

Expert Finding Systems for Organizations: Problem and Domain Analysis and the DEMOIR Approach

Dawit YIMAM-SEID, Alfred KOBSA

Department of Information and Computer Science,
University of California, Irvine, CA 92697-3425, U.S.A.
e-mail: *dseid, kobsa@uci.edu*

Computer systems that augment the process of finding the right expert for a given problem in an organization or world-wide are becoming feasible more than ever before, thanks to the prevalence of corporate Intranets and the Internet. This paper investigates such systems in two parts. We first explore the expert finding problem in depth, review and analyze existing systems in this domain, and suggest a domain model that can serve as a framework for design and development decisions. Based on our analyses of the problem and solution spaces, we then bring to light the gaps that remain to be addressed. Finally, we present our approach called DEMOIR, which is a modular architecture for expert finding systems that is based on a centralized expertise modeling server while also incorporating decentralized components for expertise information gathering and exploitation.

Keywords -- Expert finders, expert recommenders, expertise management, knowledge management, knowledge sharing, expertise modeling, domain analysis

1. INTRODUCTION

Motivated by advances in information technology, organizations are giving more emphasis to the capitalization of the increasing mass of knowledge they accumulate in the course of their business. However, as noted by Stewart [1], the attempt to put all corporate knowledge on one huge server in the style of the 18th-century French encyclopedists is doomed to fail. Stewart continues to assert that the real value of information systems is rather in connecting people to people, so they can share what expertise and knowledge they have at the moment, given that the cutting edge is always changing. Studies into the information-seeking behavior of people working in information-intensive fields also show that people searching for information commonly explore personal communications prior to using formal sources [2,3].

Thus, if technology is to foster the effective utilization of the whole range of knowledge in organizations, it has to be able to support not only access to explicitly documented knowledge but, most importantly, tacit knowledge held by individuals. By enhancing the visibility and traceability of such knowledge, technology can help catalyze collaboration and knowledge sharing among its holders, both within and between organizations. Moreover, the ability to quickly find information on expertise of people can play critical roles in fostering the formation and sustenance of virtual organizations/enterprises, communities of practice, expertise networks and the like.

The recognition of the need to foster sharing of expertise has lately spawned research efforts in, among others, the Knowledge Management and Computer Supported Collaborative Work (CSCW) communities. Concepts like expertise capitalization/leveraging, skill mining, competence management, intellectual capital management, expertise networks, knowledge sharing systems and the like are being widely discussed. However, rarely do the discussions go beyond generic assessments and touch concrete design and implementation challenges. Moreover, systems that support such processes are at a very low level of development compared to those that support access to documented information.

In this paper, we investigate the expert finding problem as well as the various design and development issues and present the solution we are pursuing. We begin with a brief analysis of the expert finding problem followed by a review of the previous works on expert finding aids. Then, after contextualizing expert finding aids in broader organizational information systems, we present the results of our domain analysis which allows for a more concrete discussion of current shortcomings and possible solutions. Finally, we describe the approach we adopted, followed by an outline of our future work.

2. PROBLEM ANALYSIS

A variety of practical scenarios of organizational situations that lead to expert seeking have been extensively presented in the literature (see e.g. [4,5,6,7,8,9]). McDonald and Ackerman [5] also provided empirical insight into the intricate organizational, personal and social processes involved in expert seeking. In our extensive survey of the knowledge management literature, we found a widespread acceptance of the fact that an effective knowledge management solution should not only aim at making knowledge that is contained in documents accessible, but also

knowledge in people's heads (e.g., [10,11]). However, still a lot remains to be understood about expert seeking so as to enable the development of effective automated aids.

2.1. What prompts expert seeking: information need and expertise need

Based on interviews we conducted with researchers at a major research institution as well as extrapolation of observations in the relevant literature, we identified two main motives for seeking an expert, namely *as a source of information* and *as someone who can perform a given organizational or social function*. This categorization, though fuzzy and sometimes overlapping, proves to be useful in analyzing the goals of automated expert finders as we will discuss later in this paper.

People may seek an expert *as a source of information* to complement or replace other sources such as documents and databases in various ways. The following cases are just a few possible scenarios that entail seeking experts as information sources:

- (1) *Access to non-documented information*. Not all information in organizations can possibly be explicitly documented; much important information can only be transferred through apprenticeship, experience and informal conversation [12]. Documented sources can hardly help in novel situations, solutions-in-progress or situations that are too focused. Besides, there are cases when the needed information is deliberately not made publicly available (in documents or on the Internet) for economic, social and political reasons [7].
- (2) *Specification need*. The information users need to solve a problem is often not well specified and requires a dialogue with an expert to precisely identify it. For example, so as to effectively retrieve information from a retrieval system, a user must be able to convert her problem into search statements. In the information retrieval literature, this conversion is referred to as a transition from the 'exploration' phase to the 'formulation' phase in the information search process [13,14]. This transition is known to be a pivotal point that is also accompanied by an emotional drift on the part of the seeker from confusion, frustration and doubt, to clarity. Naturally, people resort to seeking experts to painlessly carry them through this drift.
- (3) *Leveraging on others' expertise (group efficiency)*. Users often want to minimize the effort and time required to find some piece of information (particularly since frequently what is "lots of work" for the user is "little work" for the expert). This also includes using experts as effective and reliable information filters in selecting the useful information from the huge mass of information available.
- (4) *Interpretation need*. Users are often not interested in a piece of information *per se* but rather its situated (contextualized) interpretation or the implications that can be derived from it in the given situation. It might also be the case that the user fails to understand the information even if he/she manages to find it from documents.
- (5) *Socialization need*. Users may prefer the human dimension that is involved in asking an expert as opposed to interacting with documents and computers.

The second category of expert seeking motives, namely the need for an expert who can perform a given function, are generally characterized by a more structured search than seeking experts as information sources. In these cases, the need is for someone who possesses a certain type of expertise required to play a role in a particular undertaking. This occurs in cases when

either the involvement of the expert himself/herself in a given activity is required or a continued partnership between the seeker and the expert is intended. Hence, what is involved is more than just acquiring a piece of information. The following are just a few of such cases:

- seeking a consultant, employee or contractor;
- seeking a collaborator, team member, community or committee member, or a journal or conference reviewer;
- seeking a speaker, presenter, researcher, promoter, and expert interviewee for media.

We can also view the above two categories of the expert finding motives as related to “*information need*” and “*expertise need*”, respectively. In the former case, a user is required to specify her information need while in the latter a more complex specification of “expertise need” replete with qualitative and quantitative parameters is involved. In other words, when seeking experts as sources of information one is mainly interested in the question “who knows about topic x?”. When seeking experts who can perform an organizational function, one is equally interested in the other questions like: “how much (well) does y know about topic x?”, “What else does y know?”, “How does y compare with others in his knowledge of x?” etc. Search processes for experts who are supposed to perform an organizational function will therefore often be performed more carefully and in a more formalized manner than the search for experts who are supposed to deliver information.

Often, it is difficult to draw a sharp distinction between these two categories of motives since they may occur at the same time or an information need may later turn into an expertise need. For example, someone who served as a source of information may also be needed to act as an “ultimate advisor” in the correct interpretation of the piece of information as well as the consequences for the concrete problem of the information seeker. Nonetheless, as we will show later in this paper, studying them distinctly seems to be useful.

2.2. Internal versus external expert seeking

In organizations, expert seeking (and consequently, the benefits of finding them) can also be viewed from an *internal* and an *external* point of view. There are a number of reasons for wanting to know who knows what within an organization (as quickly as possible at that), including knowledge sharing, team formation, project launching, etc. An organization also benefits if external entities can easily discern the expertise of its staff, as this fosters collaboration, cross-organizational networking, a better image, etc. For example, many organizations can deliver efficient customer help services if the customers, or their contact points in the organization, can easily trace and direct their queries to the appropriate expert. Likewise, academic and research institutes want the industry, the public, and potential research sponsors and research collaborators to know about and make use of their staffs’ expertise.

Lewis [15] notes that most organizations consider expert-locating services as necessary for external expert seeking only, and not internal expert seeking. This is probably based on the presumption that employees know internal experts well. But, based on a recent survey, Lewis indicates that employees also benefit from support in seeking experts within the organization, as do customers. Internal expert location is more important the larger an organization is, the more geographically distributed it is, and the less homogeneous the composition of its members is (e.g., strict division borders, very different knowledge backgrounds, different histories due to company mergers, etc.) All three factors impair informal contacts, which so far are the major

ways for disseminating knowledge about others' expertise within the organization.

2.3. Automated Support: Traditional Approaches

Whatever their motives are, seekers of experts need a range of information regarding peoples' expertise. They need to know whether a person who can answer their queries or meets their criteria exists, how extensive his/her knowledge or experience is, whether there are other persons who could serve the same purpose, how s/he compares with others in the field, how the person can be accessed (contacted), etc. This, in turn, calls for a mechanism that gathers and makes such information accessible. However, doing this manually is obviously a time-consuming and laborious undertaking, making automated aids invaluable.

One way of providing automated assistance is the development of expert databases (aka "knowledge directories", "knowledge maps") through manual entry of expertise data. This is exactly what many organizations commonly did in the past. Microsoft's SPUD [16], Hewlett-Packard's CONNEX (<http://www.carrozza.com/connex/>) and the SAGE People Finder (<http://sage.fiu.edu/Mega-Source.htm>) are examples of this approach. Similarly, manual data entry is employed in skill inventory systems like Skillview (<http://www.skillview.com/>), which are common in the knowledge and human resource management domain.

However, expert databases are known to suffer from numerous shortcomings, including the following:

- (1) Manually developing the databases is a labor intensive and expensive task.
- (2) These databases are usually based on the propensity of experts to devote their time to initially provide a detailed description of their expertise.
- (3) Due to the continuous change in people's specific expertise and skills (especially organization-specific and environment-specific skills), most expert databases are quickly outdated. Updating expert databases manually is resource-consuming for the organization and time-consuming for the experts.
- (4) The expertise descriptions are usually incomplete and general. On the other hand, expert-related queries are very fine-grained and specific, and replete with various qualitative requirements [7]. Normally, identifying an individual's specific area of expertise requires a rather detailed analysis of his works, records of activity, etc.

Personal web pages, through which individuals provide information about themselves and their expertise, are also used in tracing people with needed expertise. Searching the Web or the organizational Intranet by topic may lead us to such web pages if the experts took their time to maintain up-to-date pages and if the web pages happened to contain the supplied keywords. While such pages appear to be better sources of information on experts than expert databases, simply using search engines to trace an expert is not normally an effective approach. For one thing, search on a topic returns a large number of various resources, making the selection of personal web pages currently a time consuming task. As the search process is general and based on simple keyword matching, it also may not always lead to the relevant experts. Moreover, it is entirely the task of the user to extract and compile all the required data to conduct such analyses as identifying the best expert, determining the accessibility of the expert, his/her relation with other experts, etc.

3. AUTOMATIC EXPERT FINDERS

The shortcomings of the above approaches coupled with the availability of large electronic repositories of organizational and personal records have led to the suggestion of more helpful systems known as *expert finders* or *expert/expertise recommenders*. These systems aim at mitigating the above shortcomings by trying to automatically discover up-to-date expertise information from implicit/secondary sources instead of relying on experts and/or other human sources only, although experts and their proxies can still complement/refine the automatically generated expertise information.

Attempts to develop systems that exploit implicit evidence of expertise to augment the process of finding the right expert dates back at least to the visionary work of Maron and his colleagues [17]. Their experimental system, called *HelpNet*, accepts requests for information and responds with names of people ranked by the probability that the individual will provide a satisfactory answer. This probability is computed using probabilistic models of information retrieval which combine the estimation of people's expertise in answering a question on a topic with the probability that a given user would be satisfied with the response provided by the source. To do this, the system constructs a profile by asking people to indicate their expertise by selecting from a list of topics along with a probability estimate of their ability to provide a satisfactory answer to questions on that topic. Maron and his colleagues envisioned such systems enabling the emergence of "a large, active and fruitful future network of informationally rich people providing help to one another".

The next notable system was the *Expert/Expert-Locator (EEL)* (also called "Bellcore Advisor" or "Who Knows") [18,19]. This system takes natural-language requests for technical help or advice and, using Latent Semantic Indexing (LSI), returns pointers to "nearby" research groups. The system constructs expertise index of research groups based on a representative collection of the technical documents they produce. In a test using EEL [20], descriptions of individual expert's current projects are used as query to compare how well the system managed to predict which of the 480 Bellcore work groups an expert belonged to.

ContactFinder [21,22] is an intelligent agent that monitors discussion groups and extracts contacts in some specific areas which it then uses to respond to postings that ask questions with referrals to relevant contacts. This system uses various heuristics both to extract contacts from email messages and to identify those postings that ask for technical help. Mattox *et al* [23] developed an expert finder system that exploits MITRE Corporation's corporate Intranet to enable the location of relevant experts by identifying relationships between experts and documents.

Kautz and his group approached the expert finding problem from a different perspective [7,24,25]. Their work is based on the observation that expert finding is an inherently complex process. Hence, they postulate that the best way of finding an expert is through what is called "referral chaining" whereby a seeker finds the needed expert through referral by colleagues. Schwartz and Wood [26] have also tried to enable locating people with related interests or expertise by analyzing the graph - which they called "specialization subgraph" - formed by email communication patterns (rather than their contents). They use a set of heuristic algorithms to uncover shared-interest relationships among people and derive a list of people who shared a given interest.

Vivacqua [8] describes an expert finder agent that is reminiscent of the earlier ideas in [7] in that it employs a personal agent that both profiles the expertise of a user and searches for other

experts when help is needed. This agent was developed for the domain of Java programming. Yenta [27] is also an agent-based system that creates personal profiles of peoples' interests, and uses inter-agent communication to find people that have interests close to a query as well as cluster people based on their shared interests. Expertise Recommender (ER) by McDonald and Ackerman [28,29], is a system composed of components that implement the functionality and heuristics they identified from their field study in a software development company [5]. ER provides four major components called profiling supervisor, identification supervisor, selection supervisor and interaction management that maintain and manipulate expertise profiles.

There also exist some recent commercial knowledge management systems which provide features that support expertise profiling and retrieval in organizations. *Agentware Knowledge Server*TM from Autonomy (<http://www.autonomy.com>), for example, includes a feature that identifies employees' areas of expertise based on the documents they access in and submit to the organizational Intranet. Likewise, KnowledgeMailTM from Tacit Knowledge Systems Inc. (<http://www.tacit.com/>) supports finding experts by building profiles of people's skills from their e-mail messages as well as other documents they submit. AskMe EnterpriseTM is also another similar product (<http://www.askmecorp.com>).

Also noteworthy are organizational memory systems like *Answer Garden* [30], *The Chicago Information Exchange* (CIE) system [31] and Orbital Software's (<http://www.orbitalsw.com/>) Organik that include expert finding components. These systems are basically question answering and routing systems that respond to requests for technical help by retrieving stored question and answer pairs, but also provide facilities to route unanswered questions to a defined group of experts. Similar to such systems are works that incorporate expert finding capabilities as part of peer help systems in education and training environments. PHelpS [32,33] is one such system that enables finding peer helpers among course participants or task-specific groups of people.

There are also other systems that considered people's WWW browsing patterns as indicators of expertise. The *MEMOIR* system [34] enables one to find experts in a certain domain based on a log database called "trailbase", which consists of URLs people visited as well as keywords for these pages. Likewise, the *Expert Browser* [35] enables people in organizations to use logged and indexed browse paths of experts to find relevant documents. (However, the authors only describe those features that require the user to know the expert or her group *a priori*). In general, as regards such systems, one may question their basic assumptions on the ground that browsing shows more of the *interest* of someone rather than *expertise*.

Apart from such implemented systems, there are also a few authors who tried to formulate techniques that enable finding experts by mining organizational and personal document repositories. Kimbrough and Oliver [36] suggest a retrieval method called associative retrieval that uses binary term-weight matrix-based procedures called DCB algorithms to identify resources (including people) in organizations that are closely linked to a given documented issue. Kanfer *et al* [37] present the idea of an agent called 'know-who agent' that can use TF-IDF weighting to trace relevant experts based on e-mails they receive.

4. POSITIONING AUTOMATIC EXPERT FINDERS

Before going on to the domain analysis where we consider the functionality of expert finders, let us first survey the potential application contexts for these systems, with particular reference to

other related organizational systems and services. This survey outlines the bounds of the expert finders' domain and crudely represents the 'context analysis' phase required prior to domain analysis.

As part of organizational information systems, expert finders can either stand on their own or form part of other broader organizational systems. We believe that their great potential is unleashed only when used in integration with other organizational information systems, namely knowledge management systems, recommender systems, CSCW systems, and electronic markets for human expertise.

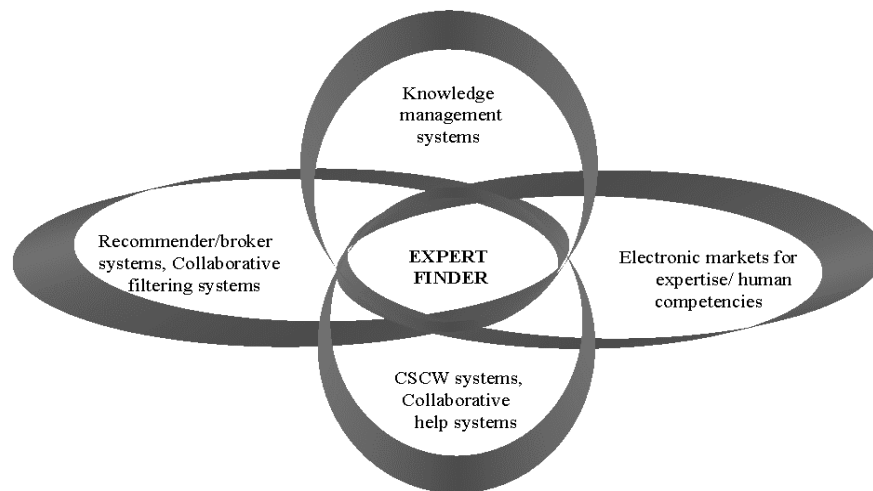


Figure 1. Expert finders integrated in organizational information systems

As mentioned above, expert finding capabilities form an important part of knowledge management systems whose aim is provide access to knowledge in all forms, including knowledge held by people. Davenport [37] called this the "hybrid approach to knowledge management". Kautz *et al* [7] also discuss the importance of integrating both the "ask a program/document" and "ask a person" paradigms into information seeking. These two approaches are mostly used in an inter-dependent manner, i.e. one is used to find the other [2]. Organizational memory systems like question answering and routing systems (e.g. Answer Garden[30]) can be cited as one realization of this paradigm which necessarily calls for expert finding capabilities.

Social recommender/collaborative filtering systems, too, can enhance their services by integrating expert finding support. For example, as Höök *et al* [38] noted, a user of such a system might be much more interested in what particular experts regard as important information than in the recommendation of a large group of peers. One can reasonably conjecture that users may want to judge the relevance of a piece of information based on its *quality* (as reflected in the expertise of the recommender) in addition to *quantity* (the number of people recommending it). Obviously, expert finders can play important roles here. For example, their integration allows

users to inquire the extent of expertise of recommenders or enable the recommender systems to factor expertise of the recommenders in their weighting schemes.

Those systems that fall under the rubric of computer supported cooperative work (CSCW) systems, too, can integrate expert finding components to permit the identification of individuals to collaborate with, in addition to supporting collaboration among those who have already identified each other [39].

With expert finding capabilities as a core, so-called “electronic markets for expertise/human competencies” wherein expertise sellers and buyers transact, are suggested by Lang and Pigneur [40]. Moukas *et al* [41] propose a software-agent based infrastructure that allows such transaction to take place. Recently booming services like XpertSite.com (<http://www.XpertSite.com>), exp.com (<http://www.advoco.com>) and ExpertCity.com (<http://www.ExpertCity.com>) where users can post their questions and get quick expert answers for pay are other examples of such a market.

5. DOMAIN ANALYSIS

Domain analysis is a method for analyzing an application domain by studying and characterizing existing systems, emerging technologies, and development histories within a domain of interest [42]. In domain analysis, common characteristics from similar systems are generalized, objects and operations that are common to all systems within the same domain and that vary from system to system are identified, and a domain model is defined to describe their relationships [43]. For our purpose, we used the analysis method based on faceted classification suggested in [43] and later improved by Birk [44].

Before proceeding, we would like to state that our aim here is far less ambitious than conducting a full-fledged domain analysis of expert finding systems (for which we concur the expert finding domain is not mature enough) as is done in software reuse efforts, where the approach is mainly employed. Our aims are (1) to come up with a systematic and structured way to characterize expert finding systems, and (2) to investigate whether better systems that can be used in a wide variety of contexts can be designed. We analyze the literature of five expert finder systems [22,23,24,8,18] to identify and define the feature space of the domain and we map both implemented and potentially applicable technology to develop an intuitive domain model. Along the way, we also introduce terminology and structure that summarize the common as well as distinct characteristics of these systems.

5.1. An Intuitive Domain Model

According to Birk [44], an intuitive domain model is a faceted classification scheme that summarizes findings of a domain analysis and describes the domain using terms from the language of software professionals. Each *facet* in the domain model is represented by a *domain factor* (attribute shared by systems). Each domain factor is defined through a set of discrete values (“*possible values*”). A facet of a concrete domain is characterized through one or more values (“*actual values*”) from the set of possible values. A domain factor together with its actual value(s) is called *domain characteristic*. For the expert finding systems domain, we identified seven facets/domain factors as shown in Table 1. Each column represents one facet/domain factor and its possible values, which are possible implementations of the respective domain factor. Hence,

we can classify a given expert finding system using its domain characteristics that consist of the domain factors along with their values.

In the following discussion of the seven domain factors, we use the term *expertise indicator* to refer to the data gathered as evidence of expertise and input to a process of analysis to generate/infer the expertise of people. By *expertise indicator source*, we refer to any source of expertise evidence that includes different types of documents, databases, individuals, etc.

a) Basis for expertise recognition. Expert finding systems use various evidences as indicators of expertise. In general, these evidences can be grouped as explicit and implicit evidences. Although explicit evidences are very helpful when available, the focus of automatic expert finders is on employing implicit evidences since they are usually far more common and do not impose burdens on the experts. Some of the implicit evidences considered as indicators of expertise are document authorship (formal and informal, individually or in group), name occurrence in non-authored documents, use of information sources or system functionality, queries sent to an information retrieval system, and the departments/projects they work in.

In theory, all documents in an organization are potential sources of expertise evidence since they are created *by* or *about* somebody. However, for various reasons including computational difficulty, privacy, coverage, etc., not all sources can actually be used. Hence, expert finding systems need to employ a procedure for recognizing and gathering potential sources of expertise indicators. We refer to this mechanism as *source recognition logic*. This can be static in that the system mines a predefined type or format of sources, or it may be designed to dynamically identify sources of expertise information from the large mass and variety of electronic as well as human resources.

b) Expertise indicator extraction. Once the sources are identified and gathered, expertise indicators are extracted from them. The expertise indicator extraction techniques that are applied on the aforementioned resources can be grouped as domain knowledge independent or domain knowledge driven. An example of the latter is described in [8], where Java documentation is employed as a domain knowledge base to extract indicators of Java programming expertise.

c) Expertise models. We use the phrase *expertise model* to refer to a meta-description of an individual's expertise, skills, etc., or a link to an expert from an independently maintained organized structure of such descriptions. Being an inferred *model*, it differs, for example, from a record in an expertise database in that it involves a degree of uncertainty about how well it matches the actual expertise of people. Similarly, we use the phrase "expertise modeling" to refer to the process of generating the expertise model through an analysis of the expertise indicators, including possible inferences on the basis of the findings. Hence, there is always a degree of uncertainty in how well the expertise model matches the actual expertise of people. An *expert model (profile)* refers to all information peculiar to an individual expert including personal details, his/her expertise model (or links to one), his/her relations with other experts, etc.

Expertise models can be dynamically generated at query time from expertise indicator sources, or extracted and stored either by personal agents or as aggregated models to which experts are associated. In *query-time generated models*, source databases are searched employing information retrieval systems, and the matching sources are mined for experts and expertise indicators (e.g. [23]). In *personal-agent based systems*, expertise modeling is distributed to self-managing agents that belong to each individual expert and reside on his/her machine. As such,

the whole system becomes a multi-agent system composed of personal agents, which carry out the dual tasks of modeling expertise (from authored documents and other sources) as well as assisting their owners in searching other experts. These agents can be endowed with mobility as well as autonomy to do all the processing and control of expertise modeling on their own. In the approach of associating experts to *aggregated expertise models*, experts are linked to a pre-constructed or dynamically generated central representation of expertise which can be a kind of knowledge model (ontology), organizational structure, etc.

d) Query mechanisms. The system can either require the user to explicitly ask the system for experts or it may infer expertise need from users' communications, activities, errors, help-seeking actions, etc.

e) Matching operations. So as to identify experts, information and/or expertise needs are matched with expertise models and/or expert profiles using retrieval techniques including exact keyword matching or similarity matching like vector-space based methods. Furthermore, some inference mechanisms can be applied to either or both queries and expertise models. This will allow such inferences as "if x is expert on topic y, then she may also/not know topic z" or "if x knows topic y, then she may know someone who is expert on topic z".

f) Output presentation. Expert finding systems may present their outputs with a varying degree of details and functionality for further exploitation. Like in information retrieval systems, some mechanism of ranking the identified experts (which we call expert differentiation) is applied. The ranking of experts may be *multidimensional* in that various criteria need to be taken into account. Apart from relevant details about each returned expert, the organizational and social context and network of the experts can be employed. Some systems also provide the sources in which experts are mentioned, leaving the task of identifying the experts in them to the user.

g) Adaptation and learning operations. These operations enable the customization of output by identifying users' preference and domain knowledge as well as gathering user feedback (both positive and negative) to the system's output. The system should not only be able to identify those people with required expertise but also be able to identify the most suitable ones meeting the needs and context of the user. For instance, the system can employ user models (including expertise models) to compare the experts' competence level with that of the user, make user-tailored rankings, and attempt to describe expertise at a level of granularity that matches queries (if an ontology is used), etc. User feedback on system's accuracy, ranking scheme etc. can be gathered and used to refine its expertise and expert models. Since expertise is in part a social construct, the system can also gather expertise evaluations and ratings from users and incorporate it in its expert differentiation scheme. Users can also be allowed bring new experts into the system. This facet is implemented in only a few of the systems we studied. In [8], an expert profile of users is employed to enable expertise level matching. Mattox *et al* [23] mention their plan to enable their system to learn over time based on results of each query to the system.

Table 1. An Intuitive Domain Model of Expert Finding Systems

Basis for Expertise Recognition	Expertise Indicator Extraction Operation	Expertise Modeling	Query Mechanisms	Matching Operations	Output Presentation	Adaptation and Learning operations
Explicit	Domain knowledge driven	Query-time generated	Explicit query for expert	Exact/overlap matching	Ranked list of names	Adaptation using user's expertise models
- Self-declaration by expert	Domain knowledge independent	Personal-agent based (distributed)	- Keyword queries	Statistical/similarity-based matching (e.g. Latent Semantic Indexing)	- Uni-dimensional ranking	User modeling
- Professional (organizational) position	- Keywords from expert (or other people)	Association to aggregated (centralized) expertise model	- Natural language queries		- Multi-dimensional ranking	
Implicit	- Name-concept co-occurrence (proximity) in documents		- Parameterized queries (e.g., keyword plus social radius of experts)	Inference matching	Ranked list plus personal details	Gathering relevance feedback
- Document authorship	- Frequently used software features	- Expert-to-canonical topics	Induce information/ expert need		Experts contextualized in their social network	Community expertise evaluation/rating
- Name-topic co-occurrence in documents		- Expert-to-publications		- Inference on relation between concepts(ontology-based)		
- Projects worked on		- Expert-to-organizational structure	- Observation by personal agent	- Inference on relations between experts		
- Frequent usage of particular service/information source/software feature		- Expert-to-profile/expertise descriptor (database, web page, etc.)	- Analysis of communication (e.g., newsgroup posting)	- Inference on expertise level of people	Documents/organizational groups containing relevant experts	
- Citations received						

5.2 Discussion

In this section, we relate the expert finding problem as discussed in the beginning of this paper with the space of solutions that we reviewed and analyzed above, and describe the gaps that came to light during this process.

As we saw in the problem analysis, a user turns to an expert as a source of information, as a candidate for a given organizational role, or a mixture of both. Each of these purposes imposes its own requirements on the system functionality. As noted before, a user with an information need tends to seek experts prior to documents. The reasons for this includes using them as filters (guides) to documents, avoiding searching for documents altogether, or when the user is interested in the situated interpretation and application of a piece of information rather than the mere information itself. Such needs imply a role for expert finders of giving the user an expert-oriented access to the information space. Given a certain information space, an expert-finding system must then enable the user to find the right expert who can be used as a gateway to, or *in lieu* of, the information space. This justifies the approach of mining the information space as an implicit source of expertise indicators as we saw in the domain analysis.

Even for those users who seek experts because they believe that the information is not documented, an expert finder that mines the information space can serve useful purposes. To begin with, it can complement information retrieval systems in ascertaining the belief that the information is actually not documented. Besides, it can be used to trace those who work in the problem domain at hand. But to help the user find an expert on something not documented, the expert finder must use other means for determining expertise beyond the organizational information resources.

The fact that users need experts as information sources also substantiates the claim that expert finding systems need to be deployed in complement to, and as part of, knowledge management, information retrieval, CSCW, collaborative filtering and related services. Furthermore, since expert-seeking arises during day-to-day work, it is beneficial if expert finding systems are embedded in the day-to-day problem solving and information search environments.

Users seeking for experts who can perform some organizational function need more than just information. These users mainly focus on evaluating and comparing the expertise of people. Viewed in the light of McDonald and Ackerman's [5] dichotomy of expert finding into two phases as *expertise identification* and *expertise selection*, these users need more support for the later phase than those who are seeking experts as information sources. However, in the works we analyzed, only expertise identification is targeted.

Specifically, while mapping the domain model to the expert finding problem, we have identified four gaps, namely (1) a source heterogeneity gap, (2) a technique optimality gap, (3) an analysis support gap, and (4) a reusability and interoperability gap. In the following, we will describe each of these shortcomings and suggest some ways to tackle them. This discussion also forms the foundations for the approach we intend to present in the next section of the paper.

5.2.1 Source heterogeneity gap

In order to adequately exploit the information space as a source of expertise indicators, expert finders need to handle the heterogeneity and distributedness of the information space. This is also reflected by the wide variety of expertise indicator sources observed in the domain analysis.

However, so far each of the current systems exploits a few types of mostly centralized sources only. Hence, architectures that are flexible enough to address this problem are needed.

A number of dimensions of heterogeneity and distributedness can be identified. One aspect of heterogeneity that needs to be addressed is the heterogeneity of sources in reflecting expertise. How well expertise indicators like terms and phrases reflect expertise is mainly a factor of how the *source* in which these indicators occur relates to the expert, irrespective of the indicators' direct statistical (or otherwise) correlation with the expert. For example, the occurrence of the term "software agents" in Dr. X's resume and its occurrence in one of her publications may not weigh the same. Moreover, the occurrence of this term in the title of a document shows different distance to her actual expertise, compared to its occurrence anywhere in the body. Therefore, expert finders need to determine how a document relates to an expert before going on to extract expertise indicators from it.

Another important aspect of heterogeneity of sources that is important in expertise recognition arises from their differing provenance. The sources include a range of organizational resources, personal information resources of experts as well as the experts themselves. This calls for functionality that specializes in monitoring each of these sources as well as mechanisms to calibrate their differing ability of indicating expertise.

There also remain a lot to be done in addressing the common situation where expertise indicators are physically distributed across the organization and stored in various formats (databases, document repositories, web sites and the like). Furthermore, exploiting personal information resources brings in a range of privacy issues. Unlike the traditional expert database approach where experts *actively* supply all information that will be accessible to the outside for evaluation, here *all* available information about the expert will be evaluated by default, unless the expert is allowed to actively *block* certain pieces of information from being scanned and evaluated. This calls for the need to enable experts to guide and control the system in monitoring their personal sources.

5.2.2. The expertise recognition methodology gap

Being a relatively nascent research area, the work in the expert finding domain until now focused more on ascertaining the feasibility of a solution to the problem than seeking optimal approaches and techniques. Consequently, the techniques employed and their underlying assumptions also fall short of realizing systems that properly address the expert finding problem.

Admittedly, expertise identification and representation poses immense challenge as the pool of expertise indicators is large and amorphous, the descriptors of expertise are difficult to extract, expert qualities are multidimensional and users' expertise needs are mostly fine-grained. The methods adopted by the existing systems to deal with these challenges are mainly based on generalizing expert finding with information retrieval assumptions and theories. Although this sort of 'reductionist' approach helps to a certain extent, we believe there is a need to explore the unique features and requirements of expertise identification and modeling, as well as expertise query formulation.

5.2.3. Expertise Analysis support gap

As already noted, current expert finders scarcely address expertise selection. Support for this can be based on the information space as well as on other sources. From the information space, an

expert finder can provide access to the output of the expert as demonstrated by documents produced, and determine the expert's influence as demonstrated by citations received, co-authorship and other relationships with other experts, etc.

While setting out to support expertise selection, one should not forget that it is the user who does the selection and that the system can only support this process by providing appropriate analysis functionality. For example, the following system capabilities may considerably enhance the selection of the right expert:

- The ability to rank and contrast experts using different user-selected criteria (rather than the mere provision of a linear listing based on a pre-determined criterion);
- Example-based search for experts using one expert's model as a query (or as a frame of reference);
- The presentation/explication of multiple relationships among experts. This can be used for visualizing (and catalyzing) collaboration, co-authorship, citation links, project groups, etc. (Note that two experts that are otherwise related may be far apart with respect to a query and, conversely, two experts may be close with respect to a query but have little relationship to each other);
- Using an ontology-based (i.e., knowledge-based) presentation/ranking of experts (subset-superset relationships, multidisciplinary and interdisciplinary expertise, etc.)

One way to help users in expertise selection is to increase the system's transparency. Viewed within the framework of our domain model, this can be done by providing interfaces to access the expertise indicator sources, the expertise model and the expert model, as well as the expertise recognition logic. This, in turn, permits the incorporation of functionality that can assist the user in evaluating and exploring (visualizing) the expertise information.

5.2.4. Reusability, interoperability and extensibility gap

Probably due to the embryonic stage of research on expert finders, current systems focus on solving a particular problem and coming up with a standalone solution (what Cohen and Northrop [45] referred to as the "single-system mentality"). However, in order to live up to the wider variety of applications where expert finders are needed, these systems should be both readily transferable from application to application and be interoperable with other systems. Moreover, since different applications and organizational constraints may require tailored implementations of the same facet, their components need to be substitutable and extensible as appropriate. Hence, as much as possible, expert finders should be designed in a generic architecture, keeping modularity and reusability in mind.

To this end, partitioning the whole system into loosely coupled modules that implement each facet appears to be promising. We believe the domain model presented in this paper will be useful in this regard.

6. THE DEMOIR APPROACH

6.1. Which Type of Expertise Model to Use ?

In the previous section we argued that there is a need for an expert finding system that can exploit heterogeneous sources of expertise indicators, that uses better techniques of expertise modeling and identification, is flexible enough to be integrated with other systems and can support analysis as well as other exploitations of expertise information. A closer look at the seven facets discussed in section 5.1 reveals that expertise models form the critical architectural foundation that determine performance and usability. Let us, therefore, briefly examine the various alternatives for implementing this facet with regard to their ability to form a foundation for a solution that attacks the aforementioned gaps.

As we indicated before, we have three alternative expertise models: query-time generated, personal-agent based and aggregated expertise models. While the query-time generated approach uses dynamically generated expertise models, the other two are *a priori* constructed. The personal agent-based models suggest a decentralized architecture while the aggregated expertise models suggest a centralized approach.

Mattox *et al* [23] argue that query-time generated expertise models have the advantage of avoiding the maintenance of (possibly outdated) information internally and would permit operation in real-time using the most recent information available to locate experts. However, not surprisingly, the authors reported that their system suffers from high latency in query processing. We also witnessed the same shortcoming from a comparable expert finding system we developed for a research department. We feel that, given the availability of search robots that can routinely gather sources of expertise indicators and can construct and update expert profiles, the advantage of query-time generated models is not important any more. Besides, this approach allows for limited manipulations on the expertise information (like analysis, visualization, browsing and comparison of expert profiles). The reason is that the approach lacks persistent expertise data, and cannot exploit all types of documents (e.g., emails and other personal sources are difficult to include). It is also difficult to exploit diverse sources of expertise data other than documents (e.g. recommendations from people, social relations, feedback, etc. are difficult to incorporate).

The tradeoffs between the other two approaches (i.e. aggregated expertise model and distributed personalized agent-based model) are more difficult to weigh. The advantages of using distributed personal agents include easy privacy maintenance through locality of expertise modeling, which in turn may give experts the feeling of being in control (although any access restrictions that may be imposed on an agent-based model can also be realized for aggregated models as well). However, this approach suffers from the limitations of relying on personal sources of information only. The fact that the expertise data is distributed results in limited accessibility and sub-optimal utilization. The approach also suffers from problems of scalability when the number of experts grows large, since having individual expertise-modeling agents for numerous people overloads the network during querying because every query is matched to all distributed personal agents.

Maintaining an aggregated expertise model overcomes the aforementioned shortcomings of the other two approaches to a great extent while introducing its own shortcomings. This approach allows for an open and multi-purpose exploitation of the expertise information, manipulation of the expertise information in the aggregate, and monitoring a wide range of organizational sources

for up-to-date expertise data. It also facilitates using both knowledge-based and statistical expertise modeling (e.g., allows the analysis and grouping of experts based on their expertise profiles, facilitating inferences on the relationships among experts, etc). However, it fails to afford localization, access to personal information sources of individual experts as well as privacy of the expertise data. A more detailed discussion of this point can be found in [46,47].

An optimal solution is one that can integrate the aggregated expertise models with personal-agent based approach. The DEMOIR Architecture (Dynamic Expertise Modeling from Organizational Information Resources), which we will briefly discuss in the following section, is motivated by the desire to implement such a hybrid and flexible solution. This approach is our attempt to lay the foundations for the development and testing of better algorithms and heuristics for expertise modeling and identification based on implicit sources of expertise evidence.

6.2. The DEMOIR Architecture

The DEMOIR Architecture, shown in Figure 1, is a modular architecture for expert finding systems that is based on centralized expertise models while also incorporating decentralized expertise indicator source gathering, expertise indicator extraction, and distributed clients. It manages to do this by dissociating functions like source gathering, expertise indicator extraction and expertise modeling and delegates them to specialized components which can be separately implemented and readily combined to suit an application environment. Specifically, this architecture has the following characteristics:

1. Configurability and extensibility at all of its modules, thereby effectively addressing the reusability and interoperability gap.
2. Expertise modeling which considers the heterogeneity of sources as valuable meta-information that must be exploited, as opposed to a hodgepodge that must be made uniform. For instance, DEMOIR tries to explicitly capture how the sources relate to experts and where the expertise evidence stems from, and factor this information into the expertise modeling process.
3. A combination of distributed and centralized monitoring of expertise data from both organizational and personal information sources. As such, it allows distributed and expert-controllable expertise indicator gathering components (like robots) which can be tailored to the specific requirements of both organizational and personal sources. This enables us to tackle some aspects of the source heterogeneity gap and privacy problems.
4. A centralized expertise information server that aggregates, stores and supplies expertise information and provides application programming interfaces (APIs) for client implementations and integration with other systems.
5. DEMOIR provides structure to integrate domain knowledge at all steps along with statistical and heuristics-based methods internal to components.
6. Delegation to clients of application-specific exploitation of the expertise information (e.g. searching, browsing, analysis, visualization, etc.).

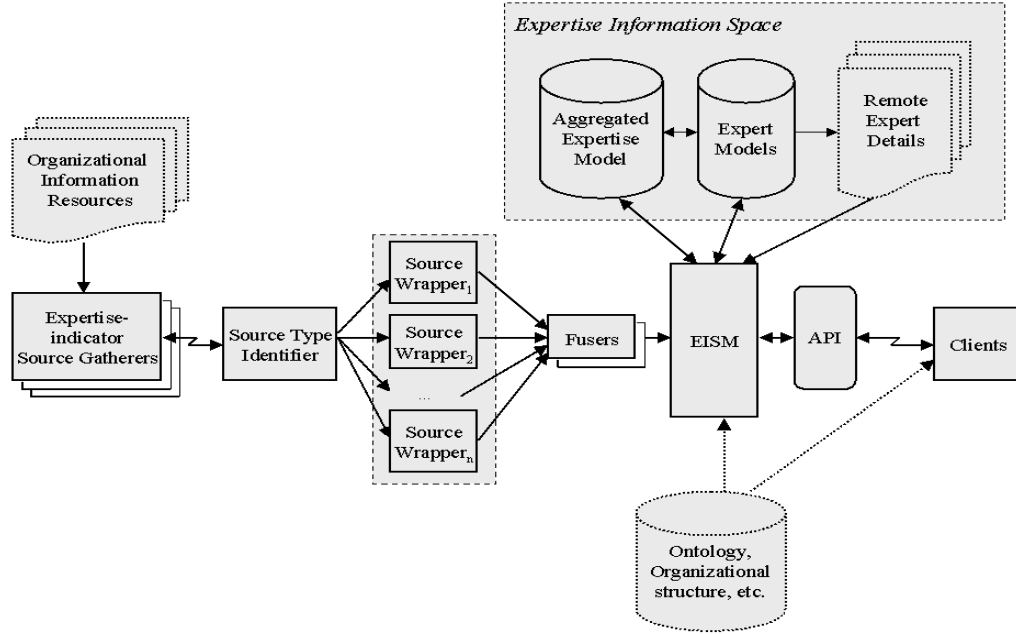


Figure 1. The DEMOIR Architecture

The above-mentioned characteristics are realized by the following DEMOIR components (also see Figure 1):

Expertise Indicator Source Gatherers. These are agents that serve as ‘informers’ to the centralized expertise modeling server by routinely gathering expertise indicator sources from various parts of the organizational information resource like web sites, databases, repositories, personal sources, etc. They are envisaged to be independent, adaptable to local constraints of the expertise indicator source they exploit, and are not to be controlled by any central system. They are equipped with heuristics to recognize potential sources of expertise data (like memorizing names of experts and resolving name variations). They can also use techniques to trace name occurrences without *a priori* knowledge of expert names using, for example, proper noun identification methods employed in natural language processing (e.g. [48]).

In addition to collecting implicit indicators of expertise using automated gatherers, an expert finding system should also be able to collect explicit expertise data from people, when available. To this end, DEMOIR aims to provide gatherers that enable people to supply expertise information and recommend experts to the system. DEMOIR also provides interfaces that allow users to monitor and edit automatically gathered expertise information. This feature is particularly useful in situations when it is critical to maintain highly accurate expertise data.

Source Type Identifier. This is a module that analyzes clues in the content and structure of sources to determine how they relate to the expert(s) whose names are contained in or linked to these sources. The output of the source type identifier is used to invoke the appropriate source wrappers (see below). Hence, the source-type identifier, along with the source wrappers, implements differentiated extraction of expertise indicators based on types of source-to-expert relations.

Source Wrappers and Fusers. Wrappers are sets of modules that handle the extraction of

expertise indicators from heterogeneous sources. Employing both statistical techniques and heuristics, each source wrapper extracts expertise indicators from a particular source type. The expertise indicators thus extracted are merged into aggregated expertise models by the fusers. The expertise models thus constructed, along with expert profiles and other related information (like the local archive of sources, expert networks, etc.) constitute what we call an *expertise information space* maintained and managed by the expertise information space manager.

Expertise Information Space Manager (EISM). This component manages the storage of, and retrieval from, the expertise information space. It specifically handles the low-level representation and maintenance of, as well as access to, the aggregate expertise model, the local expert models and the remote expert details, which are the constituents of the expertise information space. It also executes the requests from the APIs. Moreover, the EISM enables a flexible integration of domain models (ontologies) or modules to generate them using methods such as concept clustering.

As we distinguished in section 5.1, the aggregated expertise model represents the “expertise space” in a certain domain (like an organization) while expert models represent the individual experts including a link to the aggregated expertise model and other peculiar information like personal attributes and links to other experts. In other words, the expertise model represents the conceptual relation and proximity among experts while the expert models represent other relations (like social ties, organizational relations, co-authorship relations, etc.). In terms of McDonald and Ackerman’s [5] dichotomy, we can also view expertise models as supporting expertise identification (“who are the experts on topic x?”) while expert models/profiles support expertise selection (“what does expert z know? how extensive is his/her knowledge? which other experts relate to him/her?”).

The Application Programming Interface (API). This facilitates the implementation of various clients to enable search, browsing and analyses. The API also supports the integration of the DEMOIR server in other applications.

As we discussed in section 4, expert finding systems can be deployed in an organization in a variety of ways. One such application, for example, is allowing people to conduct a search for expertise based on key words or descriptions of problems. Such a system can then be made available on an organization’s Intranet, for example. The DEMOIR API can easily support the development of such applications.

7. CONCLUSION AND FUTURE RESEARCH

In this paper, we analyzed the expert finding problem by grouping what motivates people to look for experts into two groups, namely information needs and expertise needs. Then, we took a closer look at the expert finding systems domain using the approach of domain analysis and presented our findings as a domain model. By overlaying the findings of these two analyses, we discussed the gaps that remain to be addressed and argued that it is more fruitful to approach them with a relatively generic architecture. In line with this, we describe the DEMOIR approach which we employ as a modular framework for expert finders, and demonstrate how it tackles the observed gaps to a considerable extent.

As we suggested throughout this paper, a number of issues still remain open. Among the

questions that need further investigation are: How can expertise indicator sources in organizations be identified and exploited? How should expertise and expert models be structured and represented? How should one apply inference rules and algorithms on expert and expertise relationships? And how should users be supported in searching, visualizing and analyzing the expertise information space?

An investigation of the contextual and integration issues is also important for the proper positioning and deployment of expert finding systems within organizations. At the level of the individual user, adaptive features in the sense of identifying and fulfilling expertise needs from users' daily activities are necessary (see [49] for available methods). At the organizational level, the proper organization-wide applications, services and systems into which expert finders can be incorporated also need to be identified. And finally, issues of privacy and security also must be taken into account (see [50] for a reference architecture in the area of user modeling that can also be adopted in the area of expert finding).

We aim at investigating some of these issues by pursuing the development and testing of our DEMOIR Server architecture employing various techniques. Specifically, we intend to continue working on: (1) the development and testing of heuristics for source identifiers and wrappers, (2) methods of versatile representation of the expertise information space to support a range of purposes including exploitation (search and inference), visualization and analysis, and (3) evaluation of our approaches and techniques through implementation and testing with users.

ACKNOWLEDGEMENT

Much of the research described here has been carried out while the authors were affiliated with GMD – German National Research Center for Information Technology. A great deal of the ideas presented in this paper benefited from the discussions with the information brokering team members Jürgen Koenemann, Christoph Thomas, Roland Klemke and Achim Nick as well as Mark Ackerman, Wolfgang Pohl and Reinhard Oppermann.

References

- [1] T. Stewart, "Leading Lights: Author Tom Stewart on Intellectual Capital" (Interview), *Knowledge Inc.*, May 1997.
- [2] M. Hertzum, A. M. Pejtersen, "The Information-Seeking Practices of Engineers: Searching for Documents as well as for People", *Information Processing & Management*, vol. 36, no. 5, pp. 761-778, 2000.
- [3] T. Wilson, "Information-Seeking Behavior: Designing Information Systems to Meet our Client's Needs", in *ACURIL (Association of Caribbean University, Research and Institutional Libraries) XXV Conference*, San Juan, Puerto Rico, 1995.
- [4] D. Mattox, "Expert Finder", *The Edge: The MITRE Advanced Technology Newsletter*, vol. 2, no. 1, 1998. Available: http://www.mitre.org/pubs/edge/June_98.
- [5] D. W. McDonald and M. S. Ackerman, "Just Talk to Me: A Field Study of Expertise Location", *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW'98)*, Seattle, WA: ACM Press, 1998, pp. 315-324.
- [6] A. Kanfer, J. Sweet and A. Schlosser, "Humanizing the Net: Social Navigation with a 'Know-Who' Email Agent", in *The Third Conference on Human Factors & The Web*, Denver, CO, June 12, 1997. [Online]. Available: <http://www.optavia.com/hfweb/history.htm>.
- [7] H. Kautz, B. Selman and A. Milewski, "Agent Amplified Communication", in *The Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, Portland, OR, 1996, pp. 3 – 9.
- [8] A. S. Vivacqua, "Agents for Expertise Location", in *Proceedings of the AAAI Spring Symposium on Intelligent Agents in Cyberspace*, Stanford, CA, March 1999, pp. 9 – 13.
- [9] A. L. Cohen, P. P. Maglio and R. Barrett, "The Expertise Browser: How to Leverage Distributed Organizational Knowledge", presented at: *The Workshop on Collaborative and Cooperative Information Seeking in Digital Information Environments at CSCW'98*; Seattle, WA, 1998.
- [10] D. E. O'Leary, "Knowledge-Management Systems: Converting and Connecting", *IEEE Intelligent Systems & Their Applications*. May/June, pp. 30-33, 1998.
- [11] P. H. J. Hendriks and D. J. Vriens, "Knowledge-based Systems and Knowledge Management: Friends or Foes?", *Information & Management*, vol. 35, no. 2: pp. 113-125, 1999.
- [12] I. Grief, "Everyone is Talking about Knowledge Management" (Panel), *Proceeding of the ACM 1998 Conference on Computer Supported Cooperative Work, CSCW'98*, Seattle, Washington, 1998, pp. 405 – 406.
- [13] C. C. Kuhlthau, *Seeking Meaning: A Process Approach to Library and Information Services*, Norwood, NJ: Ablex Publishing, 1993.
- [14] C. C. Kuhlthau, "A Principle of Uncertainty for Information Seeking", *Journal of Documentation*, vol. 49, no. 4, pp. 339 – 355, 1993.
- [15] T. Lewis, "Where the Smart Money Is?", *IEEE Computer*, vol. 32, no. 11, p. 136, 1999.
- [16] T. H. Davenport and L. Prusak, *Working Knowledge : How Organizations Manage What They Know*, Boston, MA: Harvard Business School Press, 1998.
- [17] M. E. Maron, S. Curry and P. Thompson, "An Inductive Search System: Theory, Design and Implementation", *IEEE Transaction on Systems, Man and Cybernetics*, vol. SMC-16, No. 1, pp. 21-28, January/February 1986.
- [18] L. A. Steeter and K. E. Lochbaum, "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, pp. 345 – 349.
- [19] L. A. Steeter and K. E. Lochbaum, "Who Knows: A System Based on Automatic Representation of Semantic Structure", in *RIAO'88*, Cambridge, MA, 1988, pp. 380 – 388.
- [20] G. W. Furnas, S. T. Dumais, T. K. Landauer, R. A. Harshman, L. Y. Streeter and K. E. Lochbaum. "Information Retrieval using Singular Value Decomposition Model of Latent Semantic Structure", in *Proceedings of the 1988 ACM SIGIR Conference on Research and Development in Information Retrieval*, Grenoble, France, 1988, pp. 465 – 480.
- [21] B. Krulwich and C. Burkey, "ContactFinder: Extracting Indications of Expertise and Answering Questions with Referrals", in *Working Notes of the 1995 Fall Symposium on Intelligent Knowledge Navigation and Retrieval*, Cambridge, MA. Technical Report FS-95-03, The AAAI Press, 1995, pp. 85 – 91.

- [22] B. Krulwich and C. Burkey, "The ContactFinder Agent: Answering Bulletin Board Questions with Referrals", in *Proceedings of the 1996 National Conference on Artificial Intelligence (AAAI-96)*, Portland, OR, vol. 1, 1996, pp. 10 – 15.
- [23] D. Mattox, M. Maybury and D. Morey, "Enterprise Expert and Knowledge Discovery", in *Proceedings of the 8th International Conference on Human-Computer Interaction (HCI International'99)*, Munich, Germany, August 1999, pp. 303-307.
- [24] H. Kautz and B. Selman, "Creating Models of Real-World Communities with ReferralWeb", in *Working Notes of the Workshop on Recommender Systems held in conjunction with AAAI-98*, Madison, WI, 1998, pp. 58 – 59.
- [25] H. Kautz, B. Selman and M. Shah, "The Hidden Web", *The AI Magazine*, vol. 18, no. 2, pp. 27 – 36, 1997.
- [26] M. F. Schwartz and D. M. Wood, "Discovering Shared Interests Using Graph Analysis", *Communications of the ACM*, vol. 36, no. 8, pp. 78 – 89, 1993.
- [27] L. N. Foner, "Yenta: A Multi-Agent Referral-Based Matchmaking System", in *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, Marina del Rey, CA, February'97, 1997, pp. 301 – 307.
- [28] D. W. McDonald and M. S. Ackerman, "Expertise Recommender: a Flexible Recommendation System and Architecture", in *Proceeding of the ACM 2000 Conference on Computer Supported Cooperative Work (CSCW'00)*, Philadelphia, PA, 2000, pp. 231-240.
- [29] M. S. Ackerman, D. McDonald, W. Lutters and J. Muramatsu, "Recommenders for Expertise Management", presented at: *Workshop on Recommender Systems, Algorithms and Evaluation at the ACM SIGIR '99*, Berkeley, CA, 1999.
- [30] M. S. Ackerman and D. W. McDonald, "Answer Garden 2: Merging Organizational Memory with Collaborative Help", in *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'96)*, Boston, MA, 1996, pp. 97 – 105.
- [31] V. A. Kulyukin, K. J. Hammond and R. D. Burke, "Answering Questions for an Organization Online", in *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, Madison, WI, AAAI Press, 1998, pp. 532-538.
- [32] J. E. Greer, G. I. McCalla, J. A. Collins, V. S. Kumar, P. Meagher and J. I. Vassileva, "Supporting Peer Help and Collaboration in Distributed Workplace Environments", *International Journal of Artificial Intelligence in Education*, vol. 9, 1998, pp. 159-177, 1998.
- [33] J. E. Greer, G. I. McCalla, J. E. Cooke, J. Collins, V. S. Kumar, A. Bishop and J. I. Vassileva, "The Intelligent Helpdesk: Supporting Peer-Help in a University Course", in *Proc. of the International Intelligent Tutoring Systems Conference (ITS'98)*, San Antonio, TX, 1998, pp. 494-505.
- [34] A. Pikarakis, et al, "MEMOIR: Software Agents for Finding Similar Users by Trails", in *Proceedings of the Third International Conference on the Practical Applications of Intelligent Agents and multi-Agent Technology (PAAM-98)*, London, UK, 1998, pp. 453 – 466.
- [35] A. L. Cohen, P. P. Maglio, R. Barrett, "The Expertise Browser: How to Leverage Distributed Organizational Knowledge", presented at: *Workshop on Collaborative Information Seeking at CSCW'98*, Seattle, WA, 1998.
- [36] S. O. Kimbrough and J. R. Oliver, "On Relevance and Two Aspects of the Organizational Memory Problem", in *Proceedings of the Fourth Annual Workshop on Information Technology and Systems (WITS'94)*, Vancouver, BC, Canada, 1994, pp. 302-311.
- [37] T. H. Davenport, *Some Principles of Knowledge Management*, April 1996. [Online]. Available: <http://www.bus.utexas.edu/kman/kmprin.htm>.
- [38] K. Höök, Å. Rudström and A. Waern, "Edited Adaptive Hypermedia: Combining Human and Machine Intelligence to Achieve Filtered Information", in *Flexible Hypertext Workshop held in conjunction with the 8th ACM International Hypertext Conference (HYPERTEXT'97)*, Southampton, UK, 1997. Available: <http://www.sics.se/~kia/papers/edinfo.html>.
- [39] F. Hattori, T. Ohguro, M. Yokoo, S. Matsubara, and S. Yoshida, "Socialware: Multiagent Systems for Supporting Network Communities", *Communications of the ACM*, vol. 42, no. 3, pp. 55 – 61, 1999.
- [40] A. Lang and Y. Pigneur, "An Electronic Market of Individual Human Competencies for Team Building", *Virtual-organization.net Newsletter*, vol. 1, No. 3., pp. 4 – 13, 1997. [Online]. Available: <http://www.virtual-organization.net/files/articles/nl1-3.pdf>.
- [41] A. Moukas, G. Zacharia, R. Guttman, and P. Maes, "Agent-Mediated Electronic Commerce: An MIT Media Laboratory Perspective", *International Journal of Electronic Commerce*, vol. 4, no. 3, pp. 5 – 22, 2000.
- [42] C. Lung and J. E. Urban, "Integration of Domain Analysis and Analogical Approach for Software Reuse", in *Proceedings of the 1993 ACM/SIGAPP Symposium on Applied Computing*, New York, NY, 1993, pp. 48 – 53.

- [43] R. Prieto-Diaz, "Domain Analysis for Reusability", in *Proceedings of COMPSAC'87: The Eleventh Annual International Computer Software & Applications Conference*, Tokyo, Japan, 1987, pp. 23 – 29.
- [44] A. Birk, "Modeling the Application Domains of Software Engineering Technologies", in *Proceedings of the 12th IEEE International Conference on Automated Software Engineering*, Incline Village, NV, 1997, pp. 291-292.
- [45] S. Cohen and L. M. Northrop, "Object-Oriented Technology and Domain Analysis", in *Proceedings of the Fifth International Conference on Software Reuse*, June 2-5, 1998, Victoria, BC, Canada, 1998, pp. 86-93.
- [46] D. Yimam Seid and A. Kobsa, "Centralization vs. Decentralization Issues in Internet-based Knowledge Management Systems: Experiences from Expert Recommender Systems", presented at: *Workshop on Organizational and Technical Issues in the Tension Between Centralized and Decentralized Applications on the Internet (TWIST2000)*. Institute for Software Research, University of California, Irvine, CA, 2000. [Online]. Available: <http://www.ics.uci.edu/~kobsa/papers/2000-TWIST-kobsa.pdf>.
- [47] D. Yimam Seid and A. Kobsa, "DEMOIR: A Hybrid Architecture for Expertise Modeling and Recommender Systems", in *Proceedings of the IEEE 9th International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises*, Gaithersburg, MD, 2000, pp. 67-74. Available: <http://www.ics.uci.edu/~kobsa/papers/2000-IEEE-kobsa.pdf>.
- [48] N. Wacholder, Y. Ravin and M. Choi, "Disambiguation of Proper Names in Text", in *Proceedings of the Fifth Conference on Applied Natural Language Processing*, Washington, DC, 1997, pp. 202 – 208.
- [49] A. Kobsa, J. Koenemann and W. Pohl, "Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships", *The Knowledge Engineering Review* (forthcoming), Available: <http://www.ics.uci.edu/~kobsa/papers/2001-KER-kobsa.pdf>.
- [50] J. Schreck, "Security and Privacy in User Models". Ph.D. Thesis, Dept. of Mathematics and Computer Science, University of Essen, Germany, 2001.