Expertise Recommendation: A triangulated approach

Debbie Richards, Meredith Taylor, Peter Busch
Computing Department, Macquarie University, North Ryde, NSW, 2109, Australia
{richards,mtaylor,busch}@ics.mq.edu.au

***Abstract*— Recommender systems are becoming increasingly popular as a means for bringing products to the attention of online users. Similarly, they offer a means by which scarce resources in the form of human experts can be identified and accessed. However, if the information in the system is missing, incorrect or obsolete, recommendations will not be followed or even sought in the first place. Relying on individuals to validate and update this information is problematic. To provide automated acquisition and maintenance of information regarding who has expertise and in what areas, we employ data mining techniques. However, data mining will not provide the full picture and thus our outputs are reviewed by the experts themselves, providing a second means of validation. The third part to our triangulated approach is the use of profiles and the gathering of feedback from both searchers and experts to ensure that recommendations provided are satisfactory to both parties.**

Keywords:* *expertise recommendation; recommender systems; data mining

1. Introduction

Given the increasing recognition that an organization’s most valuable resources are its people and the knowledge they hold, expertise location is becoming an important strategy to accessing and sharing that knowledge. In contrast to expert-systems, also known as knowledge-based systems, which seek to capture what it is that the expert knows so that it can be captured and reused, an expert/ise recommender system suggests who might know about what. The goal of the system is to point someone with a question to the person who has the appropriate knowledge. In the ideal situation the system provides a two-way communication channel connecting the knowledge holder and the knowledge seeker [1]. In some cases the inquirer’s main interest is in the answer, in other cases the main interest is to find an expert who will handle the problem [2]. Knowledge about the expert’s areas of expertise is needed for such a system. To discover this knowledge it is common to use data mining techniques [3]. Another alternative is to collect this information directly from the experts via self-reporting techniques (e.g. [4]). Individually both approaches have shortcomings.

Data mining relies on the existence of data which is, or is able to be, sufficiently structured to be used as input into one or more algorithms. This raises a number of issues: expertise could be identified from many different sources (e.g. publications, webpages, press releases, projects, grants, etc); these sources will vary across individuals and organizations; the format of these sources will vary across individuals and organization; most of these sources are free-format, unstructured and unclassified (a major hurdle if one wishes to use a supervised learning technique). Not only is the input to data mining an issue, each algorithm has its own strengths and limitations typically closely tied to the structure, amount and type of data and often dependent on the (type of) domain. Furthermore, the “knowledge learnt” tends to be restricted to types of output such as association rules, clusters or classification rules. Across domains, datasets and algorithms there is variation in the definition and identification of “interesting” concepts and validation of the output.

Due to these various limitations, an alternative to data mining and other automated searching techniques frequently used in recommender systems is the use of surveys/forms to be filled in by the domain expert which usually includes the selection of keywords relating to the individual’s areas of expertise. The problem is then reduced to matching the searcher’s query terms with the expert’s keywords. This technique is often referred to as a yellow-pages approach to finding an expert as that is the way people usually find a plumber, lawyer or doctor. It is a simple and yet effective method for finding people who have certain skills. It works on the expectation that, for instance, only someone with legal training and qualifications will list themselves as a solicitor and that since they listed themselves, they are probably interested in receiving your call. Such assumptions are not always valid for recommender systems. The problems associated with the self-reporting approach include: experts failing to find the time to enter their data in the first place; data entered initially in a burst of enthusiasm by the individual or organization becoming obsolete or out-of-date; inaccurate and/or unvalidated self-reporting of expertise; and the levels of experience and degrees of currency are typically not being captured or maintained.

The approach we offer includes a combination of both techniques, together with a number of verification and validation techniques to improve the consistency, completeness and reliability of the knowledge, with the caveat that we cannot ensure consistency and completeness.

In the following section we consider related work on recommender systems. In section 3 we present two case studies conducted to elicit what is needed in an expertise recommender system. Section 4 presents the approach including results from a usability study performed to evaluate the prototype developed. Conclusions are given in Section 5.

2. Related work

Recommender systems share much in common with search engines which allow a user to enter keywords on a topic they are interested in and produce a list of links to resources based on those keywords and often also on the profile of the user. One of the most well known recommender systems is Amazon.com (http://www.amazon.com/) which allows the user to enter keywords and will search for books and other products based on those keywords. If the user does select one of the recommended products, the system will then suggest other products that it thinks the user might like based on the choices of other users who also selected the same product. That is, the system reasons that if a particular user likes the same product as 20 other users, they may also like some of the other products that those 20 users liked; this is known as collaborative filtering and is a common technique in modern recommender systems (e.g. [4]).

Recommender systems for experts work in much the same way in that they take a user's search query and try to find someone who has the expertise to answer that query. The user will then be given the expert's contact details. These systems are mainly used internally within organizations. Validation that the information provided was wanted and useful is missing in many systems. Amazon attempts to obtain feedback using the message box shown in Figure 1.

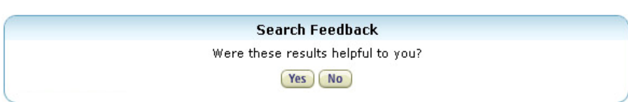

Figure.1. Search feedback from Amazon.com

Aïmeur et. al. [4] further explored the concept of validation of automatic identification of experts. They describe a recommender system called HELP which attempted to locate expertise and experts within an organization. The system included a database of questions asked previously by users and their respective answers. When a user searched for a solution to a particular problem, the question/answer database was first searched to see if their problem or a similar one had already been dealt with. If it had, the user was then presented with the solution, which they could choose to either accept or reject. If they rejected the solution, or if one wasn't found, the system would then search its user database for someone with the potential expertise to solve the problem. Users were thus required to register their own areas of expertise with the system. The rating system shown in Figure 2 served to somewhat validate the recommendation the user was given by storing the responses in the profile of the expert. If someone claimed to be an expert on a certain topic but consistently received low ratings, then they would not be recommended by the system if it was possible to recommend someone with a higher rating.

Figure. 2: Evaluation of Experts [4]

An alternative to self-reporting recommender system, are fully automatic approaches to locate experts such as Who Knows [5] and SAGE [6] using inputs such email (e.g., Agent Amplified Communications [7]), bulletin boards (e.g. Contact Finder [8]), Web pages (e.g. YENTA [9] and MEMOIR [10]), program code (e.g. Expertise Recommender [11] and Expert Finder [12]), and technical reports (e.g. KCSR Expert Finder [13]). These techniques could be applied to the artifacts of social software systems (i.e. email, WebLogs and

Wikis) to provide automatic expert location. However, a review of these systems [3] found problems related to heterogeneous information sources, expertise analysis support, reusability and interoperability.

Quickstep, described in [14], is a recommender system for online papers. It uses unobtrusive monitoring of searchers' browsing behaviour to find what kind of papers a searcher is interested in and create an interest profile for each searcher. When a searcher searches for papers, only those papers that have not already been browsed by the searcher and have a topic of interest to the searcher are returned.

Feedback forms may be provided for a searcher to provide negative or positive comments on the recommendations after the item has been received by the searcher. One major problem with relying on the searchers to actively provide feedback as in [4] is that there is never a certainty that they will do so. Aïmeur *et al.* [4] reported that most people would provide positive ratings if any at all. On the other side, if a searcher had an extremely negative experience they may provide a rating, but otherwise may not bother. Even if the system sends regular emails to the searchers to remind them to provide feedback on a recommendation they were given, it is still not guaranteed that the searcher will do so. In fact, the searcher may regard the reminder emails as an annoyance and ignore them altogether.

Any recommender system that recommends items for purchase (such as Amazon.com, for example) can, to some extent, measure how valid and useful a recommendation was by recording if the searcher decided to buy the item that was recommended. However in most of the expert recommender systems encountered in the literature, there is no way of knowing whether a recommendation of an expert caused a searcher to contact the expert. These systems provide the names and contact details of recommended experts, but then leave it to the searcher to contact the experts at their discretion; thus they not only have no idea if a contacted expert was able to help the searcher, but also have no idea if a searcher even tried to contact the expert in the first place.

The lack of feedback problem is addressed in [4] by ensuring that all searcher and expert interaction is controlled by the system and by having profiles for both searchers and experts. However, their approach is geared towards providing quick solutions to problems rather than putting people in contact with one another. Thus it seems unrealistic to insist that all contact between a searcher and an expert be through the system, but it does make sense that searchers be allowed to make initial contact with experts via the system, and experts be allowed to send the initial response through the system so the system can gather data on an expert's availability and response time.

3. Exploratory case studies

As our goal is to create a recommender system that has the confidence and support of its intended users and goes beyond the yellow-pages model, it was essential that we not just review the literature but also analyse the experiences of practitioners frequently concerned with the task of locating experts and expertise. Thus we conducted two case studies.

3.1 Case study 1

Firstly we conducted interviews with seventeen personnel during the months of September and November 2006 within a Defence R & D organization which is expertise intensive. A series of questions was presented to our interviewees in an attempt to get them to present their experiences on accessing expertise in each of their fields (e.g. "What features or criteria do you use to determine if someone is an expert?", "How do you find what projects/problems people are working on?"). The questions were also designed to elicit barriers they faced to gaining expertise/finding an expert as well as assessing the quality of the expertise they were provided with. It is this last parameter (quality) that is the most difficult to assess.

The following five main questions were asked (with subquestions associated with each to further prompt answers):

1. How do you go about finding an expert? Please give one or two actual examples of how you have done this in the past.

- Have you used any tools or software support?
- What sources of information do you look at? (documents, email, websites, databases, personal reference)
- What role has the web played in finding information about expertise inside and external to your organisation?
- What mechanisms does the organization have to identify experts?
- What features or criteria do you use to determine if someone is an expert?
- Do you consider personal characteristics? If so, how would you work these out?
- What ranking/order would you give to the criteria that you use to identify/find a person?
- Does the importance of a criteria vary for different situations? If so, can you give examples from your

experience of when certain criteria were important and which were less important or unimportant in a different situation?

- How do you determine the person's level of expertise and the currency of that expertise?
- How do you find what projects/problems people are working on?
- Is it more useful to know what problem someone has been working on or the application/domain they worked within?
- How long do you usually have to find an expert? Do you need to find an answer to a problem as quick as possible, or is it in the project planning phase when you are looking for team members/mentors/advisors?

2. Do you use a different process depending on the location of the person, their status, the department they are in, or other factor? Please give one or more examples.

- Has your strategy changed over time?
- Does this process differ for people outside of your organisation?
- Who decides who is in a team? How do they decide? Please give an example of how a recent team you were involved with was put together.

3. What are the impediments or barriers to identifying an expert?

- What are the impediments or barriers to accessing an expert?
- What are the impediments or barriers to validating an expert?

4. Is there information you are interested in gaining access to but currently you are unaware of where you could find it or whether what exists is accessible or reliable?

- Do people advertise their skills? Should they advertise? What about bidding for projects?
- What role does budget, timeframe, resources play in finding an expert and then getting access to them?

5. How is trust developed regarding expertise in your organisation?

- How do you validate that someone is an expert and is that information passed on in some way to others in the organization?
- What mechanisms does the organization have to reward experts? Is there any incentive to be recognised as an expert or does this lead to more work or less time for your own work?

It was clear from our initial investigations in 2006, that a fully automated approach which advises who to contact, or a semi-automated approach using techniques such as SNA to answer "who knows who" would not deliver an optimal or widely accepted solution. There were basically two ways of accessing 'know-how' within the organization: 1) drawing upon one's social networks (a people-centric approach); 2) reading publications/project descriptions to determine who has the relevant expertise (a more algorithmic/automated type of approach). Issues that arose in accessing expert/ise included:

- Establishing trust and quantifying levels of trust. Interviewees using social networks to locate expertise mentioned the importance of trust between the members of the social network. Also mentioned was the difficulty in determining the level of truthfulness in journal and conference papers and how much trust should be placed in the reported results.
- The organisation experiences a high turnover of experts and thus a loss of expertise.
- Access to experts across organizational boundaries.
- Currency of expertise and the impact of currency on relevancy of expertise. Although the organisation attempted to maintain the currency of the expertise held by their experts by up-dating their knowledge and ensuring that the relevant training was administered, it was mentioned that knowing something is better than knowing nothing, even if the knowledge is not current.
- Short timeframes (hours/days) for decision making related to forming a new team and project.
- Information is typically classified and only available on a need to know basis.

In this organization the key concern was "how do you find someone to work on a particular project?" Using our interview transcripts we conducted content analysis as we had done with the literature. This revealed that people and projects, and to a lesser extent, tools and groups, play central roles in identifying who is an expert. The structure of the organization also played a major role in being able to find and access an expert.

The majority of the Defense R&D personnel possessed PhDs and had a high level of technical knowledge. We found that the senior personnel tended to place more importance on the quality of a potential expert's publications than less senior staff. All personnel interviewed mentioned the importance of social networking and maintaining contact with others in their field. Senior staff tended to have more international contacts as well as larger social networks within their organization. The most junior staff member we interviewed was the least concerned with finding experts and had the least number of contacts in other departments of the organization.

Three of the personnel interviewed were liaison librarians who were often required to do literature searches for staff members. They mentioned that staff members were starting to do their own web-based subject searches rather than asking the librarians for help. The librarians personally visited a number of staff members to keep them up to date with what is happening in the organization. They also told us that many staff members were socially isolated.

At the time of the interviews, the organization had no established or frequently updated expert recommender system. We were told that there had been attempts to create an expertise database, but it was difficult to keep up to date and was not current at the time of the interviews.

3.2 Case study 2

The second case study we conducted was within our own university, another expertise intensive organization. There are four key areas within the university concerned with locating and contacting experts. These are:

- the Research Office (RO) – concerned with finding experts for research projects;
- Development and External Relations Division (DERD) (now better known as Community Engagement) – concerned with finding experts for awards and for connecting Macquarie staff with industry and the community;
- Marketing and Public Relations (MPR) – concerned with finding experts for media interviews on behalf of journalists; and
- Teaching and Learning (TL) – concerned with finding experts for guest lectures and expertise related to teaching and learning such as skills and experience with working with groups or teams, giving iLectures, handling large classes, providing feedback, etc.

We interviewed representatives (the senior executive of the RO, director and two others in DERD, two senior people in MPR and the Chair of the University Learning and Teaching Committee for TL) from these four areas, asking them the same questions as in the first case study (removing questions that were specific to the first organization) as well as asking them their general requirements for an expert recommender system. Based on the findings from the literature and our own experiences, we were particularly interested in the answer to the following questions:

1. *If there was a system that recommends experts, would you be willing to spend a small amount of time providing feedback to the system to indicate whether a recommendation given to you was useful?*
2. *What sort of searcher feedback would be useful to you?*

A few of the representatives had reservations about the feedback form (Fig 2.) presented to them as an example of what feedback might look like. TL mentioned that any negative feedback on a person's teaching style would be taken very personally and is akin to criticising the person themselves, as teaching skills are not something that can be taught easily or learned right away. While it may be useful to indicate which experts have good and bad teaching skills, it may cause people to be more reluctant to use the system or volunteer to give guest lectures. MPR mentioned that journalists would be unlikely to fill out the feedback form unless they had a negative experience and it was immediately available.

The methods university representatives preferred with respect to finding experts were quite similar to those used by staff at the Defence R&D organisation. The RO searched through publication titles in the Integrated Research Information System (IRIS) to find an expert. MPR said that they frequently used Xpertnet, Macquarie University's expert recommender system for journalists, to locate experts.

The requirements listed by the representatives included:

- Searching by age (suggested by DERD for offering age-specific awards).
- Searching for expertise in general as well as expertise in a particular field (suggested by DERD for offering non-domain-specific awards).
- Details of well-established and/or recognised experts, such as professors, and lesser established, (e.g. Masters and PhD students) experts were needed, as well as non-academics performing relevant research. (This was suggested by DERD as they have found that people studying towards a higher degree were more likely to be interested in applying for an award).
- Details about an expert's teaching awards, publications, and grants, as well as information about the type of teaching skills they possess (e.g. if they have experience in student centered learning, or team teaching). (Suggested by TL as a method of judging the quality of the expert's teaching skills and finding out, for instance, if they would be suitable to give guest lectures or mentor new staff members).

- A model illustrating the relationships between experts (Suggested by RO to show joint grants and papers).
- Certain flags to restrict searches on grant applications that are confidential (suggested by RO).
- Experts who are registered in the system should be allowed to nominate availability (suggested by MPR as the response time of an expert was an issue).

4. Approach

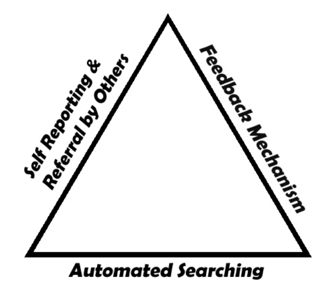

Figure.3. Our Triangulated approach

The case studies revealed that a combined automated and human-in-the-loop approach was necessary. As depicted in Figure 3, the proposed approach uses automated searching as a foundation from which the initial data is captured and against which the data is regularly cross checked. The key inputs to data mining include individual web pages, project/grant repositories, citation indexes (e.g. CiteSeer (http://citeseer.ist.psu.edu/) and publications databases. Within our university we will use the Integrated Research Information System (IRIS) as one of our data sources. IRIS consists of a collection of all publications and impact factors of individuals within the university. The latter is currently our most useful resource as it is the most structured.

In the information extraction technique we trialed [15], results were sent to each of the 20 identified experts in the Computing Department in an email giving them an opportunity to review and validate the areas of expertise found by the system. The result provided to each expert was a set of RFCD (Research Fields, Courses and Disciplines) codes as defined by the Australian Research Council (ARC). These codes were based on the expert's publications in IRIS and were used as an externally validated and publically recognized indicator of their research areas. In addition to validation of the outputs of datamining concerning their areas of expertise, we also propose that the expert would be able to provide additional information regarding their preferences as shown in Figure 4.

Using automated searching we can attach dates related to the expertise found to assist with currency. We can also keep statistics on the level of expertise using simple measures such as the number of publications in that area and Term Frequency Inverse Document Frequency (TFIDF).

This confirmation and correction by the expert of the results of automated searching provides a second dimension: *Self reporting and referral by others* which is based upon the first dimension: *Automated Searching*. By allowing experts to edit the automated results we are allowing them to self-report their areas of expertise, deleting or adding new areas as they see fit. The system would also allow for others, such as a PhD supervisor, to nominate or refer another person, such as their student.

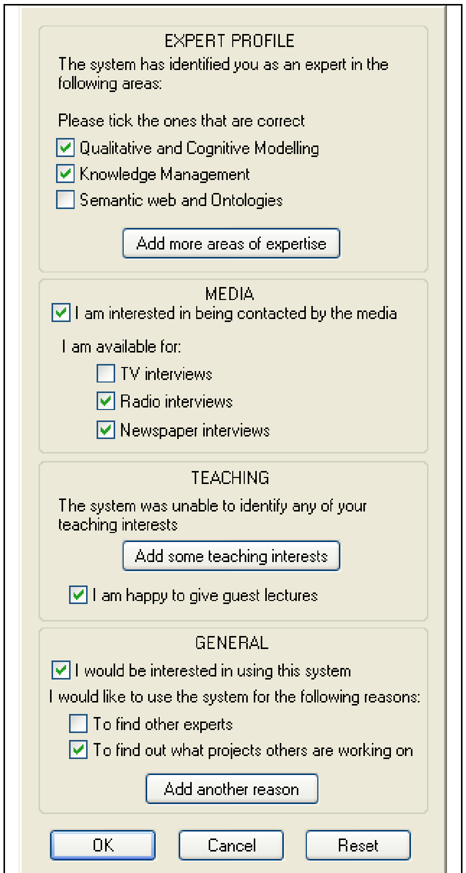

Figure 4. A simplified validation screen sent to the expert as a result of data mining from webpages, publication/citation databases, etc.

With many systems that rely entirely on self-reporting, some people will choose to simply not to have a profile rather than going to all the trouble of registering themselves and maintaining their profile. In large organizations it is also possible that some people may not even know about the system, especially if it is fairly new. To populate such a system it is most likely that some experts would need to be contacted either personally or via mass email and the rate at which experts are added to the system would be directly related to the rate at which experts (most of whom would be very busy) are able or willing to register themselves.

Rather than asking potential experts to "opt in", by using automated methods as the foundational first step in information acquisition, this system would instead ask them to "opt out" if they do not wish to be registered with the system. It would also be much easier for an expert to review a portfolio which has been automatically generated for them rather than have to format one for themselves. In addition to sharing the data validation and maintenance with the searcher, the system provides a standard way of describing its experts, which will simplify the searching process.

To address the issues of external validation, expertise currency and motivation to enter and maintain the data, individuals are sent the output from data mining at regular intervals, say twice yearly. Given the importance of reputation and track record in the university system we believe academics would be motivated to check what the system, as it reflects the data in the world about them, says and to correct any errors or omissions.

From discussions, interviews and personal experience the key question to be asked of the advice of any recommendation system is "was it useful?" As shown in Figure 3, the third support to our approach is a feedback mechanism that will allow the searchers of the system to validate the recommendations themselves.

Feedback will be gathered from both the searcher and the expert. As illustrated in the activity diagram in Figure 5, once an expert is recommended by the system the searcher can choose to contact the expert via a contact form provided by the system. The system would then email this message to the expert along with a link to a form where they can send their initial reply. In this way the expert can provide information about their availability and expertise as well as writing a personalized message to the searcher (Figure 6).

If the searcher wishes to contact the expert via another route, such as a phone call, the system would allow them to indicate this by clicking a link or button. If the searcher chooses this option, a message will be sent to the expert informing them that they may be contacted regarding the search terms entered by the searcher. The expert will be directed to a form where

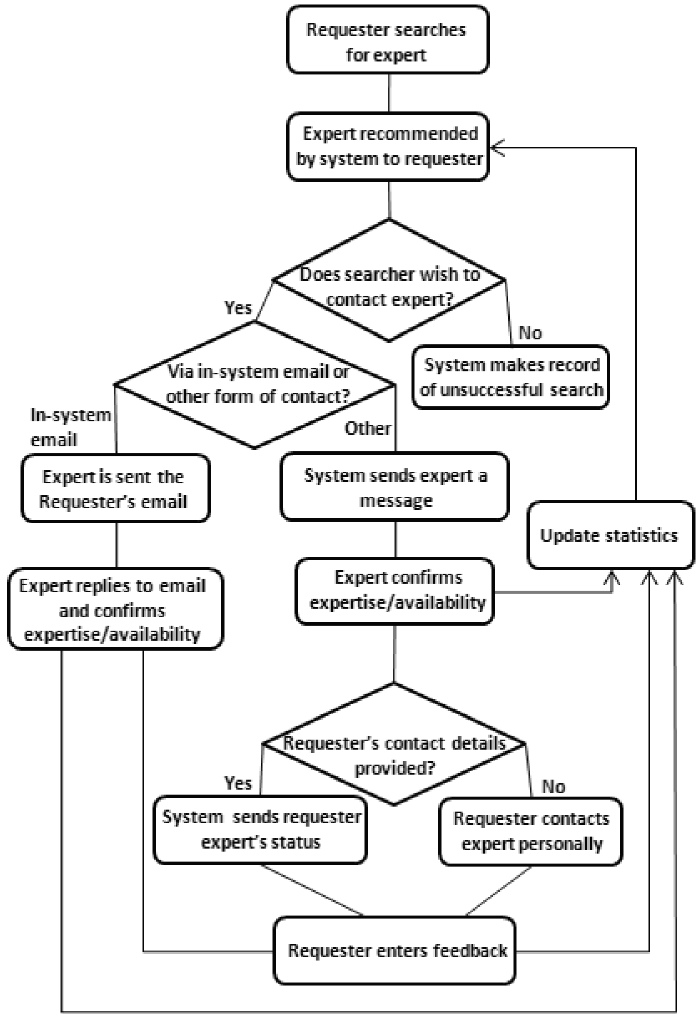

Figure 5. Activity diagram depicting search flow between Searcher, Expert and System

they can indicate if they are available and possess the expertise to address the query. This form would be similar to that in Fig 6 without the searcher's message at the top and the space to write a personalized message. The system will then be able to update the expert's statistics in the system and inform the searcher of the expert's status if the system has both searcher and expert profiles or if the searcher provides the system with their email address.

Thus the system is able to keep track of what recommendations yield success (that is, which recommendations result in an expert being contacted). If the searcher does not indicate in any way that they wish to contact the expert, the system will store the search terms used for a brief period of time and observe if similar search terms also yield unsuccessful results. If so, the system will reevaluate the profiles that are recommended for those search terms.

After interacting with a recommended expert, a searcher may provide feedback on the expert whenever they wish. Otherwise a follow-up technique such as sending an email to the searcher a couple of weeks after the searcher makes contact with an expert could be used, reminding them to provide feedback.

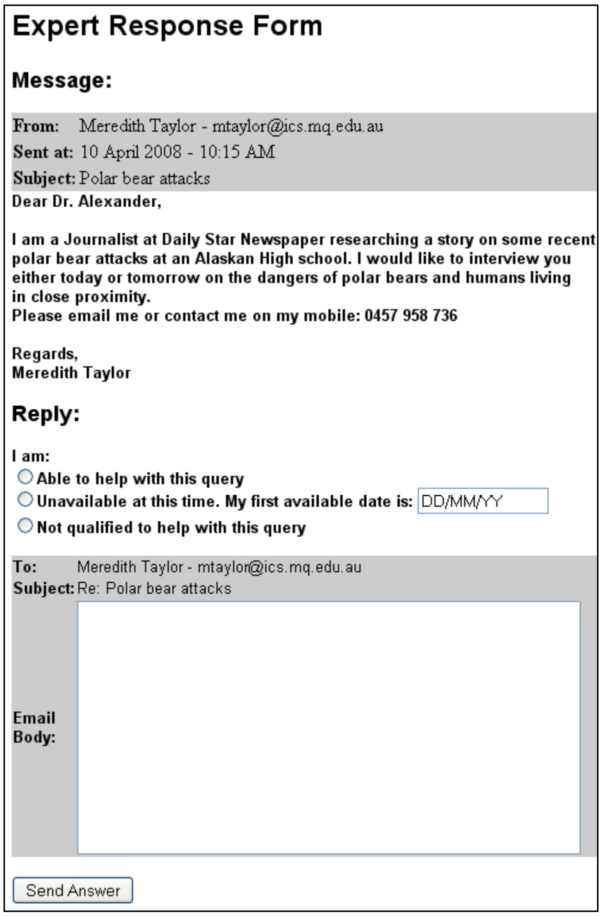

Expert Response Form

Message:

From: Meredith Taylor - mtaylor@ics.mq.edu.au
Sent at: 10 April 2008 - 10:15 AM
Subject: Polar bear attacks

Dear Dr. Alexander,

I am a Journalist at Daily Star Newspaper researching a story on some recent polar bear attacks at an Alaskan High school. I would like to interview you either today or tomorrow on the dangers of polar bears and humans living in close proximity.
Please email me or contact me on my mobile: 0457 958 736

Regards,
Meredith Taylor

Reply:

I am:
- Able to help with this query
- Unavailable at this time. My first available date is: DD/MM/YY
- Not qualified to help with this query

To: Meredith Taylor - mtaylor@ics.mq.edu.au
Subject: Re: Polar bear attacks
Email Body:

Send Answer

Figure 6. Example of an expert response form with searcher's message at the top.

The screens which make up the user interface play a critical part of the approach by allowing the development of two profiles, or user models, of the *expert* as well as the *service requester*. The two user models will provide knowledge to the system that will be used as a filter for searching and matching. The need for customizing both sides became apparent from interviews revealing that the (type of) person being sought varied for different departments. Thus it is necessary to capture the needs and preferences of the person looking for someone (e.g., they are a news journalist and need someone to provide an expert opinion for television in layman's terms in the next hour) and also the preferences, communication skills and availability of the expert (e.g. they have given radio interviews in the past but are away for two weeks). The user models can potentially be populated from data found in other places, such as from webpages maintained by the individual or corporate pages, but realize that the feasibility of this depends on the level of webpage standardization across the organization and the degree to which content management is enforced. Thus in our solution we leave this as potential future work. Statistics relating to the responses from experts regarding their expertise and availability and the requesters satisfaction with the service provided are part of the user model. The service provided has many aspects to it, some of which, such as ability to communicate and availability, can be used to rate the expert and order future results to searches. Recommended experts who frequently receive negative feedback (or poor ratings if a rating system similar to that used by [4] is used) will be recommended less (or perhaps ranked lower in the final output) by the system, than experts who have consistently received positive feedback, for example. This feedback mechanism will act as a referral system, suggested by some of the people we spoke with as the type of system they would be interested in, rather than a yellow-pages model. They didn't want contact details of anyone professing to know a particular domain, they wanted to find someone based on the experience of someone else, who they preferably knew and trusted.

While the system will use the statistics for generating and ordering recommendations in response to a query, for privacy and ethical reasons, only the expert will be able to see his/her own statistics regarding their overall usefulness and availability as perceived by the service requesters. We see this personal feedback as potentially useful for professional development and self evaluation, similar to the way in which data from student evaluations of a teacher may be used for professional development and disclosed at the experts discretion for promotion or other reasons. As represented, this additional feed-back/validation pillar provides a triangulated approach bringing together automated machine-based knowledge discovery and manual validation.

4.1 Usability study

After the requirements for the proposed system were decided upon, it was important to test that the features of the system were actually what people wanted. All too often it is the case that a system with too many features will not be easy or enjoyable to use. Features such as an in-system contact form may be viewed as an annoyance rather than a benefit. For this reason we devised a usability test on a semi-functional version of the system.

We made a prototype expert recommender system, WHOKNOWS? that we populated with a small amount of test data. Ten mock expert profiles were put into the system as well as several screens to help searchers find and contact experts. The screens consisted of an initial search screen (Figure 7), a screen to show the results of a search, a screen for the searcher to contact the recommended experts, and a screen for the searcher to provide feedback on an expert. Screens showing the profiles of the experts were also included as well as a screen for the expert to respond to a searcher's request (Figure 6), although this screen was not included in the test.

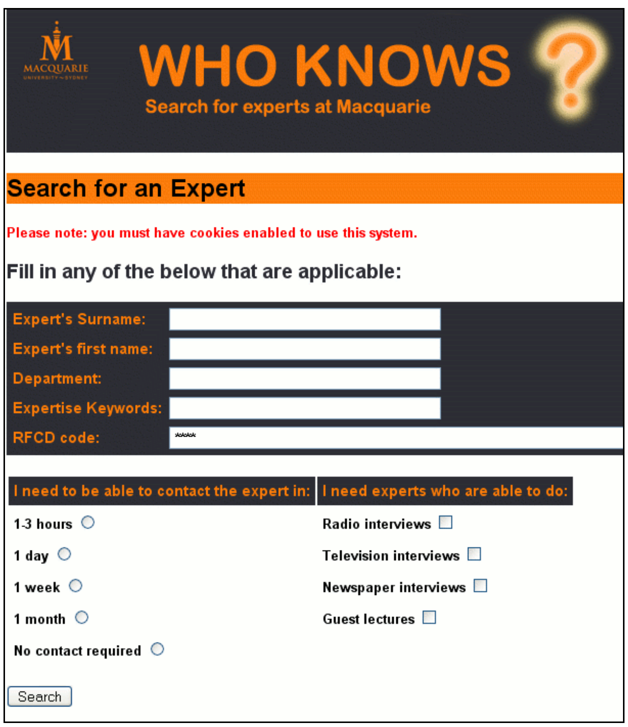

Figure 7. Initial search screen in test system

WHOKNOWS? did not allow for any searcher profiles. Instead the initial search screen (Figure 7) allowed users to specify what time frame they had to contact the expert in and whether they required an expert for a radio, television, or newspaper interview, or a guest lecture.

As the system did not contain any real expert profiles, no automated searching was done by the system, rather we assumed that this stage had already been completed, and the resulting expert profiles had been stored in the system. An evaluation of automated searching and expert validation was performed on members of the Computing department at Macquarie University using publication data found in IRIS and is described in [15].

Scenario 1 Part 1.

You are a Journalist working for the Daily Star Newspaper. You are researching a story about a polar bear attacking some school children outside an Alaskan high school. You wish to find an expert on polar bears to interview for your article. You need to have the article ready in two days. Use the system to find and contact the expert.

Scenario 1 Part 2.

After contacting the expert, you receive the following reply:

Hi,
I am too busy to give any interviews at the moment. However, you may try James Paterson (james@email.com) as he will be only too happy to grant you an interview.

You contact James and he responds immediately. You are able to get an interview with him that day.

Use the system to provide feedback for the expert you initially contacted.

Figure 8. First scenario in usability test

The algorithm for ranking the recommended experts returned by a search was also implemented so it could be assessed by the participants. The algorithm contained the following steps:

1. all experts to whom the criteria entered by the searcher was not applicable were removed.
2. if the searcher entered expertise keywords, the remaining experts were ranked on how many keywords were found in their listed areas of expertise. The experts for whom no keywords were found were removed.

The remaining experts were ranked on their combined availability and searcher feedback scores.

WHOKNOWS? and the usability test were both made available online. Two scenarios were given to the participants (who responded to an emailed advertisement). Each scenario gave the participant a job occupation and a task that involved using the system to search for an expert (Figure 8). After each scenario had been completed, the participants were asked to fill in a questionnaire. The questionnaire asked the participants to rate how easy or difficult it was to complete the task, and whether the layout of the system helped or hindered them. The participants were asked to evaluate the algorithm the system uses to rank the experts, as well as their opinion on certain components of the system (such as the contact and feedback pages).

4.1.1 Participants

Thirty-eight participants responded to the emailed advertisement. However, as the usability test was online, it was not possible to make sure that the participants completed the whole test. As a result only 28 participants completed all steps which were part of the first scenario and filled out the survey, and 22 of those went on to complete the second scenario and fill out the associated survey.

	Gender	
Age Range	Female	Male
15-19		1
20-24	2	1
25-29	2	7
30-34	3	
35-39	2	2
40-44	1	1
45-49		3
50-54		
55-59		
60-64		1
65-69	1	1
Total	11	17

Table 1. Number of participants in each age range by gender

A summary of the biographical details of the 28 participants can be found in Tables 1 and 2. Participants were both male and female and a range of ages. About half the male participants and a third of the female participants were employed in a job where they needed to find experts. These included members of the Research Office and Media and Public Relations Office at Macquarie University who had participated in our interviews in section 3.2.

	Involved in finding experts		
Gender	**No**	**Yes**	**Total**
Female	7	4	11
Male	8	9	17
Total	**15**	**13**	**28**

Table 2. Number of participants who are and are not involved in finding experts for their profession by gender

4.1.2 Results and Discussion

Results were gathered for each part of the test system: the search page, results page, expert profile page, contact page and feedback page. Each of these is outlined in the subsections below. The system's ranking algorithm was also evaluated by participants, although the results are not discussed here and will be presented in a future publication.

The questionnaire given to each participant after they completed the tasks in each scenario contained statements about each screen in the system. The participants indicated their level of agreement with each of the statements on a 5-value Likert scale. The statements for each page were the following:

Search Page

S1. I found it easy to understand how to search for the expert on the search page

S2. The search options on the search page were specific enough for me to search for the expert I needed

Results Page

S1. The experts' details on the results page were sufficient for me to tell if I needed to contact them.

S2. It was clear to me how I could use the system to contact the experts on this page

S3. After reading this page I understood how to provide feedback on an expert.

Expert Profile Page

S1. The details of the expert listed on this page were sufficient for me to tell if I needed to contact them.

S2. I found the categories on this page easy to understand.

S3. It was clear to me how I could use the system to contact the expert on this page.

Contact Page

S1. It was clear to whom the email was being sent.

S2. I found it easy to understand how the text boxes should be filled in.

S3. I would use this feature in the future if available.

Feedback Page

S1. I was able to adequately express my feelings about the expert on this page.

Table 3 shows the percentage of participants that either agreed or strongly agreed with each statement for each scenario.

Search Page	Scen1	Scen2
S1	71.43%	95.45%
S2	78.57%	90.91%
Results Page		
S1	89.29%	95.45%
S2	85.71%	95.45%
S3	60.71%	86.36%
Expert Profile page		
S1	92.86%	95.45%
S2	71.43%	90.91%
S3	89.29%	95.45%
Contact page		
S1	85.71%	95.45%
S2	85.71%	90.91%
S3	89.29%	90.91%
Feedback page		
S1	60.71%	63.64%

Table 3. Percentage of participants who agreed or strongly agreed with statements about each screen in the system after completing tasks in the first and second scenarios

From Table 3 we can see that the percentage of people who agreed or strongly agreed with the statements after completing the task in scenario 2, is higher in every case than the percentage who agreed or strongly agreed after completing the task in scenario 1.

The percentage increase can be explained partly by the fact that the participants would have a better grasp of how the system works after completing the first scenario and starting the second; and partly by the fact that six of the participants who completed the first scenario did not complete the second. Table 4 shows the percentage of the 22 participants who completed both scenarios who either agreed or strongly agreed with each statement for each scenario

After removing the 6 participants who didn't continue to the second scenario, we can see in Table 4 that the percentage differences are smaller in general than in Table 3. Percentage increases are recorded for both statements about the Search page; statement 3 about the Results page; and statement 2 about the Expert Profile page. This increase can most likely be attributed to the participants gaining experience in using the system after they completed the second scenario

Search Page	Scen1	Scen2
S1	81.82%	95.45%
S2	86.36%	90.91%
Results Page		
S1	95.45%	95.45%
S2	95.45%	95.45%
S3	68.18%	86.36%
Expert Profile Page		
S1	100.00%	95.45%
S2	81.82%	90.91%
S3	100.00%	95.45%
Contact Page		
S1	95.45%	95.45%
S2	90.91%	90.91%
S3	90.91%	90.91%
Feedback Page		
S1	68.18%	63.64%

Table 4. Percentage of participants who agreed or strongly agreed with statements about each screen in the system after completing tasks in the first and second scenarios with participants who did not complete both tasks removed.

A few participants found the layout of the search page confusing initially, and some said that they would have preferred fewer options and an "advanced search" option instead.

A relatively low percentage of people (68.18%) agreed or strongly agreed with the third statement about the Results Page (*after reading this page I understood how to provide feedback on an expert*) after completing the first task. Many participants thought the instructions on how to submit feedback were not very clear initially. To rectify this, there should be a separate button for each expert that, when clicked, would take the searcher immediately to the feedback page for that expert. In reality, however, a searcher is probably not likely to provide feedback on an expert immediately, but rather after some time has elapsed and they have been sent a reminder email by the system that includes a link to the feedback page for the expert they contacted.

One participant commented that they were not able to discern how helpful an expert was going to be by viewing their profile. This is a difficult problem to fix, as the feedback information the system uses to rank an expert is not displayed for ethical and practical reasons. Many experts would not be happy with their details in a system that displays to the public what other people think of them. If an expert saw that they had an average feedback score of 20%, for instance, they may become upset and ask for their profile to be removed from the system. While an expert should be allowed to know what their feedback score is, showing this information to all users of the system would not be appropriate. Showing an expert's availability information to the public, however, may be acceptable, as this is not based on people's opinions, but on facts.

It would be beneficial to both the searcher and the expert to have the expert's availability information displayed, as the searcher will know that they might not have much luck if they try to contact the expert, and the expert will not have to be constantly rejecting requests for help from searchers.

The percentage of people who agreed or strongly agreed with the statements about the Contact page remained the same for both tasks. The most promising result was the high percentage of participants who indicated for both scenarios that they would use the feature in the future if available (statement 3). None of the participants disagreed with this statement, although one participant commented that they could imagine copying the email address and sending their own email to the expert.

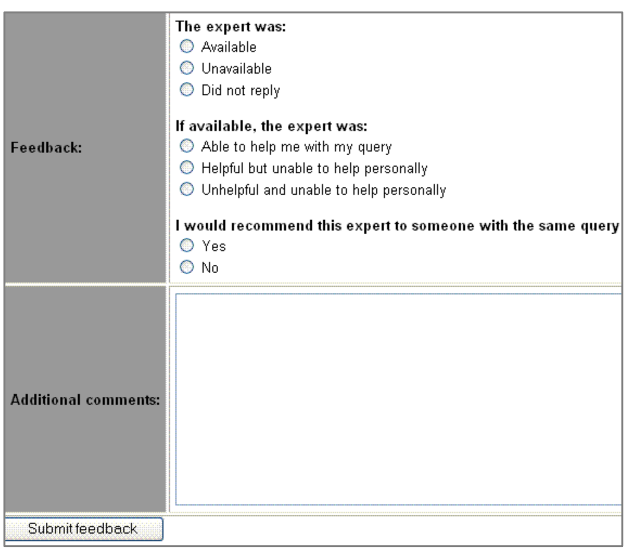

Figure 9. Feedback form

The statements about the feedback page generated the lowest percentage of agreement (68.18% after the first task and 63.64% after the second task). Some participants thought the feedback options available on this page (Figure 9) were too rigid, especially for scenario 1, when the recommended expert actually recommended another expert, but wasn't of any help otherwise. The section of the feedback form that requires a Yes/No answer (*I would recommend this expert to someone with the same query*) would be especially hard to fill out in a situation such as this. Another participant commented that the additional comments section was the only place where an expert's performance could be evaluated (with the other sections evaluating the expert's immediate response and availability). The feedback page was structured in this way to avoid making people fill out too many sections, as they would be unlikely to provide feedback regularly if this was the case.

Adding another section to indicate how satisfactory an expert's performance was could be a step towards solving this problem. It could allow the user to give a Yes/No response to the statement "I was satisfied with the expert's performance" or have them rate their satisfaction with the expert's performance on a scale of 1-5 with 5 being very satisfied and 1 being very dissatisfied. There are some issues with this method, however. A person's satisfaction with another person's performance can be very subjective. One person may think an expert performed excellently, while another may think they performed poorly, even if they gave the same performance in both cases. If free text was used to evaluate an expert's performance, the searcher could choose the comments to be sent to the expert so they can see exactly what the searcher was dissatisfied with. A Yes/No response, or a rating out of 5 would not give the expert a good idea of exactly what the searcher thought, and would therefore not be able to improve.

A second option would be to show each expert their feedback scores and comments on a private part of their profile. This may encourage them to improve their performance if the searcher was not satisfied. This would require comments to be heavily moderated, however, to ensure that searchers are not allowed to submit abusive feedback.

5. Conclusion

Rating systems, such as we propose and that used by [4], raise several ethical issues. For instance many people may object to the concept of rating another person and may refuse to participate. On the other end of the scale, some users may give recommended experts an unnecessarily bad rating simply because they don't like them on a personal level. In addition, the possibility that a person's personal or teaching skills could be criticized would be a sensitive issue for many and may result in a large number of experts refusing to be registered in the system. In our approach we aim to try different methods of user feedback as well as limiting the visibility of an expert's feedback results and preserving the anonymity of the users who provide the feedback in an attempt to avoid the ethical issues. In addition, we will also consider the use of more personalised feedback (e.g. a reporter wishing to interview an expert will be rating them on different criteria than someone wishing to work with the expert on a project).

As a key part of our approach is the combination of self-reporting/referral and automated searching through available data. Some data can only be obtained via self-reporting (e.g. indicating if you are available to do media interviews or guest lectures). However, information about which units one teaches, expertise areas, grants, etc. can be gained from personal websites and internal databases. An outstanding issue would be how to reconcile differences between these sources and between the outputs of automated searching and self-reporting.

A usability study was performed on the test expert recommender system we developed – WHOKNOWS? The participants included both males and females across a range of ages, a number of whom are concerned with finding experts in their professions. The participants completed two tasks using the test system and filled in questionnaires about the system's various features. A large majority of the participants responded favourably to the system and provided valuable feedback that resulted in a re-evaluation/design of the system's ranking algorithm, search options and feedback form.

Recommender systems greatly speed up and simplify the searching process, whether the item being searched for is a book, film, or another person. This project is interested in a recommender system which maintains user profiles to match experts with service requesters. Validation of the user profiles and the system recommendations using a combination of automated and human-based techniques seeks to combine and reinforce these two main approaches which individually have numerous weaknesses. Closing the loop between the seeker and the sought is aimed at providing both parties with confidence and motivation to use the system. We anticipate that our findings will be of use to other recommender systems and search engines, such as Google, in general. Finally, to provide a generalized framework for expertise location, we will consider what modifications are necessary to allow other knowledge-intensive organizations to use the framework and toolkit.

6. References

[1] D. Richards, M. Taylor, and P. Busch, "Expertise Recommendation: A two-way communication channel", Proc. 4th International Conference on Autonomic and Autonomous Systems (ICAS'08), March 16-21, 2008, Gosier, Guadeloupe, pp. 35-40.

[2] D. Yimam-Seid, and A. Kobsa, "Expert Finding Systems for Organizations: Problems and Domain Analysis and the DEMOIR approach" Jrnl Org.l Comp. & Electronic Commerce vol. 13, 2003, pp. 1-24.

[3] Y. Sim, R. Crowder and G. Wills, "Expert Finding by Capturing Organizational Knowledge from Legacy Documents" Proc. IEEE ICCCE '06, Kuala Lumpur, Malaysia, 2006.

[4] E. Aïmeur, F. Onana, F. S. Mani and A. Saleman, "HELP: A Recommender System to Locate Expertise in Organizational Memories" Proc. AICCSA 2007. pp. 866-874.

[5] L. Streeter, and Lochbaum, K. "An Expert/Expert Locating System Based on Automatic Representation of Semantic Structure" Proc. 4th IEEE Conf. on AI Appl. Comp. Soc. of IEEE, San Diego CA, 1998, pp. 345-349.

[6] I. Becerra-Fernandez, "The Role of Artificial Intelligent Technologies in the Implementation of People-Finder Knowledge Management Systems" Knowledge-Based Systems, vol. 13, 2000, pp. 315-320.

[7] H. Kautz, B. Selman, and M. Shah, "Referral Web: Combining Social Networks and Collaborative Filtering" Communications of ACM, vol. 40(3), 1997, pp. 63-65.

[8] J. Lave, and E. Wenger, Situated Learning: Legitimate Peripheral Participation, Cambridge University Press, Cambridge, U.K, 1991.

[9] L. Foner, "Yenta: A Multi-Agent Referral-Based Matchmaking System" Proc. 1st Int.l Conf. on Autonomous Agents, Marina del Rey California, 2002, pp. 301-307.

[10] A. Pikrakis, T. Bitsikas, S. Sfakianakis, M. Hatzopoulos, D. DeRoure, S. Reich, G. Hill, and M. Stairmand, "MEMOIR – Software Agents for Finding Similar Users by Trails" Proc. 3rd Intl. Conf. on Practical Application of Intelligent Agents and Multi-agents, London, UK, 1998, pp. 453-466.

[11] D. McDonald, and M. Ackerman, "Expertise Recommender: A Flexible Recommendation System and Architecture" Proc. ACM2000 CSCW Conf., Philadelphia PA, 2000, pp. 231-240.

[12] A. Vivacqua, “Agents for Expertise Location” in Proc. AAAI Spring Symp. on Intelligent Agents in Cyberspace, Stanford, CA, 1998, pp. 9-13.

[13] R. Crowder, G. Hughes and W. Hall, “An agent based approach to finding expertise” in Proc. 4th Int.l Conf. on Practical Aspects of Knowledge Mgt Heidelberg Germany , 2002, pp. 179-188.

[14] S. E. Middleton, D. C. D. De Roure, and N. R. Shadbolt, “Capturing knowledge of user preferences: ontologies in recommender systems”. Proc. International Conference on Knowledge Capture (K-CAP’2001), ACM Press, New York, NY, USA, 2001, pp. 100-107.

[15] M. Taylor, and D. Richards, Discovering Areas of Expertise from Publication Data, In B-H. Kang and D. Richards, Proceedings of Pacific Knowledge Acquisition Workshop in conjunction with PRICAI'08, December 16-17, Hanoi, Vietnam, 2008, pp. 173-186.