Research Repository UCD

Title	The altruistic searcher
Authors(s)	Smyth, Barry, Coyle, Maurice, Briggs, Peter
Publication date	2009-08
Publication information	Smyth, Barry, Maurice Coyle, and Peter Briggs. "The Altruistic Searcher." IEEE Computer Society, 2009.
Conference details	Paper presented at the International Symposium on Social Intelligence and Networking (SIN09) in conjunction with the 2009 IEEE International Conference on Social Computing, 29th - 31st August 2009, Vancouver, Canada
Publisher	IEEE Computer Society
Item record/more information	http://hdl.handle.net/10197/1663
Publisher's version (DOI)	10.1109/CSE.2009.202

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

The Altruistic Searcher

Barry Smyth, Maurice Coyle, Peter Briggs

CLARITY: Centre for Sensor Web Technologies
School of Computer Science and Informatics,
University College Dublin,
Belfield, Dublin 4,
Ireland
Email: {firstname.lastname}@ucd.ie

Abstract—Recently researchers have argued that the prevailing view of Web search, as a solitary activity, is flawed: that, in reality, Web search can be an inherently collaborative task. In this paper we describe and evaluate an approach to collaborative Web search that seeks to enhance mainstream search engines by harnessing the past search experiences of communities of like-minded searchers in order to adapt the result-lists of traditional search engines so that they reflect the niche interests of community members.

I. INTRODUCTION

Notwithstanding the success of mainstream search engines, recent research has highlighted how even the leading search engines suffer from low success rates when it comes to delivering relevant results to the average searcher. In large part this problem is as much due to the searcher as it is the search engine: our search queries tend to be vague and under-specified, and rarely provide a clear indication of our search needs [1]–[5]. This means that Web search is far less efficient than it should be: recent studies suggest that among information workers 10% of salary costs are lost due to wasted search time [6]. In this regard there are at least two active areas of research that have the potential to improve upon contemporary Web search. They question some key assumptions of mainstream Web search and suggest important adaptations to conventional Web search approaches.

The first assumption concerns the *one-size-fits-all* nature of mainstream Web search — two different users with the same query will, more or less, receive the very same result-list, despite their different preferences. Recent work has argued that this approach is overly simplistic, and that Web search needs to become more *personalized* so that the implicit needs and preferences of searchers can be accommodated. There are many different approaches to personalizing Web search by harnessing different types of user preference and context information to influence the search experience; see for example [7]–[17].

The second assumption that is increasingly questioned concerns the *solitary nature* of Web search. By and large Web search takes the form of a isolated interaction between individual searcher and search engine. However, recent research has suggested that there are many circumstances where the search for information has a distinctly collaborative flavour,

with groups of searchers (e.g., friends, colleagues, classmates) cooperating in various ways as they search for and share results. We will describe recent work in the area of *collaborative information retrieval*, which attempts to capitalize on the potential for collaboration during a variety of information seeking tasks; see for example, [18]–[25].

In this paper we will focus on the second of these lines of research – collaborative information retrieval – to describe an approach to community-based Web search that harnesses the latent collaboration that exists in many search scenarios. We propose that it is possible to deliver a much more collaborative form of Web search in which individual searchers play a more altruistic role as their search experiences are harnessed to the benefit of other searchers within a given search community. Specifically we will describe a particular approach to community-based Web search, known as *collaborative Web search*, in which the search experiences of a community of searchers are pooled and used as the basis of recommendations during future searches for related queries. We go on to describe a recent live-user trial which integrates collaborative Web search with Google so that searchers can benefit from community recommendations while they search as normal.

II. MOTIVATIONS

Failed searches are often due to the mismatch between the query-space of the searcher and the document-space of the search engine index: users are prone to submit vague queries that often contain terms that are different from those used to index documents. The result is a significant *vocabulary gap* that leads to poor search performance. To understand the scale of this vocabulary gap, we submitted just under 7,700 queries to the three leading search engines (Google, Yahoo, and MSN) to locate a particular target page for each query; admittedly a particularly tough measure of relevance, but one that was straightforward to objectively measure. We evaluated the effectiveness of each search engine in terms of the percentage of times that the target page was retrieved in the top-ten results returned. The results, (Figure 1) highlight how all 3 search engines struggle to perform, at best retrieving the target results in their top-ten less than 14% of the time. Indeed it is interesting to note how all three search engines perform best for queries with 3 terms, suggesting that modern

Fig. 1. The vocabulary gap in Web search: while, initially, longer queries may help to guide more successful searches, beyond a certain query size (typically 3+ terms) searchers frequently use terms that are less helpful in retrieving what they need.

search engine technology has been optimized for typical query lengths [26]. Importantly, we can also see how retrieval effectiveness increases, as query size grows from 1 to 3 terms, supporting the view that search engine performance can be improved if users provide more detailed queries. However, this is true only to a point. For queries with more than 3 terms we see a decline in retrieval effectiveness because these extra terms are often less than helpful when it comes to identifying a target document; for example, users frequently chose very specialised terms that do not even occur in the target document.

There are many scenarios in which search can be viewed as a community-oriented activity. For example, the employees of a company will act as a type of search community with overlapping information needs. Similarly, students in a class may serve as a search community as they search for information related to their class-work. Visitors to a themed website (e.g., a wildlife portal or a motoring portal) will tend to share certain niche interests and will often use their site's search facilities to look for related information. And of course, groups of friends on a social networking site may act as a community with shared interests.

We are interested in these emergent *search communities* because of the potential for patterns to exist between the search behaviours of community members. For example, Figure 2 shows the results of a 17-week study of the search patterns for 70 employees of a local software company. During the study we examined more than 20,000 individual search queries and almost 16,000 result selections. Just over 65% of queries shared at least 50% (> 0.5 similarity threshold) of their query terms with at least 5 other queries; and more than 90% of queries shared at least 25% of their terms with about 10 other queries. Thus, searchers within this ad hoc corporate search community do search for similar things in similar ways, much more so than in generic search scenarios, where we typically find much lower repetition rates of about 10% at

Fig. 2. An analysis of query overlap among a community of work colleagues highlights a high degree of query overlap.

the 0.5 similarity threshold [26]. This is an important result, which is supported by similar studies on other communities of searchers [26]. It tells us that, in the context of communities of like-minded searchers, the world of Web search is a repetitive and regular one. It hints at a type of *social search knowledge* that is generated from the search experiences of individuals as they search. This suggests that it may be possible to harness this search knowledge by facilitating the sharing of search experiences among community members in a way that supports individual searchers. In this way, as individuals search, for their own information gain, their search patterns have the potential to help other searchers in the future who will likely search for similar information in similar ways.

III. A REVIEW OF COLLABORATIVE INFORMATION RETRIEVAL

The more social perspective of Web search alluded to above has motivated researchers to consider the implications of a more collaborative form of information retrieval and Web search; see, for example, [18]–[21]. Indeed, recent work by [27] highlights the inherently collaborative nature of more general purpose Web search. For example, during a survey of just over 200 respondents, clear evidence of collaborative search behaviour emerged. More than 90% of respondents indicated that they frequently engaged in collaboration at the level of the *search process*. For example, 87% of respondents exhibited “back-seat searching” behaviours, where they watched over the shoulder of the searcher to suggest alternative queries. Thus, despite the absence of explicit collaboration features from mainstream search engines there is clear evidence that users implicitly engage in many different forms of collaboration as they search, although, as reported by [27], these collaboration “work-arounds” are often frustrating and inefficient. Naturally, this has motivated researchers to consider how different types of collaboration might be supported by future editions of search engines.

The resulting approaches to *collaborative information retrieval* can be usefully distinguished in terms of two important

dimensions, *time* — *synchronous* versus *asynchronous* search — and *place* — *co-located* versus *remote* searchers. Co-located systems offer a collaborative search experience for multiple searchers at a single location, typically a single PC (e.g. [25], [28]) whereas remote approaches allow searchers to perform their searches at different locations across multiple devices; see e.g. [22]–[24]. The former enjoy the obvious benefit of an increased faculty for direct collaboration that is enabled by the face-to-face nature of co-located search, while the latter offer a greater opportunity for collaborative search. Alternatively, synchronous approaches are characterised by systems that broadcast a “call to search” in which specific participants are requested to engage in a well-defined search task for a well defined period of time; see e.g. [28]. In contrast, asynchronous approaches are characterised by less well-defined, ad-hoc search tasks and provide for a more open-ended approach to collaboration in which different searchers contribute to an evolving search session over an extended period of time; see e.g. [22], [29].

A good example of the co-located, synchronous approach to collaborative Web search is given by the work of [25]. Their CoSearch system is designed to improve the search experience for co-located users where computing resources are limited; for example, a group of school children having access to a single PC. CoSearch is specifically designed to leverage peripheral devices that may be available (e.g. mobile phones, extra mice etc.) to facilitate distributed control and division of effort, while maintaining group awareness and communication. The CoSearch interface also provides various opportunities for users to associate notes with pages. Interesting pages can be saved and as users collaborate a *search summary* can be created from the URLs and notes of saved pages. In the case where observers have access to mobile phones, CoSearch supports a range of extended interface features to provide observers with a richer set of independent functionality via a bluetooth connection. In this way, observers can download search content to their mobile phone, access the page queue, add pages to the page queue and share new pages with the group.

The work of [30] focuses on image search tasks using a table-top computing environment. Once again, preliminary studies speak to the potential for such an approach to improve overall search productivity and collaboration at least in specific types of information access tasks. A variation on these forms of synchronous search activities is presented in [28], where the use of mobile devices as the primary search interface allows for a remote form of synchronous collaborative search. The iBingo system allows a group of users to collaborate on an image search task with each user using a ipod touch device as their primary search/feedback device (although conventional PCs appear to be just as applicable). Interestingly, where the focus on CoSearch is largely on the division of search labour and communication support, iBingo offers the potential to use relevance feedback from any individual searcher to the benefit of others.

Remote search collaboration (whether asynchronous or syn-

chronous) is the aim of SearchTogether, which allows groups of searchers to participate in extended shared search sessions as they search to locate information on particular topics; see also [22]. In brief, the SearchTogether system allows users to create shared search sessions and invite other users to join in these sessions. Each searcher can independently search for information on a particular topic, but the system provides features to allow individual searchers to share what they find with other session members by recommending and commenting on specific results. In turn, SearchTogether supports synchronous collaborative search by allowing searchers to invite others to join in specific search tasks, allowing cooperating searchers to synchronously view the results of each others’ searches via a split-screen style results interface. As with CoSearch above, one of the key design goals in SearchTogether is to support a division of labour in complex, open-ended search tasks. In addition, a key feature of the work is the ability to create a shared awareness among group members by reducing the overhead of search collaboration at the interface level. SearchTogether does this by including various features such as integrated messaging, query histories, and recommendations arising out of recent searches.

Overall, the collaborative information retrieval systems we have so far examined have been largely focused on supporting collaboration from a division of labour and shared awareness standpoint, separate from the underlying search process. These systems have assumed the availability of an underlying search engine and provided a collaboration interface that effectively *imports* search results directly, allowing users to share these results. As noted by [31], one of the limitations of these approaches is that collaboration is restricted to the interface: while individual searchers are notified about the activities of collaborators, they must individually examine and interpret these activities in order to reconcile their own activities with their co-searchers. Thus, the work of [31] describes an approach to collaborative search that is more tightly integrated with the underlying search engine resource so that the operation of the search engine is itself influenced by the activities of collaborating searchers in a number of ways. For example, mediation techniques are used to prioritise, as yet, unseen documents, while query recommendation techniques are used to suggest alternative avenues for further search exploration.

IV. A MODEL OF COLLABORATIVE WEB SEARCH

In this paper we will focus on a community-based approach to *collaborative Web search* (CWS) in which the *asynchronous* search experiences of communities of like-minded *remote* searchers are harnessed to provide an improved search experience that is more responsive to the learned preferences of the community. In short, results are recommended during search that other community members have found to be useful for similar queries in the past. In this section we will describe a system architecture that facilitates integration with a mainstream search service. We will also present some sample screen-shots of the resulting integrated service and explain how results are recommended and promoted to searchers.

Fig. 3. The collaborative Web search architecture supports the maintenance of community-based search knowledge (H) as a source of promotions for some underlying search engine or as part of an enhanced meta-search service.

A. System Architecture

Earlier versions of CWS have been implemented as a form of meta-search engine, post-processing the results of some underlying search engine(s) and presenting these re-ranked results through a dedicated search interface [26], [32]. The problem with this approach has been the need to convince users to try an alternative search service and to learn a new search interface, which most users are unwilling to do.

The architecture presented here facilitates a very different form of integration by using a proxy-based approach (see Figure 3) to intercept queries on their way to the underlying search engine, and manipulating the results that are returned from this engine back to the searcher. Users continue to use their favourite search engine in the normal way, but community recommendations/promotions are incorporated into the result-lists directly via the proxy.

Briefly, the *proxy agent* (currently a standard Squid¹ proxy) operates to transparently filter URL requests from all configured browsers. In this case, the proxy redirects Google requests (e.g. query submissions, next page requests, etc.) to the CWS server. For example, consider a user U_i submitting query q_T to Google. This request is redirected to the CWS whereupon two things happen. First, the query is passed on to Google and the result-list R_S is returned in the normal way. Second, in parallel the query is also used to access the CWS *hit-matrix* for U_i 's community to generate a ranked set of promotion candidates, R_P , as outlined above. These promotion candidates are annotated by the *explanation engine* (see Section IV-B) to present the searcher with a graphical representation of their community history. Result-lists R_P and R_S are merged and the resulting list R_{final} is returned to the user; typically this merge involves promoting the k (e.g., $k = 3$) most relevant promotions to the head of the result-list.

¹<http://www.squid-cache.org>

B. The CWS Interface

To better explain the end-user experience under CWS we present Figures 4 and 5 as example screen shots for the result-list returned for the query 'Michael Jordan'. In the case of Figure 4 we see the default Google result-list, with results for the basketball star clearly dominating. In Figure 5, however, we see a result-list that has been modified by our proxy-based version of CWS, trained by a community of artificial intelligence and machine learning researchers. The results are presented through the standard Google interface but now we see that the top 3 results are promotions for the well-known Berkeley professor.

Fig. 4. The result-list returned by Google in response to the query 'michael jordan', at the time of writing.

In addition to recommending particular results for the community there are further benefits to be gained by exposing the searcher to the search history of a promoted result, to help them better understand the reasons behind its promotion. Thus, we introduce an explanation engine to the CWS architecture, which annotates promoted results with one or more explanation icons designed to capture different aspects of the result's community history. Specifically, three basic types of explanation icon are used: *popularity* icons convey the overall selection popularity of a result among the target community; *recency* icons indicate how recently the result has been selected by a community member and can help the searcher to judge how current a particular promotion might be; finally, *related queries* icons indicate whether the result in question has been previously selected by the community for any queries that are related (similar) to the current target query.

In Figure 5 we can see how each of the top 3 promotions are annotated with different subsets of the above icons, depending on the selection history of the result in question. Moreover, 'mousing-over' an explanation icon will reveal additional explanatory detail to the searcher; see Figure 5. For example, mousing-over the related queries icon will present a list of the most common related queries that have led to selections for the result in question. These related queries can themselves be selected by the searcher to initiate a new search.

Fig. 5. The result-list returned by CWS in response to the query 'michael jordan' issued within a community with a shared interest in computer science. The extra explanation information available by mousing over each promoted result icon type is also shown.

C. A Case-Based Approach to Recommendation

Ultimately the task of our collaborative Web search engine is to complement the results of an underlying search engine (such as Google in this example) with community search results that have been drawn from the relevant past search experiences of a community of like-minded searchers. To do this we adopt a case-based reasoning [33] perspective. Case-based reasoning is an approach to reasoning and problem solving that emphasises the role of past experiences during problem solving tasks. In short, rather than attempting to solve some new problem from scratch, for example by using some complex set of rules or inference model, a case-based reasoning system will base its solution on a modified version of a solution to a very similar problem that has already been solved. Thus, a case-based reasoning system relies on a collection of past problem solving experiences called a *case base*. Each problem experience is a *case* and the experience is represented as a combination of *problem specification* and *solution*. A case-based reasoning system will use similarity-based retrieval techniques to identify a relevant case and may then automatically adapt or refine its solution to fit the current target problem. This problem solving technique has enjoyed considerable success in many practical problem solving scenarios including design, planning, classification, and recommendations tasks.

In our model of collaborative Web search, the search history of a community is stored as a case base of search cases with each search case made up of a specification part and a solution part; see Equation 1. The *specification* part (see Equation 2) corresponds to a given query. The *solution* part (see Equation 3) corresponds to a set of selection-pairs; that is, the set of page selections that have been accumulated as a result of past uses of the corresponding query. Each selection-pair is made up of a result-page id and a hit-count for the number of times that the given page has been selected by community members

in response to the given query.

$$c_i = (q_i, (p_1, r_1), \dots, (p_k, r_k)) \quad (1)$$

$$Spec(c_i) = q_i \quad (2)$$

$$Sol(c_i) = ((p_1, r_1), \dots, (p_k, r_k)) \quad (3)$$

Given a new target query, q_T , CWS must identify a set of *similar* search cases from the community's search case-base. A standard term-overlap metric (Equation 4) is used to measure query-case similarity so that past search cases can be rank-ordered according to their similarity scores. Some or all of these similar cases may then be reused by promoting them in the result-list returned to the user.

$$Sim(q_T, c_i) = \frac{|q_T \cap Spec(c_i)|}{|q_T \cup Spec(c_i)|} \quad (4)$$

Consider a page, p_j , that is associated with query, q_i , in some search case, c_i . The relevance of p_j to c_i can be estimated by the relative number of times that p_j has been selected for q_i ; see Equation 5.

$$Rel(p_j, c_i) = \frac{r_j}{\sum_{r_m \in Sol(c_i)} r_m} \quad (5)$$

Then, the relevance of p_j to some new target query q_T can be estimated as the combination of $Rel(p_j, c_i)$ values for all cases $c_1, \dots, c_n$ that are deemed to be similar to q_T , as shown in Equation 6. Each $Rel(p_j, c_i)$ is weighted by $Sim(q_T, c_i)$ to discount the relevance of results from less similar queries; $Exists(p_j, c_i) = 1$ if $p_j \in Sol(c_i)$ and 0 otherwise.

$$WRel(p_j, q_T, c_1, \dots, c_n) = \frac{\sum_{i=1 \dots n} Rel(p_j, c_i) \cdot Sim(q_T, c_i)}{\sum_{i=1 \dots n} Exists(p_j, q_i) \cdot Sim(q_T, c_i)} \quad (6)$$

This weighted relevance metric, $WRel$, is used to rank-order search results from the community case-base that are

Fig. 6. Success rates by session.

Fig. 7. Success rates by promotion type.

promotion candidates for the new target query. The top ranked candidates are then listed ahead of the standard meta-search results to give R_T ; see [34] for further details on the search interface and result presentation.

V. EVALUATION

So far we have described the CWS technique for adapting the results of an existing search engine(s) to conform to the preferences of a community of searchers. In this section we will describe the results emerging from a recent CWS trial in a corporate context, pointing out how CWS helped employees to search more successfully as a result of the sharing of community search knowledge.

A. Trial Setup

The trial included 70 employees from a local Dublin software company where CWS was deployed for an initial period of 10 weeks as the primary search engine covering more than 12,600 individual search sessions. Essentially all requests for Google were temporarily redirected to the CWS server for the period of the trial. From a user perspective, the standard Google interface was adapted to accommodate CWS promotions. During the course of this initial 10 week trial approximately 25% of search sessions included CWS promotions. We refer to these as *promoted* sessions. The remaining 75% of search sessions carried the standard Google result-list. We refer to these as *standard* sessions. While it was not feasible to elicit direct relevance feedback from trial participants, one useful indicator of search performance is to look at the frequency of so-called *successful sessions*.

B. Promoted vs Standard Sessions

A search session is successful if at least one result is selected by the searcher. This is a very crude measure of performance — often result selections are good indicators of at least partial relevance but sometimes they are not — but the lack of any result selections is a good indication that no relevant results have been noticed. Importantly, when we analysed the success rates of trial search sessions we found marked differences between the promoted and standard sessions. For

example, Figure 6 shows an average success rate of $\sim 48\%$ for standard Google searches, compared to a success rate of $\sim 62\%$ for promoted sessions; a relative advantage due to CWS promotions of approximately 29%. In other words, when community promotions were available they were found to help users search more successfully.

C. Sharing Search Experiences

Sharing is a key theme in CWS; past search experiences are shared by community members through result promotions. These promotions can come from two very different sources. On the one hand a promotion may come from the past history of the current searcher: today I might search using a query that is similar to queries that I have used in the past, and I will receive promotions based on my own previous selection history. We call these *self promotions* and they are useful when it comes to helping searchers to *recover* results that they have previously encountered. On the other hand, a promotion may come from a different community member altogether. We call this a *peer promotion*, and when I receive peer promotions I am sharing the search experiences of other community members. Peer promotions are especially useful when it comes to helping me *discover* new results, and they potentially help me to draw on the experiences of more informed searchers within the community.

During the trial, it was possible to reconstruct information about the origin of promotions in order to investigate any differences between the behaviour of users when it came to sessions made up of self and peer promotions. The results of this analysis are revealing; see Figure 7. We see that promoted sessions containing only self promotions have an average success rate of just under 60%. By comparison, sessions with only peer promotions have a success rate of 66%, while *mixed* sessions, containing both self and peer promotions, have an average success rate of more than 70%. Clearly searchers do benefit from the search experiences of others. Indeed sessions containing peer promotions (Figure 8) have higher ‘click-thru’ rates (60-70%) than sessions containing only self promotions (30%).

Fig. 8. Click-thru rates by promotion type.

VI. CONCLUSIONS

We have considered how a more social perspective on Web search can lead to new ways to improve search quality. We have described our approach to collaborative Web search which seeks to harness the search experiences of communities of like-minded users. The results highlight the inherently collaborative nature of search communities with evidence about how the searches of an individual can help others to search more effectively.

As our research has taken shape, new opportunities have emerged to extend the basic CWS concept. For example, the work of [34] looks at how useful search communities can be automatically identified and how their promotions can be combined to deliver further improvements in search quality. Another issue concerns the susceptibility of CWS to spamming by malicious users. In this regard, recent research [35] quantifies how CWS offers some level of protection from such malicious users. Moreover, the work of [36] explains how a model of user-trust can be incorporated into CWS to offer further protection from malicious users by weighting community promotions according to the trustworthiness of the community members who originally selected them. Finally, it is worth highlighting the work of [37], which looks at profiling result selections according to their snippet terms, with significant improvements in promotion quality accruing to this richer representation format. In addition, this extended approach to CWS has led to a novel technique for generating community-based result summaries which have been shown to offer better precision/recall characteristics than more conventional summarisation techniques; see [38].

ACKNOWLEDGMENT

This material is based in part on works supported by ChangingWorlds Ltd., and on works supported by Enterprise Ireland and Science Foundation Ireland.

REFERENCES

- [1] A. Spink, D. Wolfram, M. B. J. Jansen, and T. Saracevic, "Searching the Web: the Public and their Queries," *Journal of the American Society for Information Science*, vol. 52, no. 3, pp. 226–234, 2001.

- [2] A. Spink, J. Bateman, and M. B. Jansen, "Searching heterogeneous collections on the web: behaviour of excite users," *Information Research: An Electronic Journal*, vol. 4, no. 2, 1998.
- [3] A. Spink and B. J. Jansen, "A study of web search trends," *Webology*, vol. 1, no. 2, p. 4, 2004. [Online]. Available: <http://www.webology.ir/2004/v1n2/a4.html>
- [4] S. Lawrence and C. L. Giles, "Accessibility of Information on the Web," *Nature*, vol. 400(6740), pp. 107–109, 1999.
- [5] B. Smyth, "A community-based approach to personalizing web search," *IEEE Computer*, vol. 40, no. 8, pp. 42–50, 2007.
- [6] S. Feldman and C. Sherman, "The High Cost of Not Finding Information," in *IDC White Paper*. IDC Group, 2000.
- [7] P.-A. Chirita, D. Olmedilla, and W. Nejdl, "Pros: A personalized ranking platform for web search," in *AH*, 2004, pp. 34–43.
- [8] L. A. Granka, T. Joachims, and G. Gay, "Eye-tracking analysis of user behavior in www search," in *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*. New York, NY, USA: ACM, 2004, pp. 478–479.
- [9] M. Speretta and S. Gauch, "Personalized search based on user search histories," in *WI '05: Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 622–628.
- [10] F. A. Asnicar and C. Tasso, "Ifweb: a prototype of user model-based intelligent agent for document filtering and navigation in the world wide web," in *Proceedings of the workshop "Adaptive Systems and User Modeling on the World Wide Web, Sixth International Conference on User Modeling"*, 1997, pp. 3–11.
- [11] Z. Ma, G. Pant, and O. R. L. Sheng, "Interest-based personalized search," *ACM Trans. Inf. Syst.*, vol. 25, no. 1, p. 5, 2007.
- [12] C. Makris, Y. Panagis, E. Sakopoulos, and A. Tsakalidis, "Category ranking for personalized search," *Data Knowl. Eng.*, vol. 60, no. 1, pp. 109–125, 2007.
- [13] B. Zhou, S. C. Hui, and A. C. M. Fong, "An effective approach for periodic web personalization," in *WI '06: Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 284–292.
- [14] P. A. Chirita, W. Nejdl, R. Paiu, and C. Kohlschütter, "Using odp metadata to personalize search," in *SIGIR '05: Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*. New York, NY, USA: ACM Press, 2005, pp. 178–185.
- [15] A. Pretschner and S. Gauch, "Ontology based personalized search," in *ICTAI '99: Proceedings of the 11th IEEE International Conference on Tools with Artificial Intelligence*. Washington, DC, USA: IEEE Computer Society, 1999, p. 391.
- [16] X. Shen, B. Tan, and C. Zhai, "Implicit User Modeling for Personalized Search," in *Proceedings of the Fourteenth ACM Conference on Information and Knowledge Management (CIKM 05)*, 2005.
- [17] J. Budzik and K. J. Hammond, "User interactions with everyday applications as context for just-in-time information access," in *IUI '00: Proceedings of the 5th international conference on Intelligent user interfaces*. New York, NY, USA: ACM, 2000, pp. 44–51.
- [18] M. C. Reddy and P. Dourish, "A finger on the pulse: temporal rhythms and information seeking in medical work," in *CSCW*, 2002, pp. 344–353.
- [19] M. C. Reddy, P. Dourish, and W. Pratt, "Coordinating heterogeneous work: Information and representation in medical care," in *ECSCW*, 2001, pp. 239–258.
- [20] M. C. Reddy and P. R. Spence, "Collaborative information seeking: A field study of a multidisciplinary patient care team," *Inf. Process. Manage.*, vol. 44, no. 1, pp. 242–255, 2008.
- [21] M. C. Reddy and B. J. Jansen, "A model for understanding collaborative information behavior in context: A study of two healthcare teams," *Inf. Process. Manage.*, vol. 44, no. 1, pp. 256–273, 2008.
- [22] M. R. Morris and E. Horvitz, "S³: Storable, shareable search," in *INTERACT (1)*, 2007, pp. 120–123.
- [23] —, "Searchtogether: an interface for collaborative web search," in *UIST*, 2007, pp. 3–12.
- [24] B. Smyth, P. Briggs, M. Coyle, and M. P. O'Mahony, "Google? shared! a case-study in social search," in *User Modeling, Adaptation and Personalization*. Springer-Verlag, June 2009.
- [25] S. Amershi and M. R. Morris, "Cosearch: a system for co-located collaborative web search," in *CHI*, 2008, pp. 1647–1656.
- [26] B. Smyth, E. Balfé, J. Freyne, P. Briggs, M. Coyle, and O. Boydell, "Exploiting query repetition and regularity in an adaptive community-based

- web search engine." *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, vol. 14, no. 5, pp. 383–423, 2004.
- [27] M. R. Morris, "A survey of collaborative web search practices," in *CHI*, 2008, pp. 1657–1660.
 - [28] A. F. Smeaton, C. Foley, D. Byrne, and G. J. F. Jones, "ibingo mobile collaborative search," in *CIVR*, 2008, pp. 547–548.
 - [29] B. Smyth, E. Balfe, J. Freyne, P. Briggs, M. Coyle, and O. Boydell, "Exploiting query repetition and regularity in an adaptive community-based web search engine," *User Model. User-Adapt. Interact.*, vol. 14, no. 5, pp. 383–423, 2004.
 - [30] A. F. Smeaton, H. Lee, C. Foley, and S. McGivney, "Collaborative video searching on a tabletop," *Multimedia Syst.*, vol. 12, no. 4-5, pp. 375–391, 2007.
 - [31] J. Pickens, G. Golovchinsky, C. Shah, P. Qvarfordt, and M. Back, "Algorithmic mediation for collaborative exploratory search," in *SIGIR*, 2008, pp. 315–322.
 - [32] B. Smyth, E. Balfe, O. Boydell, K. Bradley, P. Briggs, M. Coyle, and J. Freyne, "A Live-user Evaluation of Collaborative Web Search," in *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI '05)*. Morgan Kaufmann, 2005, pp. 1419–1424, edinburgh, Scotland.
 - [33] A. Aamodt and E. Plaza, "Case-based reasoning: Foundational issues, methodological variations, and system approaches." *AI Communications*, vol. 7, pp. 39–59, 1994.
 - [34] J. Freyne and B. Smyth, "Cooperating Search Communities," in *Proceedings of the 4th International on Adaptive Hypermedia and Adaptive Web-based Systems*, Dublin, Ireland, 2006.
 - [35] M. P. O'Mahony and B. Smyth, "Collaborative web search: A robustness analysis," *AI Review*, vol. Forthcoming, 2008.
 - [36] P. Briggs and B. Smyth, "Provenance, trust, and sharing in peer-to-peer case-based web search," in *ECCBR*, 2008, pp. 89–103.
 - [37] O. Boydell and B. Smyth, "Enhancing case-based, collaborative web search," in *Proceedings of the International Conference on Case-Based Reasoning*. Springer, 2007, pp. 329–343.
 - [38] —, "From social bookmarking to social summarization: an experiment in community-based summary generation," in *Intelligent User Interfaces*. ACM, 2007, pp. 42–51.