

Evaluation of an Approach to Expertise Finding

Yee-Wai Sim and Richard Crowder

Intelligence Agents Multimedia Group
School of Electronics and Computer Science,
University of Southampton, Southampton SO17 1BJ, UK,
`rmc@ecs.soton.ac.uk`

Abstract. This paper presents an approach to locating an expert through the use of existing organizational information. This approach was realised through an Expert Finder framework developed by the authors. The framework enables the relationships of heterogeneous information sources to experts to be factored in to the modelling of an individuals' expertise. The framework also provides an architecture that can be easily adapted to different organizations. The framework has been applied to a real world application and been evaluated using the notions of precision and recall.

1 Introduction

Many organisations require systems that gather corporate information and make it available to employees in order to resolve specific problems or queries. This allows organizations to capitalise on their best knowledge, leading to higher levels of productivity and competency, [1]. When employees collaborate, they initially rely on their own experiences, but as knowledge transfer occurs during task activities, a dynamic knowledge creation process takes place. With the correct tools organisations can acquire and analyse knowledge from individuals, in order to learn from their successes and failures. As businesses and companies become larger and more geographically dispersed, this collaboration space has now becoming more virtual than physical; hence it is becoming increasingly difficult for organisations to know where their best knowledge is and even more difficult for them to *know who knows what*. A survey of a number of different approaches to People-Finding is to be found in the paper by Beccerra-Fernandez, [2]

The recognition for the need to augment the expertise finding behaviour has resulted in the development of a number of Expert Finders. An Expert Finder is considered to be repositories that attempt to manage knowledge by, holding pointers to experts, who possess specific knowledge. As discussed latter in this paper there are problems of maintaining and retrieving expertise in these systems, relating to the exploration of heterogeneous information sources, the methodology used to analyse expertise, and system interoperability.

In many organisations there is a need to develop a framework for providing up-to-date information for expertise modelling. This is because people in organisations accumulate knowledge through task achievements [1] and this output

is a valuable source for capturing knowledge that relates to an individual's expertise. In practice this output is the documentation (reports, memorandums, e-mails, etc.) generated during a project or similar activity. Exploiting meta data information from these documents can draw inferences to derive or update the knowledge about expertise associated of an individual.

2 Approaches to Expert Finding

Any Expert Finder requires a rang of information relating to an individual, for example;

- Level of knowledge or experience possessed.
- Performance of an individual compared to others in a related field.
- Current availability of an individual.
- Contact information of individuals who posses the required knowledge.

To manage an Expert Finder, there is a need for tools that gathers and consolidates this information in a form that is accessible by the Expert Finder. The availability of large electronic repositories in organisations, have led to the development of a autonomous approach to collect and analyse information when locating experts. The literature details are number of systems that undertake a fully automatic approach to expert finding, including, *Who Knows* [3], *Agent Amplified Communication* [4], *ContactFinder* [5], *Yenta* [6], *MEMOIR* [7], *Expertise Recommender* [8], *Expert Finder* [9], *SAGE* [2] and the *KCSR Expert Finder* [10]. In a review of these systems, problems related to heterogeneous information sources, expertise analysis support, and interoperability were identified.

Heterogeneous Information Sources. Experts can be used as effective and reliable information filters to locate useful information. This implies a role for Expert Finders in providing the expert-oriented access to an organisation's information system, thus mining an organisation's information system is a suitable path for an implicit source for expertise evidence.

In order to effectively explore the organisational information space for expertise evidence, Expert Finders need the ability to handle the heterogeneity of the widely distributed information sources. This is reflected by the wide variety of expertise evidence used, including, emails [4, 11], bulletin boards [5], program codes [8, 9], Web pages [4, 6, 7, 12], and technical reports [3, 10, 13]. Hence, a framework that is flexible enough to address this problem is required.

The systems proposed by Crowder [10] and Mattrox [13], used a raw document indexing technique to capture a documents' concepts. These systems ignored the documents' structural elements (e.g. title, abstract, etc.) and treated all the text as a *bag of words*. Since this technique only captures the concept-to-document relationships, a different approach is needed to capture concepts related to expertise, as in an Expert Finder, the analysis of the documents

should focus more on how they relate to experts rather than how they relate to the concepts they contain.

The authors propose that the heterogeneity of information sources should be used as an indicator to reflecting experts' competencies. How well these expertise indicators (e.g. indexed terms) reflect expertise is mainly a factor of how the source in which these indicators occur relates to the expert. This idea is based on the assumption that terms found in different types of documents indicate expertise differently, irrespective of their statistical traits. For example, the occurrence of a term in the title of the document shows a different distance to a persons actual expertise, compared to its occurrence anywhere in the body of the document. Therefore, the relationship of expert-to-document needs to be determined before extracting indexed terms from the document.

Expertise Analysis Support. The system developed by McDonald [8] incorporates a user-customisable expertise filtering process. Such a process takes a set of unfiltered names recommended by the system, and then reorders and removes items to generate a refined recommendation. Support for this filtering process can be provided by accessing the database that maintains personal and organisational relevant data, e.g. departmental affiliation.

A user seeking individuals as sources of information and/or as individuals who can perform given organisational functions, imposes their own requirements on the Expert Finder's functionality. Hence it is the users who should select the appropriate filters depending on their needs, and the system should only support the expert finding process by providing analysis functionality. This means that including the ability to rank experts using different user-customisable criteria (rather than the mere provision of a linear listing based on pre-determined criteria) can considerably enhance the expert finding applications. For example, a filter which represents the working relationships among experts can be used for catalysing collaboration in co-authoring documents or launching projects. In cases where users require knowledge that can only be shared through the human dimension, a filter based on the physical distance between departments can be employed.

Another approach that can support users in expertise analysis is to increase the system's transparency. This can be achieved by providing interfaces to access the expertise evidence, together with the expertise recognition logic. For example, the system can supply the scores for ranking the experts along with the associated expertise evidence. It allows users to evaluate the validity of the system's recognition logic. This in turn permits the incorporation of functionality, which can assist the users in evaluating and exploring the expertise evidence, i.e. spotting anomalies in the expert finding process, giving users greater trust in the system.

Interoperability. The majority of the Expert Finders surveyed focus on solving a particular problem, resulting in a closed standalone solution. However, Expert Finders can and should be integrated into other organisational systems

and should be readily transferable from application to application and be interoperable with other systems, [14].

3 A Framework for an Expert Finder

In order to address the problems discussed above, the authors propose the framework, shown in Figure 1, for an Expert Finder. The framework consists of a collection of components for expertise evidences extraction, expertise modelling, querying, expertise matching and user interface. These components are flexible enough to address different organisational environments. The functionality of the proposed frame work is as follows:

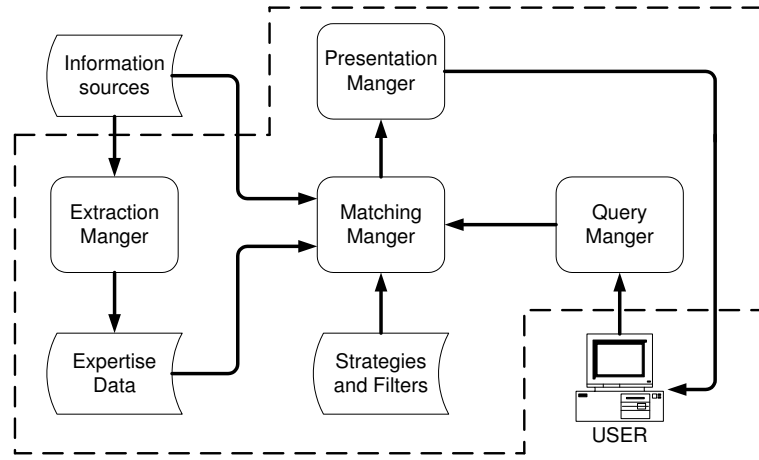


Fig. 1. The framework for the Expert Finder described in this paper.

Expertise Extraction. The extraction manager is responsible for identifying potential information sources containing expertise evidence. This will typically include shared or personal workspaces, document storage systems, or e-mail archives. It is recognised that each repository attached to the Expert Finder will require a customised API. However this is less complex than a custom system for each organisation or part of an organisation.

Expertise Modelling. Regardless of the origin of the information sources, i.e. shared or personal, the extracted evidence will need to be stored on a server for use by the Expert Finder. However, in order to address the heterogeneous information sources problem presented above, the process of building the expertise model needs to detail the relationship of a given source to an expert. For instance, the occurrence of a term in the title of a document should be given a higher weighting, compared to its occurrence in the body of a document, this approach is discussed in Section 4.1.

Expertise Matching. The matching manager is initiated by a user request. Through the query manager, the user can select the strategies and/or filters required in selecting expertise for their needs. In order to achieve the goal of selection, the matching manager provides access to the information space that maintains personal and other data, (i.e. departmental affiliation) which can be used as criteria. Based on these criteria, the set of candidate experts are re-ordered to a specified rank and/or deleted from the set to produce a refined list of experts. As an example, users may wish to remove individuals who left the organisation. This can be achieved by using the human resource database to identify dates of employment.

User Interface. At the conclusion of the process a ranked list of experts coupled with evidence (e.g. documents) retrieved for expertise modelling is submitted to the user. As described earlier, increasing the system's transparency can provide by providing the expertise recognition logic and expertise evidence to the users, this in turn increases the users' trust in the system.

4 A Prototype Expert Finder

The prototype Expert Finder was developed as part of an ongoing project at Southampton whose objective was to develop tools to support the activities of the design engineering, particularly in areas of knowledge capture, sharing and reuse. The *Relational Expert Finder System* based on the framework discussed earlier has been implemented in Java, the key features are summarised below.

4.1 Expertise Extraction

To implement the prototype system, we used data supplied by a major UK manufacturer, in the form of their internal publication database. This corresponds to a total of over 170,000 entries, covering a time period of fifty years. All the records in the database were entered manually, and ranged from technical reports to departmental memoranda. In practice neither the data source could be guaranteed to be correctly maintained nor can the entries in any fields be guaranteed to be valid and/or consistent. The supplied data was placed into a database management system with Structured Query Language (SQL) support. Careful design was required due to the database size, otherwise, accessing the resource could be very time consuming and the responses resulting from queries could easily overwhelm the system. A considerable amount of effort was taken to remove inconsistent or duplicated entries from the database, to improve the performance of the expertise modelling.

Expertise Modelling. Expertise models were created using text modelling algorithms based on the vector space model. The vector space model has advantages over other text modelling methods in overcoming overfitting, and dimensionality large problems, [15]. This means the model can easily adapt to dynamic

environments where additions to the database are frequent. TFIDF is a popular function employed by most vector space modelling applications, [16].

However, the TFIDF function ignores the document's structural elements, i.e. title or body, and treats all the text contained in that document as a bag of words, e.g. unstructured text. In view of this, the authors modified the TFIDF function to account for the structural elements,

$$w'(t_k, d_j) = w(t_k, d_j) + s_j$$

where $w'(t_k, d_j)$ denotes the weight of the term t_k in a document d_j , $w(t_k, d_j)$ denotes the weight of t_k in d_j calculated by the TFIDF function, and s_j denotes the weight of the structural element in which t_k occurs at least once. The value of s_j was determined heuristically as part of the development process. As part of the process the raw documents were processed by a number of tools to extract the text under the various headings, i.e. project name, authors, report abstract, etc., [17]. The equation above reflects the fact that the more often a term occurs in a document's major structural elements, such as title, abstract, etc. the more it is representative of its contents, and the more documents in which the term occurs, the less discriminating it is. The $w'(t_k, d_j)$ term is normalised for the purpose of comparison during the expertise matching process. It should be noted that, in this application, stemming is not performed on the text in the data source, since it produces ambiguous results, particularly when dealing with technical terms. The finalised expertise models will contain indexes resulting from the calculation of term space coverage and application of dimension reduction mechanisms, and are stored prior to expertise matching.

Expertise Matching. The querying process as currently implemented uses the AND Boolean operator. Using the query terms together with the AND operator, expertise models containing the query terms are identified, and then combined by taking the intersection of the sets of expertise models retrieved earlier. In order to simplify the query input interface, a query is made without any formal syntax, hence if *crack propagation* is entered it was interpreted as being a request that matches *crack* AND *propagation*.

In using the Expert Finder the users should select the appropriate expert finding strategies according to their needs, and the Expert Finder should only provide analysis functionality. This can be achieved by providing an interface to access the expertise recognition logic, i.e. how expertise is identified using various strategies. The system implementation includes two expert finding strategies. The first strategy is based on the concept of organisational awareness [18], in which the system only considers an individual as an expert if they are linked to a large number of relevant documents. However such an approach tends to reflect the interests of experts instead of their competency levels, and is only appropriate for users who seek individuals as sources of information.

The second strategy identifies experts by the importance of terms (supplied by the user) in documents. The importance of such terms is calculated using a vector space algorithm, as defined by the equation above. This algorithm not

only reflects the importance of terms in relation to their occurrences in a set of documents, it also indicates the terms' importance in relation to the document structures. Therefore, the computed terms' importance can then be used as indicators for experts' competency. This strategy, coupled with or without the first strategy, can be used to find individuals who can perform given organisational functions, such as collaboration in co-authoring documents or launching projects.

Two filters were implemented to refine the list generated from the above strategies. If the number of experts returned by the system overwhelms the user, then a filter is used to display only the top twenty experts in the recommendation list. Since the database covers over fifty years of work, the information will be out-dated for certain queries or requirements. In view of this, users are also given the option to filter the list of experts by document publication dates.

User Interface. After the expertise matching process has been completed, the ranked experts' names with their associated scores are displayed. The scores for ranking the experts are supplied to assist the users in analysing the recommendation. An interface is also provided to access the list of documents selected for expertise recognition, and the documents are grouped by authors for browsing purposes. This approach not only satisfies users' requirements in locating experts as information sources, it also allows users to evaluate the expertise recognition logic for themselves, hence, giving them greater trust as regards the recommendation.

4.2 Evaluation

Expertise retrieval effectiveness can be measured in terms of the information retrieval notions of *precision* and *recall*. In this paper, *precision* is defined as the probability that a random expert suggested in response to a query is correct; analogously, *recall* is defined as the probability that, if a random expert is recommended in response to a query, this decision is accepted. Neither precision nor recall make any sense in isolation, as it is widely recognised that higher levels of precision may be obtained at the price of a low recall. Hence, the evaluation requires a combined effectiveness measure, determined by precision and recall. The question then arises as to how the combined measures can be used to determine whether *system A* is better than *system B*. Generally, *system A* is assumed to be better than *system B* if, at every recall point, *system A*'s precision value is higher than *system B*'s, [16]. If this does not hold, then the precision values for selected recall values are averaged and compared.

Effectiveness is computed as an *11-point interpolated average precision* [19], this measure is widely adopted in the case of document-pivoted categorisation. For obtaining estimates of precision and recall relative to multiple decisions when two or more queries are submitted, microaveraging was adopted as a global evaluation method, [20].

The evaluation allowed us to compare the effectiveness of the proposed system against the system previously reported by Crowder and Hughes, [10]. However,

this system approach the modelling of expertise differently as the *KCSR Expert Finder* ignores such structural information by representing its expertise models using full-text indexes. In order for the experimental results on the two Expert Finders to be directly comparable, the experiments were performed using identical source databases and queries.


Test Data The effectiveness of an Expert Finder can be evaluated by users relative to specific contexts. The most likely context is their experience accumulated from the workplace. In order to measure the effectiveness of the *KCSR Expert Finder* and the *Relational Expert Finder System* in retrieving experts, a set of questions that can provide contextualised problem statements was obtained. A total of nine possible users, drawn from the upper half of the company’s engineering career structure, were interviewed to obtain sample questions and names of the individuals who they believed to have the expertise in answering these questions. The sample queries, Table 1, were chosen, as they gave a good spread of topics, ranging from the general to the highly specific.

Table 1. Exemplar questions obtained from the user community. For each question the users also provided a number of recommendation of whom they considered to be experts across the company.


Number	Sample Question
1	What are the operational issues of a specific systems?
2	What is the condition of specific parts at overhaul?
3	How does oil behave in hydraulic seals?
4	What are the positive characteristics to be noted when checking a simulation?
5	The procedure for modelling a turbine component for analysis.
6	Defining a turbine component ratio.
7	Requirements for an impact containment casing.
8	What are the properties of this material.
9	What are the resistant loads in designing a specific component?

Results The experts recommended by the *KCSR Expert Finder* and *Relational Expert Finder System* were compared with those identified by the user community. It was noted that:

- In the initial evaluation queries based on questions 1, 2 and 6 failed to return any matches with both systems.
- Queries based on questions 3,4,5 and 8 failed to return any matches with the *KCSR Expert Finder*.
- Matches were obtained between both Expert Finders and the user community for Questions 7 and 9. Figure 2, show how the expertise retrieval precision varies with recall thresholds.



(a) Question 7



(b) Question 9

Fig. 2. Precision-recall curves showing comparative results between the *KCSR Expert Finder* – dotted line – and the *Relational Expert Finder System* – solid line, for two specific searches.

In order to compare the systems' effectiveness in finding the overall values for each of the precision-recall curves were calculated, Table 2. Additionally, the global effectiveness values using the microaveraging method, values of **0.12** for the *KCSR Expert Finder* and **0.28** for the *Relational Expert Finder System* were obtained. These values were used to determine whether the result of the comparison between the two Expert Finders were significant.

Table 2. Average interpolated precision values for specific queries based on the sample questions

Specific query based on question number:	Average Interpolated Precision	
	Expert Finder	Relational Expert Finder System
1	0	0
2	0	0
3	0	0.17
4	0.16	0.45
5	0	0.17
6	0	0
7	0.55	0.61
8	0	0.15
9	0.2	1.0

In the precision-recall curves shown in Figure 2, the *Relational Expert Finder System's* precision values are higher than the *KCSR Expert Finder's* at every recall threshold. In question 9, the difference in the precision values between the two Expert Finders is most evident. As seen in Table 2, the *Relational Expert Finder System* has a higher success rate in retrieving experts who were recognised by the users community. Using the *KCSR Expert Finder's* microaveraged effectiveness value as the reference in identifying experts who have the relevant expertise for a specific interpretation of the sample questions, the improvement is 16%, which is statistically significant.

5 Concluding Comments

We have presented an approach to using capturing organisational knowledge for expert finding. A framework for Expert Finders is proposed by the authors that extends the relationships between information sources and expertise models. We have noted two significant activities for why experts need to be located, either users seek experts as sources of information or as collaborators in specific activities. Hence, we suggested that the Expert Finder should provide analysis functionality, since it is the users who selects the appropriate expert finding strategies depending on their needs. This was incorporated into the Expert Finders by providing interfaces to access the expertise recognition logic and evidence.

In contrast to many Expert Finders that were designed to solve a particular problem within a specific organisational environment, our framework is

both flexible and modular; so that its components can be easily extended and replaced depending on requirements. The framework enables an implemented Expert Finder to be interoperable with other organisational systems via appropriate API. This approach allows generated expertise data to be shared across across an organisation.

The Expert Finder developed by the authors has been compared with a similar system based on full-text indexing system which ignores structural information when analysing the source documentation. Although both of these systems shared the same data source, in initial testing the system developed by the authors outperformed the full-text indexing system in terms of expert retrieval effectiveness for a limited number of test cases. It is our view that the improvement in locating experts through the methodologies embodies in the *Relational Expert Finder System* will translate to a reduction in costs within an organisation as the correct expert is located more rapidly.

Although the approach demonstrates that the mechanism for extracting expertise data can be automated, it however trades the problems of increased workload and subjective self-assessments with the problem associated with ‘dirty’ data. The authors invested a considerable amount of time and effort to formulate techniques for validating data and folding it to a format that can be processed by the systems. As such, the validity and consistency of the data plays an important role in determining the performance of expertise retrieval, and should be considered when deploying an Expert Finder.

6 Acknowledgement

The authors acknowledge the studentship for Yee-Wai Sim from the Faculty of Engineering, Science and Mathematics, University of Southampton.

References

1. Shadbolt, N.R., O’Hara, K.: Actuality: An overview of the aims, ambitious and assumptions of the advanced knowledge technologies interdisciplinary research collaboration. In Shadbolt, N.R., ed.: *Advanced Knowledge Technologies: Selected Papers 2003*. AKT (2003) 1–11
2. Becerra-Fernandez, I.: The role of artificial intelligence technologies in the implementation of People-Finder knowledge mangement tools. *Knowledge Based Systems* **13** (2000) 315–320
3. Streeter, L.A., Lochbaum, K.E.: An expert/expert locating system based on automatic representation of semantic structure. In: *Proceedings of the Fourth IEEE Conference on Artificial Intelligence Applications*, Computer Society of the IEEE, San Diego, CA (1988) 345–349
4. Kautz, H.A., Selman, B., Shah, M.: Referral Web: Combining socail networks and collaberative filtering. *Communications of ACM* **40** (1997) 63–65
5. Krulwich, B., Burkey, C.: The ContactFinder agent: Answering bulletin board questions with referrals. In: *Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference*. Volume 1., Portland, OR (1996) 10–15

6. Foner, I.N.: Yenta: A multi-agent referral-based matchmaking system. In: Proceedings of the First International Conference on Autonomous Agents, Marina del Rey, CA (1997) 301–307
7. Pikrakis, A., Bitsikas, T., Sfakianakis, S., Hatzopoulos, M., DeRoure, D., Hall, W., Reich, S., Hill, G., Stairmand, M.: MEMOIR - software agents for finding similar users by trails. In: Proceedings of the Third International Conference and Exhibition on the Practical Application of Intelligent Agents and Multi-agents, London, UK (1998) 453–466
8. McDonald, D.W., Ackerman, M.S.: Expertise recommender: a flexible recommendation system and architecture. In: ACM 2000 Conference on Computer Supported Cooperative Work. ACM, New York (2000) 231–40
9. Vivacqua, A.S.: Agents for expertise location. In: Proceedings of the AAAI Spring Symposium on Intelligent Agents in Cyberspace, Stanford, CA (1999) 9–13
10. Crowder, R., Hughes, G., Hall, W.: An agent based approach to finding expertise. In Karagiannis, D., Reimer, U., eds.: Proceedings Fourth International Conference on Practical Aspects of Knowledge Management. Volume LNAI 2569. Springer-Verlag, Berlin Heidelberg (2002) 179–88
11. Kanfer, A., Sweet, J., Schlosser, A.: Humanizing the net: Social navigation with a 'know-who' e-mail agent. In: Proceedings of the Third Conference on Human Factors and the Web, Denver, Co (1997) Available at <http://archive.ncsa.uiuc.edu/edu/trg/human/>.
12. Cohen, D., Prusak, L.: Networks and communities. In Cohen, D., Prusak, L., eds.: Good Company: How Social Capital Makes Organizational Work. Harvard Business School Press (2001)
13. Mattox, D., Maybury, M., Morey, D.: Enterprise expert and knowledge discovery. In: Proceedings of the 8th International Conference on Human-Computer Interaction (HCI International'99), Munich, Germany (1999) 303–7
14. Crowder, R., Bracewell, R., Hughes, G., Kerr, M., Knott, D., Moss, M., Clegg, C., Hall, W., Wallace, K., Waterson, P.: A future vision for the engineering design environment. A future sociotechnical scenario. In: International Conference on Engineering Design, ICED2003, Stockholm (2003)
15. Joachims, T.: Estimating the generalization performance of a SVM efficiently. In: Proceedings of the Seventeenth International Conference on Machine Learning, Stanford, CA (2000) 431–438
16. Sebastiani, F.: Machine learning in automated text categorization. *ACM Computing Surveys* **34** (2002) 1–47
17. Sim, Y.W.: Capturing Organisational Knowledge from Documentation for Expert Finding. PhD thesis, School of Electronics and Computer Science, University of Southampton (2004)
18. Maybury, M., Amore, D., House, R.: Awareness of organizational expertise. Technical papers, MITRE (2000)
19. Salton, G.: Advanced information-retrieval models. In Salton, G., ed.: Automatic Text Processing. Addison-Wesley Publishing Company (1989)
20. Lewis, D.D.: Evaluating text categorization. In: Proceedings of Speech and Natural Language Workshop. (1991) 312–318