4%)	13 (11%)
3	1 (0%)	5 (4%)
4	0 (0%)	0 (0%)

6.3 Discussion

The results of the comparison indicate that the incorporation of recommendations into the OPAC is extremely limited. The instances of recommendations at UK academic institutions (11%) fall far below Yang and Hofmann's figure of 34%. While it seems plausible that North American libraries have been quicker to implement next-generation features, it is also suggested that the inclusion of functions such as "See more by the same author" within the category of "Recommendations" has falsely inflated the occurrence figures – and that the instances of true recommendations within OPACs are low on both sides of the Atlantic. There are a number of potential reasons for this. Many libraries, both public and academic, are bound to LMS contracts, and as such have limited opportunities for developing advanced functionality on legacy systems. While some libraries have opted to add a discovery interface to their existing LMS, there is naturally an incremental cost associated both with the technology itself and the resources required to implement and maintain it. Furthermore some discovery tools themselves do not support recommendations. Although the University of Huddersfield has demonstrated that in-house implementations are possible, not all institutions have the staff expertise to follow suit. The comparison also indicated that the potential for generating CF based recommendations is limited by the scarcity of ratings functions, and the likelihood of users accessing the OPAC without logging in. This means that any CF recommendations must rely on implicitly generated data, which in turn complicates the development of truly powerful recommendation algorithms.

The discrepancies between the results in the University and Public sectors are also worthy of comment. While the overall prevalence of recommender systems is low for both groups, it is perhaps surprising that a higher proportion of University libraries are found to offer the feature. This is not the case for other features, with both reviews (25% v 10%) and ratings (6% v 3%) more widely found in the public library sector. It must be noted that the incidences of all functions are low, and the results might be explained in part by the differing market shares of competing Library Management Systems (with differing "standard" features) within each sector. Nevertheless, these results do suggest that public libraries place a greater emphasis on features that assist resource evaluation (ratings and reviews) than discovery (tags and recommendations), with the opposite true of academic libraries.

Comment [SW2]: Added discussion of public v academic sector results

Library professionals now seem to broadly accept the potential benefits to users of next generation OPAC features, but are perhaps yet to be persuaded that the effectiveness of recommendation functions is great enough to warrant their implementation becoming a priority. Although the results of

this study, when compared to Tam et al.'s 2009 findings, show a rise of almost 10% in University Libraries that offer recommendations (from 2% to 11%), this increase can be explained almost entirely by the eight Encore OPACs that recommend recently added items. While this is potentially a useful service, a content-based approach that employs acquisition date as the sole arbiter of utility might be seen as a relatively primitive form of recommendation. The prevalence of more advanced recommender systems (those able to build on past behaviour and the collective preferences of users) in commercial environments suggests such systems can add value, and the examples of OPAC recommendation systems outlined above appear to provide a useful and respected service.

7. Conclusion and Future Research

While our research demonstrates the limited uptake of recommender systems within library OPACs, these results alone do not yet allow for broad conclusions about the role such systems can or should play as a means of aiding resource discovery. Fundamentally, a persuasive case for prioritising recommender system implementation in the OPAC has yet to be made. For those institutions offering recommendations, it remains to be established exactly how users utilise suggestions, and what precise value they offer to the various stakeholders. How can different information-seeking contexts be reconciled within a single system, particularly when searches are conducted without users logging in? Are recommendations more suitable for academic or public libraries, and are different types of recommendation more suitable for each sector? In the absence of ratings and universal user profiles, what data might best be used to determine context and generate recommendations? Do content-based systems, which align more closely with traditional library services, offer a better fit than collaboratively filtered recommendations? When and where should recommendations appear? How can we best measure the value of a library recommendation? These questions form the basis of a research project currently underway at the University of Sheffield. In partnership with OCLC we are taking a user-centered approach to the design of a recommender system for a library catalogue – in this case the aggregated catalogue of OCLC member libraries known as WorldCat (http://www.worldcat.org). Through extensive qualitative research combined with log and circulation records analysis, we hope to better establish the role that recommendations can and should play in a user's functional relationship with an OPAC. From this platform we will seek to place users' interactions with recommendations within a broader Information Behaviour framework, and finally develop a prototype system that most

appropriately serves what is undoubtedly a diverse user base. It is also anticipated that this project has the potential to inform wider debates about the role of adaptive technologies and information seeking within the library space.

References

Adomavicius, G. and Tuzhilin, A. (2005), "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions", *Transactions on Knowledge and Data Engineering*, Vol. 17 No. 6, pp. 734-749.

Baez, M., Birukou, A., Casati, F. and Marchese, M. (2010), "Addressing Information Overload in the Scientific Community", *Internet Computing*, Vol. 14 No. 6, pp. 31 – 38.

Bonhard, P. and Sasse, M. A. (2006), "'Knowing me, knowing you' — using profiles and social networking to improve recommender systems", *BT Technology Journal*, Vol. 24 No. 3, pp. 84-98.

Burke, R. (2002), "Hybrid Recommender Systems: Survey and Experiments", *User Modeling and User-Adapted Interaction*, Vol. 12 No.4, pp. 331-370.

Burke, R. (2007), "Hybrid Web Recommender Systems", in Brusilovsky, P., Kobsa, A., and Nejdl, W. (Eds.), *The Adaptive Web*, Springer-Verlag, Berlin, pp. 377-408.

Chen, L. and Pu, P. (2007), "Preference-based Organization Interfaces: Aiding User Critiques in Recommender Systems", *Lecture Notes in Computer Science*, Vol. 45 No. 11, pp. 77-86.

Cho, Y. H., Kim, J. K., Kim, S., H. (2002), "A personalized recommender system based on web usage mining and decision tree induction", *Expert Systems with Applications*, Vol. 23 No. 3, pp. 329-342.

Cosley, D., Lam, S. K., Albert, I., Konstan, J. A., and Riedl, J. (2003), "Is Seeing Believing? How Recommender Interfaces Affect Users' Opinions", *Proceedings of the SIGCHI conference on Human factors in computing systems, Ft. Lauderdale,FL*, ACM, New York, NY, pp. 585-592.

Gretzel, U. and Fesenmaier, D. (2005), "Persuasion in Recommender Systems", *International Journal of Electronic Commerce*, Vol. 11 No. 2, pp. 81-100.

Harden, S. & Harden, R. (1994), *UK Public Libraries* [online] (updated 25th March, 2008), available at http://ds.dial.pipex.com/town/square/ac940/ukpublib.html (accessed, 20 April, 2011).

Herlocker, J. L., Konstan, J. A., and Riedl, J. (2000), "Explaining Collaborative Filtering Recommendations", *Proceedings of the, 2000 ACM conference on Computer supported cooperative work, Philadelphia, PA*, ACM, New York, NY, pp. 241-250.

Herlocker, J. L., Konstan, J. A., Terveen, L. G. and Riedl, J. (2004), "Evaluating Collaborative Filtering Recommender Systems", *ACM Transactions on Information Systems*, Vol. 22 No 1, pp. 5-53.

Hervas-Drane, A. (2007). "Word of Mouth and Recommender Systems: A Theory of the Long Tail." working paper 07-41, NET Institute, Universitat Pompeu Fabra, Barcelona, Spain, 12 August, available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1025123## (accessed 14 March, 2011).

Lee, T. Q., Park, Y., and Park, Y-T. (2008), "A time-based approach to effective recommender systems using implicit feedback", *Expert Systems with Applications*, Vol. 34 No. 4, pp. 3055–3062.

McNee, S. M., Lam, S. K., Guetzlaff, C., Konstan, J. A. and Riedl, J. (2003), "Confidence Displays and Training in Recommender Systems", *Proceedings of the 9th IFIP TC13 International Conference on Human Computer Interaction, Zurich, Switzerland*, IOS Press, Amsterdam, The Netherlands, pp. 176-183.

McSherry, D. (2005), "Explanation in Recommender Systems", *Artificial Intelligence Review*, Vol. 24 No. 2, pp. 179–197.

Mendes, L. H., Quinonez-Skinner, J. and Skaggs, D. (2009), "Subjecting the catalog to tagging", *Library Hi Tech*, Vol. 27 No. 1, pp.30 – 41.

Mercun, J. and Zumer, M. (2008), "New generation of catalogues for the new generation of users: a comparison of six library catalogues", *Program*, Vol. 42 No. 3, pp. 243-261.

Monnich, M. and Spierling, M. (2008), "Adding Value to the Library Catalog by Implementing a Recommendation System", D-Lib Magazine, Vol. 14 No. 5/6, available at: http://www.dlib.org/dlib/may08/monnich/05monnich.html (accessed 15 March, 2011).

O'Donovan, J. & Smyth, B. (2005). "Trust in Recommender Systems", *Proceedings of the 9th International Conference on Intelligent User Interfaces, 19-12 January, 2005, San Diego, CA*, ACM, New York, NY, pp. 167-174.

Parsons, J., Ralph, P., and Gallager, K. (2004), "Using Viewing Time to Infer User Preference in Recommender Systems", *Proceedings of the AAAI Workshop on Semantic Web Personalization, 25-26 July,, 2004, San Jose, CA*, available at: http://maya.cs.depaul.edu/~mobasher/swp04/accepted/parsons.pdf (accessed 10 March, 2011).

Pattern, D. (2005), "Using "circ_tran" to show borrowing suggestions in HIP", blog, 17 November, available at: http://www.daveyp.com/blog/archives/49 (accessed 3 April, 2011).

Pattern, D. (2009), "The impact of book suggestions / recommendations?", blog, 29 March, available at: http://www.daveyp.com/blog/archives/729 (accessed 3 April, 2011).

Perugini, S., Goncalves, M. A., and Fox, E. A. (2004), "Recommender Systems Research: A Connection-Centric Survey", *Journal of Intelligent Information Systems*, Vol. 23 No. 2, pp. 107–143.

Prekopcsák, Z. (2007). *Content organization and discovery: state-of-the-art and new ideas for P2P-Fusion*, Budapest University of Technology and Economics, Budapest, Hungary, available at: http://prekopcsak.hu/papers/preko-2007-p2pfusion.pdf (accessed 10 April, 2011).

Rashid, A. M., Karypis, G., and Riedl, J. (2008), "Learning Preferences of New Users in Recommender Systems: An Information Theoretic Approach", *SIGKDD Explorations*, Vol. 10 No. 2, pp. 90-100.

Resnick, P. and Varian, H. R. (1997), "Recommender Systems", *Communications of the ACM*, Vol. 40 No. 3, pp. 56-58.

Schafer, J. B., Frankowski, D, Herlocker, J. and Sen, S. (2007), "Collaborative Filtering Recommender Systems" in Brusilovsky, P., Kobsa, A., & Nejdl, W. (Eds.), *The Adaptive Web*, Springer-Verlag, Berlin, pp. 291–324.

Schafer, J. B., Konstan, J., and Riedl, J. (1999), "Recommender Systems in E-Commerce", *Proceedings of the 1st ACM Conference on E-Commerce, 3-5 November, Denver, CO,* ACM, New York, NY, pp. 158-166.

Shani, G., Heckerman, D., and Brafman, R. I. (2005), "An MDP-Based Recommender System", *Journal of Machine Learning Research*, Vol. 6, pp. 1265-1295.

Sinha, R. and Swearingen, K. (2002). "The Role of Transparency in Recommender Systems.", *Proceedings of the Conference of Human Factors in Computing Systems,, 20-25 April,, 2002, Minneapolis, MN*, ACM, New York, NY, pp. 830-831.

Tam, W., Cox, A. and Bussey, A. (2009), "Student user preferences for features of next-generation OPACs: a case study of University of Sheffield international students", *Program*, Vol. 43 No. 4, pp. 349-374.

The Guardian (2010), *University guide, 2011: University league table* [online], available at: http://www.guardian.co.uk/education/table/2010/jun/04/university-league-table (accessed 5 April, 2011).

Tintarev, N. and Masthoff, J. (2007), "A Survey of Explanations in Recommender Systems", *Proceedings of the, 2007 IEEE 23rd International Conference on Data Engineering Workshop, Istanbul, Turkey*, IEEE Computer Society, Washington, DC, pp. 801-810.

Webster, J., Jung, S. and Herlocker, J. (2004), "Collaborative Filtering: a new approach to searching digital libraries", *New Review of Information Networking*, Vol. 10 No. 2, pp. 177 – 191.

Westcott, J., Chappell, A. and Lebel, C. (2009), "LibraryThing for libraries at Claremont", *Library Hi Tech*, Vol. 27 No. 1, pp.78 – 81.

Yager, R. R. (2003), "Fuzzy logic methods in recommender systems", *Fuzzy Sets and Systems*, Vol. 136 No. 2, pp. 133–149.

Yang, S. Q. and Hofmann, M. A. (2011), "Next Generation or Current Generation? A Study of the OPACs of 260 Academic Libraries in the United States and Canada", *Library Hi Tech*, Vol. 29 No. 2.

Yang, S. Q. and Wagner, K. (2010), "Evaluating and comparing discovery tools: how close are we towards next generation catalog?", *Library Hi Tech*, Vol. 28 No. 4, pp.690 – 709.