Fazel-Zarandi, M.; Fox, M.S.; Yu, E. (2012). "Ontologies in Expertise Finding Systems: Modeling, Analysis, and Design." Ontology-Based Applications for Enterprise Systems and Knowledge Management, IGI Publishing.

Ontologies in Expertise Finding Systems: Modeling, Analysis, and Design

Maryam Fazel-Zarandi
Department of Computer Science, University of Toronto, Canada

Mark S. Fox
Department of Mechanical and Industrial Engineering, University of Toronto, Canada

Eric Yu
Faculty of Information, University of Toronto, Canada

ABSTRACT

Knowledge Management Systems that enhance and facilitate the process of finding the right expert in an organization have gained much attention in recent years. This chapter explores the potential benefits and challenges of using ontologies for improving existing systems. A modeling technique from requirements engineering is used to evaluate the proposed system and analyze the impact it would have on the goals of the stakeholders. Based on the analysis, an ontology-based expertise finding system is proposed. This chapter also discusses the organizational settings required for the successful deployment of the system in practice.

INTRODUCTION

Expert profiling and identification are important to both organizations and knowledge workers. In today's competitive business environment, companies need to understand the skills and competency of their human resources in order to best utilize them. This is particularly important for organizations that engage with multiple and changing clients such as consulting firms and software development companies since these organizations need to be able to flexibly respond to internal and external demands for skills and competencies. From a knowledge worker's perspective, finding individuals with appropriate skills and knowing who to go to are important for accomplishing knowledge intensive tasks and solving complex problems. For these purposes, people often rely on their past experiences, existing documents, and others who have the needed expertise.

Two main motives for seeking an expert are 1) finding a source of information and 2) finding someone who can perform a given task (Yimam-Seid and Kobsa, 2003). Finding an expert for these purposes, however, might not be an easy task for many reasons. Expertise is highly dynamic (Maybury, 2006), difficult to quantify (Earl, 2001), and varying in level (Earl, 2001). When expert involvement in a given activity is required, it is also necessary to know if the expert is actually competent enough to perform the task in addition to being knowledgeable in the field. It is also difficult to validate other people's expertise (Maybury, 2006) and to distinguish a good expert from a bad one. Furthermore, due to the complexity of some problems, the assistance of multiple experts may be required (Earl, 2001). The difficulty of locating experts increases in larger and more geographically distributed organizations and communities.

In order to augment and assist the process of locating expertise within an organization, the study and development of special Knowledge Management Systems that suggest people who have some expertise in a given area, has received the attention of both researchers and organizations. The resulting systems either rely on individuals to provide accurate and comprehensive profiles of their competences and experiences (Earl, 2001), or use mechanisms to automatically discover up-to-date expertise information from secondary sources such as articles, email communications, and forums (Stankovic, Wagner, Jovanovic, and Laublet, 2010; Yimam-Seid and Kobsa, 2003). In a review of these systems in (Yimam-Seid and Kobsa, 2003), however, problems related to heterogeneous information sources, expertise analysis support, and interoperability were identified.

The common solution to the problems related to heterogeneous information sources and interoperability is to formally specify the meaning of the terminology of each system and to define a mapping between these terminologies. In other words, use ontologies to provide a shared common understanding of the structure of information among systems and software agents. In addition, because of their powerful knowledge representation formalism and associated inference mechanisms (Razmerita, Angehrn, and Maedche, 2003), ontologies can also be incorporated to address problems related to expertise analysis support.

Taking these facts into account, it would seem natural to expect that Expertise Finding Systems (EFS) can benefit from the use of ontologies. However, there are various potential difficulties and challenges associated with the use of ontologies that may cause the system to fail. In this chapter, we are interested in investigating the circumstances under which an ontology-based EFS might or might not work. More specifically, we want to systematically explore and analyze how ontologies might be used in EFS, before creating a prototype and conducting case-studies. For this, we use a modeling technique from requirements engineering to evaluate the proposed ontology-based EFS and analyze the impact that it would have on the goals of the stakeholders. This chapter extends our previous work in (Fazel-Zarandi and Yu, 2008) with a more extensive literature review, a description of our proposed ontology-based EFS based on the analysis, and a more detailed discussion of areas of future work.

The organization of this chapter is as follows: A brief presentation of related work on Expert Finding Systems and the role that ontologies can play in such systems is followed by knowledge management analysis and a discussion of knowledge processes, knowledge markets between individuals, and role of technology supported by goal models developed using the i* notation. The required organizational settings for a successful deployment of the technology are also elaborated on in this section. Our proposed ontology-based EFS is then presented. Finally, the chapter concludes with a discussion of contributions made and areas of future work.

BACKGROUND

Expertise Finding Systems are a subset of recommender systems where experts are the "items" being recommended (Hansen, Khopkar, and Zhang, 2010). Comprehensive surveys of this literature are available elsewhere (Adomavicius and Tuzhilin, 2005; Burke, 2005; Schafer, Frankowski, Herlocker, and Sen, 2007). We focus instead on aspects of the literature that motivate and inform our model and analysis, and distinguish between non-ontology-based and ontology-based systems.

Both categories of EFS rely on implicitly or explicitly provided data about individuals' skills and competencies to create profiles of experts (Stankovic, et al., 2010; Yimam-Seid and Kobsa, 2003) and recommend individuals based on those profiles. Initially, the evidences considered were content created by the individual within the organization, enrollment in learning activities, and/or experiences related to the workplace. With the growth of the World Wide Web, however, digital media and communication networks have become an important medium for enabling new levels of interactions in organizations and communities. Many online communities and interactive collaboration spaces (such as forums and wikis) evolve into large-scale knowledge networks which are context-dependent and multi-dimensional (Y. Huang, Contractor, and Yao, 2008) and provide additional evidence for expertise identification. In recent years, enterprise social networking such as LinkedIn has also been considered as a different approach to employee profiles for expert finding and community formation. Some of these platforms allow individuals to create profiles of themselves and indicate their connections to other users. Others, such as IBM's Fringe Contacts, allow individuals to describe their colleagues by tagging them with keywords on their expertise and interests, thus, creating a publicly visible tag cloud which characterizes the individual employee (Braun and Schmidt, 2008). In addition, the growing number of data published on the Web according to Linked Data principals and using unambiguous vocabularies (Stankovic, et al., 2010) make automated reasoning possible for improving expertise identification. In short, expertise can be declared by individuals about themselves or by others, and/or can be derived from 1) activities performed by the individual either online or offline including enrolment in learning activities, experiences related to the workplace, generating content both within the organization and on the Web, and question-answering in online forums; 2) recommendations and "wisdom of the crowd"; and 3) assessments in the form of various tests or evaluation in the form of 360 reviews or performance appraisal (Figure 1).

Insert Figure 1 Here

Figure 1. Sources of expertise information

Non-ontology-based approaches usually use databases as skill repositories in which user profiles are expressed by data structures or vectors of terms (Colucci, Di Noia, Di Sciascio, Donini, and Ragone, 2007). To evaluate possible matches in non-ontology-based systems, information retrieval techniques such as database querying and similarity between weighted vectors of terms have been used (Veit, Müller, Schneider, and Fiehn, 2001). Another approach used in these systems is to model skill matching as a bipartite graph in which the first set of vertices includes assignees and the second one includes tasks to be performed, and edges link people to tasks (Saip and Lucchesi, 1993 as cited in Colucci, et al., 2007). By assigning a cost to each edge a weighted bipartite graph is achieved and techniques for solving the

Assignment Problem can be used to solve it. The main limitations of non-ontology-based systems are 1) lack of mechanism for reasoning about individual's skills and competencies, and 2) lack of a common understanding of the terminology used in the organization.

Ontology-based approaches have been proposed to address these limitations. An ontology is a formal description of a set of objects, concepts, and other entities that are assumed to exist in a domain of interest along with their properties and relationships that hold among them (Gruber, 1993) and the constraints that exist over them. Ontologies provide a shared common understanding of information among people or software agents (Fox, Barbuceanu, and Gruninger, 1996) and enable the reuse of domain knowledge (Guarino, 1998). They can be used to capture and represent data semantics and by assuming deductive capability as provided by an inference engine, ontologies provide the means for deduction and automated reasoning in order to generate further knowledge (Fox, et al., 1996) (i.e., knowledge that is not explicitly known but that can be deduced based on the general knowledge of the domain). In addition, ontologies can also be used for information integration, i.e., the merging of information from different sources despite differing conceptual and contextual representations (Guarino, 1998).

In the context of expertise finding, different types of ontologies may prove to be beneficial for different purposes. A domain (expertise) ontology which would capture and represent the terminology and concepts used in the domain is of central importance and has been considered in previous research. For example, EFS use different evidences as indicators of expertise (Yimam-Seid and Kobsa, 2003), some of which include self declarations of expertise and artifacts created by experts (Maybury, 2006). A domain ontology can be used to automatically annotate existing information resources and to perform automated reasoning to improve expertise indicator detection and extraction mechanisms. Also, since sources can possibly be physically distributed across the organization and stored in different formats, domain ontologies can be used to automate information integration. Furthermore, for self declarations of expertise, domain ontologies can be used to structure the different characteristics of users and their relationships, as well as help the users in specifying their goals and competences. Another useful ontology is the organization ontology (Fox, et al., 1996) which formalizes the organizational structure and can be used to infer expertise based on the roles that the agents play and the communications that occur among them. The Organization Ontology was developed as part of the Toronto Virtual Enterprise (TOVE) project (Fox and Gruninger, 1998) in cooperation with several companies with the goal of providing a basis for enterprise modeling. The knowledge provenance and trust ontologies presented in (J. Huang and Fox, 2006) are other examples of ontologies which can improve expertise finding. These ontologies can be used to formally define the semantics of information sources, information dependencies, relationships between information sources and experts, and trust relationships to improve expertise recognition and extraction.

Previous works on ontology-based expertise finding have primarily focused on building and maintaining skill catalogs in a domain of interest. In the KOWIEN project, for example, practical fundamental research to build up and maintain a detailed ontology-based skill catalog is carried out (Dittmann, 2003). Mochol, et al. (2007) and Gomez-Perez, et al. (2007) develop HR ontologies by integrating existing standards and classifications for supporting the recruitment process. Biesalski and Abecker (2006) discuss the integration of HR processes with ontologies in a project at DaimlerChrysler AG. The results are modeled in a competence catalog that represents knowledge over all areas of production, management, and administration. Schmidt and Kunzmann (2006) and Dorn, et al. (2007) describe ontologies that integrate concepts from skill management and learning, and Aleman-Meza, et al. (2007) describe the integration of existing vocabularies for expertise finding. Techniques for ontology-based skill-profile matching have also been considered. Lau and Sure (2002) propose an ontology-based skill management system for eliciting employee skills and searching for experts within an insurance company. Liu and Dew (2004) present a system which integrates the accuracy of concept search with the flexibility of keyword search to match expertise within academia. Colucci, et al. (2003) propose a semantic based approach to the problem of skills finding in an ontology supported framework. They use description logic inferences to handle the background knowledge and deal with incomplete knowledge while finding the best individual for a given task or project, based on profile descriptions sharing a common ontology.

Ontologies can be used to address the problems related to heterogeneous information sources, expertise analysis support, and interoperability which were identified in (Yimam-Seid and Kobsa, 2003). In order to have a successful system, however, it is essential to take into account the potential difficulties and challenges in using formal ontologies for EFS. Developing a domain ontology that is agreed upon by the members of the organization is often a difficult task (for example, a term may have several widely accepted definitions or none at all). Achieving interoperability between different systems and integrating information from different sources are other highly challenging issues. Incompatibilities may arise due to different vocabularies and differences in the expressiveness power of the ontologies (or simply terminologies) used for different systems. This is usually rooted in the fact that different engineering teams and domain experts are involved in creating different systems. These issues are the focus of many on-going studies in the ontology engineering literature. Once the required ontologies are developed and the desired level of interoperability is achieved, it is important to acknowledge that the ontologies should be incorporated in the daily activities and used for performing workflows. Otherwise, they will decay over time and would not be able to keep up with the dynamic nature of organizations. Specifications for ontologies often need to be changed to reflect changes in the real world; therefore, ontologies have to be maintained and modified frequently, and the maintenance process needs to be viewed as an organizational process (Staab, Studer, Schnurr, and Sure, 2001). A group of knowledge engineers should be responsible for ontology maintenance, and a set of rules for making changes to the ontologies must be present. When making changes to an ontology, ontology versioning must be taken seriously, and the impact of the changes to the overall architecture must be considered. Otherwise, changes may result in incompatibilities between different system components, and also may change the semantics of the data. Having to rely on a group of knowledge engineers for maintenance is by itself another potential difficulty.

KNOWLEDGE MANEMENT ANALYSIS USING AGENT-ORIENTED MODELING

The analysis of the benefits and difficulties of using ontologies for improving EFS raises two important questions: 1) under which conditions can the ontology-based EFS fail? 2) how would failure affect the goals of corresponding stakeholders? One way to go about answering these questions is to construct prototypes and conduct case-studies. However, we are more interested to see if it is possible to find a solution by conducting a systematic analysis of the expertise finding problem. The final goal of such an analysis would be the determination of a set of steps for systematically guiding the design of EFS.

To answer the above questions, we use an agent-oriented modeling approach, in which the EFS can be regarded to represent an "intentional actor" (Strohmaier, Yu, Horkoff, Aranda, and Easterbrook, 2007) – intentional because it pursues assigned goals, and actor because it can exhibit active behavior to a certain extent. The benefits of using this approach are (Strohmaier, et al., 2007): 1) Making intentions of EFS explicit aids in reasoning and arguing about it; 2) Reasoning about the EFS allows for the evaluation of different degrees of goal satisfaction among different actors, thus takes situational context of knowledge transfer into account; and 3) By making relations between different actors' goals and the EFS explicit, the how and why the EFS works or fails can be made visible.

For the construction of agent-oriented models, the i* framework (Yu, 2009; Yu, Giorgini, Maiden, and Mylopoulos, 2011) was chosen in this chapter. This framework allows for clear and simple representation of actors' goals and dependencies among them by means of strategic dependency (SD) and strategic rationale (SR) models. SD models describe the network of dependency relationships among various actors in an organizational context. SR models, on the other hand, contain goals, tasks, resources and softgoals of specific actors that are related to each other through task-decomposition and means-ends links. In other words, SR models provide a more detailed level of modeling compared to SD models, by considering internal, intentional relationships. In i*, actors are represented as agents, roles or positions and the framework has the ability to model common concepts such as goals, softgoals, tasks and resources. In addition, i* also gives the ability to reason about modeled goals by means of goal evaluation algorithms (Horkoff, 2006). Figure 2 shows some of the elements of the i* framework and their corresponding graphical representations. See (Yu, 2009; Yu, et al., 2011) for a more comprehensive background information about the i* framework.

Insert Figure 2 Here

Figure 2. Selected elements of the i framework*

Let us illustrate with a simple example. Figure 3 depicts a simplified SR model of the high-level goals, processes, and intentional dependencies of the roles of expertise seeker and provider. In this model, And, Help, Make, and Some+ are contribution links used to link an element to a softgoal for modeling how the element contributes to the satisfaction or fulfillment of the softgoal[i]. As can be seen in the Figure 3, the expertise provider has the top level goals of Keep the job and get promotion. In order to achieve the former goal s/he needs to Finish his/her own required functions on time. Satisfying management can also have a positive impact on both of the top level goals and to achieve this goal the expertise provider needs to Help expertise seekers in addition to Finish his/her own required functions. The expertise seeker, on the other hand, depends on the expertise provider to provide him/her with the required information or knowledge and/or teach him/her the required skills needed for completing a project. Note that the means-ends link is used to indicate that the goal of Help expertise seeker can be achieved by Provide information or Teach expertise seeker.

Insert Figure 3 Here

Figure 3. Simplified model of intentional relationships between expertise seeker and provider

In the following subsections, we first analyze the expertise finding problem in a general setting using the i* modeling notation. Then, we add the EFS to the model and analyze the impact it would have on the goals of the stakeholders. Finally, we look at the development and maintenance of the ontologies for EFS and analyze the interactions of ontology engineers with other stakeholders.

Expertise Finding without EFS

To model and analyze the expertise finding problem we make use of existing studies on information and expertise seeking. McDonald and Ackerman (1998) report that participants in a medium-sized software firm use complex, iterative behaviors to minimize the number of possible expertise sources, while at the same time, provide a high possibility of garnering the necessary expertise. They distinguished two steps in finding expertise within organizations: 1) Expertise identification: the problem of knowing what information or special skills other individuals have, and 2) Expertise selection: the problem of appropriately choosing among people with the required expertise. Other studies also support this distinction (Borgatti and Cross, 2003; Casciaro and Lobo, 2005). These studies reveal that in addition to expertise seeker's awareness of a potential source's expertise, other factors such as timely access to the source, a degree of safety in the relationship, and willingness of an expertise provider to cognitively engage in problem solving, all play an important part in determining whom the expertise seeker chooses to go to. Other criteria for not selecting an expert include cultural differences, language problems, or a lack of experience in a related but necessary discipline (McDonald and Ackerman, 1998).
Figure 4 shows a SR model of the high-level goals, processes, and intentional dependencies of the roles of

Insert Figure 4 Here

Figure 4. Intentional relationships between expertise seeker and provider using the i notation. For simplicity dependums are omitted.*

expertise seeker and provider and existing IT systems. For simplicity, other goals and tasks of these roles which are not related to expertise seeking and providing are omitted from the figure. For example, the Finish a project on time goal of the expertise seeker may involve finding information from explicit sources, but since this is not related to expertise finding, it is not shown in the figure. As can be seen in the model, the Obtain required information or knowledge from others goal of the expertise seeker is decomposed into two tasks: Identifying experts and Selecting experts. The former task is further decomposed into Searching existing documents, Social networking, and using Experiences from previous interactions. The expertise selection task is performed by Social networking and using pervious experiences. The model in Figure 4 can also help in better understanding the knowledge market that may exist in an organization between expertise seekers and providers. The expertise seeker depends on the expert to provide him/her with adequate information or knowledge, teach necessary skills, or perform a task that the seeker is incapable of performing on his/her own. On the other hand, the expertise provider would want to Keep his/her job and Satisfy management in order to Get a promotion. Gaining reputation would also be helpful in Getting promotions, in addition to reducing the chances of being laid off and thus keeping the job. This goal can be (partially) achieved by helping expertise seekers. However, this task (Help expertise seekers) has a negative contribution to the goal of Not to be bothered too much which helps to get his/her own work done on time. Therefore, if the expertise provider is reputable enough, has a heavy workload, or even in case of competition between members, s/he may not be motivated or willing to provide sufficient information, teach, or perform a task for the expertise seeker. Of course, there could also be personal satisfaction in helping others which could motivate the expertise provider to help the seeker. Considering the goals of the expertise provider, therefore, it is possible to see the importance of introducing proper incentives to encourage individuals to share their knowledge with others. Depending on the culture of the organization, incentives can include individual recognition, status upgrades, or even monetary prizes.

An organization may also have IT systems and technology-oriented interfaces that actors depend on for their daily tasks. These may include search and retrieval tools, storage mechanisms, and technologies for communication and collaboration such as groupware and Wiki. The technologies, on the other hand, depend on the expertise seeker and provider (amongst others) to contribute contents to the systems. Thus, human actors and IT systems can also have knowledge relationships, although the knowledge market between them would depend on the embedded knowledge and adjusted settings provided by the developers and maintainers.

Expertise Finding Systems – Benefits and Pitfalls for Various Stakeholders

In this section, the analysis of the capabilities and functionalities of EFS is presented. The key requirements of an EFS, as pointed out in (Maybury, 2006), typically include the ability to identify experts, classify the type and level of expertise, validate the breadth and depth of expertise, and recommend experts. Based on this and the material presented in the previous sections, the functionalities and goals of the EFS are shown in Figure 5 without the consideration of the knowledge management needs that it can fulfill in an organization.

Insert Figure 5 Here

Figure 5. Goals, capabilities, and functionalities of the EFS technology

The EFS should be able to interact and work with other existing IT systems in an organization in order to be more effective. For automatic detection of user expertise, the system depends on other IT systems to provide it with information sources that can be indexed and parsed to extract expertise. Such information sources may include documents generated by users, past projects individuals worked on, contents of emails, forums, and bulletin boards, and traces of online or offline activities. In order to analyze the roles and knowledge processes required to make EFS successful in an organization, the interactions of the expertise seeker and provider with the EFS are illustrated in Figure 6. This model expresses the tasks that the roles expect the EFS to do, as well as the required tasks that need to be performed by expertise seeker and provider in order to make the EFS successful.

As can be seen in Figure 6, the introduction of the technology introduces some new tasks (shown with oval around them) for the expertise seeker and provider to do. For example, the new task of Create and maintain user profiles which is added to the responsibilities of the expertise provider and also Query EFS which is created for the expertise seeker would not have existed before. However, the new task of Provide feedback about interactions may have existed before and done somewhat informally, but with the introduction of the new technology its nature changes.

Insert Figure 6 Here

Figure 6. Interactions of expertise seeker and provide with the EFS

Now that the EFS is conceptualized as an agent and strategic dependencies are made explicit, it is possible to do goal evaluation to see if stakeholder goals are achieved and to determine the viability of the proposed solution. For example, the goal evaluation algorithm of (Horkoff, 2006) can be used. This algorithm assigns qualitative evaluation labels to the elements of the i* model according to a six point scale that range from satisfied, partially satisfied, conflict, unknown, partially denied to denied. The algorithm starts by assigning initial evaluation values and then continues by propagating these initial values through the network of actors using a combination of guidelines and human judgment (Strohmaier, et al., 2007). By conducting these kinds of analysis, it is possible to identify criteria for success and effects of failure of the EFS.

When, for example, the expertise provider does not perform the Create and maintain user profile task, then depending on whether s/he performs one or more of the other tasks satisfying Help expertise seekers or not[ii], this may or may not result in the propagation of the negative effect to Help expertise seekers, and in return have an effect on Satisfy management and Get promotion goals. However, this initial assignment would result in denying the resource dependency User Profiles and in return result in denying the Use profiles to determine expertise task of the EFS. Now if the ontologies are properly maintained and the resource dependency Information sources is satisfied, then this would have a positive impact on Use existing artifacts to determine expertise and would stop the negative impact of the denial of Use profiles to determine expertise to propagate any further. However, if the ontology engineers fail to Maintain Ontologies, then the resource dependency Shared Ontologies is denied. This in turn would have a significant negative impact on the ability of the EFS to identify and recommend experts and would prevent it from achieving its goal Improve and facilitate expert finding. As such it is important to note the importance of using multiple sources of information for identifying experts.

On the other hand, if the expertise seeker fails to Provide feedback, in addition to having a negative impact on his/her own goal Future need, this action denies the resource dependency Feedback from interactions. This in turn has a negative impact on Validate the breadth and depth of expertise task of the EFS, and this negative impact propagates to the root and finally prevents the EFS from achieving its goal Improve and facilitate expert finding. In addition, the assignment of denied to the resource dependency Feedback from interactions, would also have a negative impact on the Gain reputation goal of the expertise provider. Thus, expertise seekers should be motivated to use the system and provide feedback. The ease of use, simplicity, and familiarity of the system, trust in that content is up-to-date, and the accuracy of expertise identification, are all factors that encourage the use of the system. Promoting the use of the system through the use of offline and online communications that highlight the features of the system can also be effective in encouraging individuals to use the EFS.

Role of Ontology Engineers

The shared ontologies (see Figure 5) that are to be used to improve the EFS depend on ontology engineers to create and maintain them. This is an additional role that is needed if the system is to benefit from the use of ontologies. These engineers rely on domain experts, end-users, and management to provide them with the knowledge about the domain and existing standards that are used within the organization. They may also require the specifications of exiting IT systems and their information sources in order to create ontologies that are compatible with those systems. These knowledge dependencies, along with high level goals, capabilities, and responsibilities of ontology engineers are shown in Figure 7. In this model we can see the important role that tacit knowledge about the domain plays in the design of high quality shared ontologies. Of course, depending on the expertise domain, the amount and quality of the tacit knowledge that is required to be transferred to ontology engineers varies. In situations where there is an ongoing need for interactions with various actors to gain the required knowledge, the development of the ontology may not be successful. Actors may not be motivated enough to transfer their knowledge, especially since what they would be gaining in return may seem vague and non-immediate. To create a standard ontology, thus, it is important that the contributing actors have a positive attitude towards the ontology.

Insert Figure 7 Here

Figure 7. Goals, capabilities, and functionalities of Ontology Engineers

In dynamic environments, the maintenance of ontologies may become an important concern. The changes may occur due to changes in the environment the system is operating in or changes to the requirements of the users. To have a successful and up-to-date system, the necessary changes should be applied to the ontologies frequently. Ontology engineers depend on domain experts, end-users, and management to inform them of the new requirements and possible changes in the environment. These actors should be aware of the importance of ontology maintenance and its implications for the whole system. Having a trained team of individuals in the organization can improve and facilitate the maintenance process.

As already mentioned, designing large ontologies and insuring their consistency is an extremely complex and knowledge intensive task. Ontology engineers depend on ontology design tools for the creation and maintenance of ontologies. These tools typically use reasoners or theorem provers to provide feedback to the user about the logical implications of their design such as highlighting inconsistencies and redundancies (Horrocks, 2007). If the ontology maintenance is to be partly supported by individuals within the organization, then it is necessary for these actors to become familiar with these tools. It is important to note that if the ontologies are designed as modular as possible – i.e., if the ontology is divided into an upper level which rarely changes and an organic lower level that can be modified and changed by individuals – then ontology maintenance becomes more manageable.

EXPERT PROFILING AND RECOMMENDATION

We now briefly present the design and development of a system based on the analysis in the previous section by stating the key requirements of a successful ontology-based EFS and following each with a brief description. The interested reader is referred to (Fazel-Zarandi, Devlin, Huang, and Contractor, 2011; Fazel-Zarandi and Fox, 2010, 2011) for further details.

- EFS should use multiple sources of expertise information.

The analysis showed that if the important goal of Identify experts is denied, then the EFS would not be able to achieve its goal of Improve and facilitate expert finding. This in turn will discourage indiviudals from using the system. As such, the ability to construct accurate and complete expert profiles which can change over time using multiple sources of information is crucial to the success of an EFS. However, most existing systems use only one source of competency information. To this end, we use self-declarations of expertise in addition to expertise suggested from credentials and work experiences to create an initial model of the individual. We then make use of recommendations, "wisdom of the crowd", feedback from interactions, observations of online and offline activities, and content generated by the individual both within the organization and on the Web, in order to reason about individual's skills and expertise in a dynamic environment. These sources are used for validating the breadth and depth of expertise over time in order to improve profiles and provide reliable recommendations. Declarations of skills and expertise statements about individuals should be guided by a domain ontology in order to ensure common understanding of the terminology within the organization.

- EFS should identify experts who are neither under or over qualified.

Expertise is varying in level and often the most qualified expert is not needed. As such, classifying the level of expertise is important for identifying individuals who are qualified enough so that the most qualified experts do not get too many expertise requests. However, most of the existing systems do not take this factor into account. Level of expertise, or proficiency, may depend on different factors such as familiarity with the subject, the span of the activities that one can perform, how much experience one has in performing the activities, etc. To specify proficiency, we use ideas from the Measurement Ontology of (Kim, Fox, and Gruninger, 1999). We consider proficiency as determined by attributes related to that skill or expertise that can be measured and the span of activities that can be performed. For example, recency, years of experience, and average number of errors found are some of the attributes that can be used to measure proficiency in programming. This provides a very general representation which allows for measuring a variety of attributes. These attributes are input to the model (primitive types), can be objective or subjective and categorical or numeric. Each attribute takes up values from a specification set defined by its elements and an ordering between the elements. All measured attributes must have a specification set, and the value for an attribute must be an element of this set. We consider five different proficiency levels: novice, advanced beginner, competent, proficient, and expert and an ordering between them. For each skill, the required value of a particular measured attribute related to it is identified by domain experts for different levels of proficiency. Individuals are then assessed and/or observed and relevant attributes are measured. In other words, proficiency assessment is done through a series of activities that perform measurement.

Different matchmaking strategies exist for matching individuals to requirements for expertise finding. In general, matchmaking strategies based on purely logic deductive facilities present high precision and recall, but are often characterized by low flexibility (Bianchini, De Antonellis, and Melchiori, 2008). Flexibility refers to the ability to recognize the degree of similarity when an exact match does not exist. On the other hand, similarity-based approaches are characterized by high flexibility, but limited precision and recall (Bianchini, et al., 2008). To take advantage of the benefits of logic-based and similarity-based approaches, we first use a deductive model to determine the kind of match between an individual and requirement set, and then based on the kind of match determine the similarity measure to use in order to rank the individuals with partial match. The details of this procedure are out of the scope of this document. The interested reader is referred to (Fazel-Zarandi and Fox, 2009).

- Ontology design and maintenance should be a joint effort by ontology engineers and individuals within the organization.

As pointed out in the analysis, the shared ontologies should be organized into sub-ontologies or modules. At a high level, we divide them into upper level which rarely changes and an organic lower level which can be modified and changed by individuals within the organization. In particular, for integrating information from different sources and inferring and validating expertise over time, we reuse and extend: 1) the Process Specification Language (PSL) (Gruninger and Menzel, 2003) which is a first-order language for modeling processes comprised of a layered collection of families of axioms; 2) the Organization Ontology (Fox, et al., 1996) which formalizes the organizational structure; and 3) the Trust Ontology of (Huang & Fox, 2006). PSL provides predicates and axioms that enable representation of and reasoning about fluents[iii], activities, activity-occurrences, and values of fluents before and after activity-occurrences. For example, the Activity-Occurrence Extension of PSL defines relations that allow the description of how activity-occurrences relate to one another with respect to the time at which they start and end; and the State Extension introduces the concept of state (before an activity-occurrence) and post-state (after an activity-occurrence). We also assume a taxonomy of skills and expertise in the specific domain of interest. This taxonomy will ensure a common understanding of the terminology used and will help guide skill and expertise declarations. Individuals are allowed to modify or add terms to the taxonomy and the change will be communicated to members of the organization.

- EFS should be interoperable with the existing IT infrastructure.

The EFS should be able to interact and work with other existing IT systems in an organization in order to be more effective. To this end, the incoming data from divergent sources are mapped to an occurrence of a particular activity and added to the knowledgebase. Examples of such activities include: *performs* denoting an individual *x* performing an activity in the workplace, *communicates* denoting *x* communicating information to another individual, and *creates-content-on* denoting *x* has created content related to a particular skill or knowledge field. To map incoming data to a particular activity, we use existing vocabularies and ontologies. For example, we use the Dublin Core (Nilsson, Powell, Johnston, and Naeve, 2008) which represents publication metadata for describing physical resources such as books, digital materials such as text files and source codes, and composite media such as web pages. In particular, the attributes *creator*, *contributor*, *subject* (represents the topic of the resource), and *type* (represents the nature or genre of the resource) are used to generate an occurrence of the *created-content-on* activity for the *creator*. Another example is the tagging ontologies of (Kim, Passant, Breslin, Scerri, and Decker, 2008) and (Passant and Laublet, 2008) which can be used for declarations of expertise by others. In this case, an occurrence of the *declares* activity is added to the knowledge-base. In the absence of meta-data, an information retrieval component relates documents and online activities to the domain of expertise and skills that they identify using typical information retrieval techniques (e.g., lexical pattern matching and keyword extraction, indexing). For example, if this component relates document *d* created by individual *x* to domain *y*, then an occurrence of the activity *created-content-on* is added to the knowledge-base. Based on this added activity and prior information, further knowledge about *x*'s skills and proficiency is then inferred and stored in the database for future querying.

- EFS should use social and contextual information for recommending experts.

From the perspective of the expertise provider, Help expertise seeker and Finish own required functions are conflicting goals. As stated in the analysis section, if the expertise provider is reputable enough, has a heavy workload, or even in case of competition between members, s/he may not be motivated or willing to provide sufficient information, teach, or perform a task for the expertise seeker. As such, it is important to use social and contextual information to select among qualified experts. To this end, contextual information such as workload and availability are factors that can be taken into account. Leveraging social science insights on how and why people collaborate with each other successfully is also useful. Considering previous interactions between individuals, for example, it is possible to estimate the willingness of the expertise provider to help the expertise seeker in terms of the supply and demand of resources that each had to offer. The interested reader is referred to (Fazel-Zarandi, et al., 2011).

- EFS should provide an easy to use feedback mechanism.

In order to guarantee reliable recommendations, an EFS should provide an easy to use feedback mechanism. This would be similar to online question answering forums where the receiver of the information can provide feedback on the usefulness of the information provided and there is usually a rating of the answer that can be used to imply proficiency. In addition to verifying the depth and breadth of competency and knowledge of the expertise provider, such feedback mechanism encourages expertise providers to help others by gaining them indivdiual recognition.

FUTURE RESEARCH DIRECTIONS

Three promising research directions for further extension of ontology-based EFS are as follows. First, future work should focus on the evaluation of ontology-based EFS in real world settings. In particular, an interesting topic is to evaluate the expert profiles by measuring how close they are to reality, under what circumstances the profiles converge into the real evaluation, and how important a source is in increasing the quality of the information that is known about an individual. In addition, further work should be done on evaluating the usefulness of ontology-based EFS in different domains and studying how structure within a particular domain relates to the general model.

Second, in most organizations individuals are usually members of teams set up to pursue specific projects. Teams are temporary in nature and are set up when needed. Selecting teams is a complex problem because of the importance of a number of different variables. When considering individuals for teams, for example, complexities may arise due to fitness between team members. In addition, research in social sciences on how and why people form teams has shown that goals and context such as exploring new ideas and resources or exploiting existing resources and capabilities have different levels of impact within and across communities (Contractor, Wasserman, and Faust, 2006). For example, heterogeneous well-balanced teams are more effective than homogenous groups on complex projects and innovative problem solving tasks (Krass and Ovchinnikov, 2006). Furthermore, project requirements, personal and technical characteristics of human resources and their availability are important factors that influence the success of a team. Most of the existing EFS focus on identifying single individuals and do not recommend teams of experts. It would be interesting to see how EFS could be extended to address these complexities.

Finally, a third area of future research should focus on the problem of *selecting* expert. Although over the past decade or so much effort has gone into creating techniques to increase and evaluate the recommendation quality for objects such as books and movies, the personalized search for subjects such as experts in a particular field has not so far received much attention (Hansen, et al., 2010; Malinowski, Keim, Wendt, and Weitzel, 2006). Previous studies that have attempted to provide personalized search for subjects can be divided into two groups. Techniques in the first approach consider user preferences and make recommendations based on the degree of similarity between expert profiles and user preferences. Systems in the second approach, on the other hand, consider social relations and network statistics. Examples include Referral Web (Kautz, Selman, and Shah, 1997) which uses coauthoring and co-citation relationships and Expert Recommender (McDonald and Ackerman, 1998) which considers friendship and departmental relations. Useful recommendation of experts, however, depends significantly on the motivation of the user in seeking the recommendation (Contractor and Monge, 2002; Fazel-Zarandi, et al., 2011). For example, looking for a quick answer versus a potential collaborator will require very different recommendation strategies.

CONCLUSION

This chapter presented the analysis of the expertise finding problem in a general setting using the i* modeling notation. This method allows for the systematic modeling of EFS and its effects on the goals of the stakeholders. Based on the findings, apart from the benefits and challenges of incorporating ontologies in systems, the social and organizational factors, such as having positive attitude towards the ontology, interactions between different actors, and having a trained team of individuals responsible for ontology maintenance, are important for the successful deployment of an EFS.

The chapter also presented the design of a EFS based on the findings. The system uses multiple sources of information to gather expertise information. Starting with less than accurate profiles of individuals, additional skills and expertise are inferred based on the activities individuals participate in and the contents they generate. Breadth and depth of expertise are then validated over time by using recommendations, "wisdom of the crowd", and peer-reviews. The profiles are then used to provide reliable recommendations for improving and facilitating expertise finding.

The role of the executive leadership in the application of these findings is critical. They can help motivate users to participate, establish realistic expectations, communicate progress and successes, and overcome barriers. They can encourage a culture of mutual support and knowledge sharing, and help change corporate culture from being competitive to being cooperative.

REFERENCES

Adomavicius, G., & Tuzhilin, A. (2005). Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Transactions on Knowledge and Data Engineering, 17*(6), 734-749.

Aleman-Meza, B., Bojars, U., Boley, H., Breslin, J., Mochol, M., Nixon, L., et al. (2007). *Combining RDF Vocabularies for Expert Finding*. Paper presented at the 4th European Semantic Web Conference (ESWC2007).

Bianchini, D., De Antonellis, V., & Melchiori, M. (2008). Flexible semantic-based service matchmaking and discovery. [Article]. *World Wide Web Journal, 11*(2), 227-251. doi: 10.1007/s11280-007-0040-y

Biesalski, E., & Abecker, A. (2006). *Skill-Profile Matching with Similarity Measures*. Paper presented at the 8th International Conference on Enterprise Information Systems (ICEIS), Paphos, Cyprus.

Borgatti, S. P., & Cross, R. (2003). A relational view of information seeking and learning in social networks. *Management Science, 49*(4), 432-445.

Braun, S., & Schmidt, A. (2008). *People Tagging & Ontology Maturing: Towards Collaborative Competence Management*. Paper presented at the Proceedings of the 8th International Conference on the Design of Cooperative Systems (COOP), Carry-le-Rouet, France.

Burke, R. (2005). Hybrid systems for personalized recommendations. In B. Mobasher & S. S. Anand (Eds.), *Intelligent Techniques for Web Personalization* (Vol. 3169, pp. 133-152). Berlin: Springer-Verlag.

Casciaro, T., & Lobo, M. S. (2005). Competent jerks, lovable fools, and the formation of social networks. *Harv Bus Rev, 83*(6), 92-99, 149.

Colucci, S., Di Noia, T., Di Sciascio, E., Donini, F. M., Mongiello, M., & Mottola, M. (2003). A formal approach to ontology-based semantic match of skills descriptions. *Journal of Universal Computer Science, 9*(12), 1437-1454.

Colucci, S., Di Noia, T., Di Sciascio, E., Donini, F. M., & Ragone, A. (2007). Semantic-based Skill Management for Automated Task Assignment and Courseware Composition. *Journal of Universal Computer Science, 13*(9), 1184-1212.

Contractor, N. S., & Monge, P. R. (2002). Managing Knowledge Networks. *Management Communication Quarterly, 16*(2), 249-258.

Contractor, N. S., Wasserman, S., & Faust, K. (2006). Testing multitheoretical, multilevel hypotheses about organizational networks: An analytic framework and empirical example. *Academy of Management Review*, 681-703.

Dittmann, L. (2003). Towards Ontology-based Skills Management *Project KOWIEN Report*: University Duisburg-Essen.

Dorn, J., Naz, T., & Pichlmair, M. (2007). *Ontology Development for Human Resource Management*. Paper presented at the 3rd International Conference on Knowledge Managements.

Earl, M. (2001). Knowledge management strategies: Toward a taxonomy. *Journal of Management Information Systems, 18*(1), 215-233.

Fazel-Zarandi, M., Devlin, H. J., Huang, Y., & Contractor, N. S. (2011). *Expert Recommendation based on Social Drivers, Social Network Analysis, and Semantic Data Representation.* Paper presented at the 2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec 2011), Chicago, USA.

Fazel-Zarandi, M., & Fox, M. S. (2009). *Semantic Matchmaking for Job Recruitment: An Ontology-Based Hybrid Approach*. Paper presented at the 8th International Semantic Web Conference.

Fazel-Zarandi, M., & Fox, M. S. (2010). *Reasoning about Skills and Competencies*. Paper presented at the 11th IFIP Conference on Virtual Enterprises (PRO-VE 2010), Saint-Etienne, France.

Fazel-Zarandi, M., & Fox, M. S. (2011). *Constructing Expert Profiles over Time for Skills Management and Expert Finding*. Paper presented at the 11th International Conference on Knowledge Management and Knowledge Technologies (i-KNOW 2011), Graz, Austria.

Fazel-Zarandi, M., & Yu, E. (2008). Ontology-Based Expertise Finding (pp. 232-243).

Fox, M. S., Barbuceanu, M., & Gruninger, M. (1996). An organisation ontology for enterprise modeling: Preliminary concepts for linking structure and behaviour. *Computers in Industry, 29*(1-2), 123-134. doi: 10.1016/0166-3615(95)00079-8

Fox, M. S., & Gruninger, M. (1998). Enterprise Modelling. *AI Magazine*, 109-121.

Gomez-Perez, A., Ramirez, J., & Villazon-Terrazas, B. (2007). *Reusing Human Resources Management Standards for Employment Services*. Paper presented at the Workshop on the First Industrial Results of Semantic Technologies (FIRST'2007) at ISWC/ASWC'2007.

Gruber, T. R. (1993). Toward Principles for the Design of Ontologies Used for Knowledge Sharing *KSL* (pp. 93-04): Knowledge Systems Laboratory, Stanford University.

Gruninger, M., & Menzel, C. (2003). Process Specification Language: Theory and Applications. *AI Magazine, 24*, 63-74.

Guarino, N. (1998). *Formal Ontology and Information Systems*. Paper presented at the Formal Ontologies in Inforamtion Systems (FOIS'98), Trento, Italy.

Hansen, D. L., Khopkar, T., & Zhang, J. (2010). Recommender Systems and Expert Locators. *Encyclopedia of Library and Information Sciences, Third Edition*, 4433 - 4441.

Horkoff, J. (2006). *Using i* Models for Evaluation.* Master's Thesis, University of Toronto, Toronto, Canada. Retrieved from http://www.cs.utoronto.ca/~jenhork/MScThesis/Thesis.pdf

Horrocks, I. (2007). *Semantic Web: The Story So Far*. Paper presented at the International Cross-Disciplinary Conference on Web Accessibility (W4A).

Huang, J., & Fox, M. S. (2006). *An Ontology of Trust - Formal Semantics and Transitivity.* Paper presented at the the 8th International Conference on Electronic Commerce.

Huang, Y., Contractor, N. S., & Yao, Y. (2008). *CI-KNOW: recommendation based on social networks.* Paper presented at the 2008 International Conference on Digital Government Research, Montreal, CA.

Kautz, H., Selman, B., & Shah, M. (1997). Referral web: Combining social networks and collaborative filtering. [Article]. *Communications of the Acm, 40*(3), 63-65.

Kim, H., Passant, A., Breslin, J. G., Scerri, S., & Decker, S. (2008). *Review and alignment of tag ontologies for semantically-linked data in collaborative tagging spaces.* . Paper presented at the 2nd International Conference on Semantic Computing, San Francisco, USA.

Kim, H. M., Fox, M. S., & Gruninger, M. (1999). An Ontology for Quality Management - Enabling Quality Problem Identification and Tracing. *BT Technology Journal, 17*(4), 131-140.

Krass, D., & Ovchinnikov, A. (2006). The University of Toronto's Rotman School of Management Uses Management Science to Create MBA Study Groups. *Interfaces, 36*(2), 126-137.

Lau, T., & Sure, Y. (2002). *Introducing Ontology-Based Skills Management at a Large Insurance Company*. Paper presented at the Modellierung.

Liu, P., & Dew, P. (2004). *Using Semantic Web Technologies to Improve expertise Matching within Academia*. Paper presented at the 11th International Conference on Knowledge Management and Knowledge Technologies (I-KNOW), Graz, Austria.

Malinowski, J., Keim, T., Wendt, O., & Weitzel, T. (2006). *Matching People and Jobs: A Bilateral Recommendation Approach*. Paper presented at the 39th Annual Hawaii International Conference on System Sciences (HICSS'06), Kauai, HI.

Maybury, M. T. (2006). Expert finding systems (pp. The MITRE Corporation). Bedford, Massachusetts, USA: MITRE Center for Integrated Intelligence Systems

McDonald, D. W., & Ackerman, M. S. (1998). *Just talk to me: a field study of expertise location*. Paper presented at the 1998 ACM Conference on Computer Supported Cooperative Work, Seattle, Washington, United States.

Mochol, M., Wache, H., & Nixon, L. (2007). *Improving the Accuracy of Job Search with Semantic Techniques*. Paper presented at the 10th International Conference on Business Information Systems, Poznan, Poland.

Nilsson, M., Powell, A., Johnston, P., & Naeve, A. (2008). Expressing Dublin Core metadata using the Resource Description Framework (RDF), from http://www.dublincore.org/documents/dc-rdf/

Passant, A., & Laublet, P. (2008). *Meaning of a tag: A collaborative approach to bridge the gap between tagging and Linked Data.* . Paper presented at the WWW 2008 Workshop Linked Data on the Web (LDOW2008), Beijing, China.

Razmerita, L., Angehrn, A., & Maedche, A. (2003). *Ontology-based User Modeling for Knowledge Management Systems*. Paper presented at the 9th International Conference on User Modeling.

Schafer, J. B., Frankowski, D., Herlocker, J., & Sen, S. (2007). Collaborative filtering recommender systems. In B. Peter, K. Alfred & N. Wolfgang (Eds.), *The adaptive web* (pp. 291-324): Springer-Verlag.

Schmidt, A., & Kunzmann, C. (2006). Towards a Human Resource Development Ontology for Combining Competence Management and Technology-Enhanced Workplace Learning. *Lecture Notes in Computer Science, 4278*, 1078-1087.

Staab, S., Studer, R., Schnurr, H. P., & Sure, Y. (2001). Knowledge Processes and Ontologies. *IEEE Intelligent Systems, 16*(1).

Stankovic, M., Wagner, C., Jovanovic, J., & Laublet, P. (2010). *Looking for Experts? What can Linked Data do for You?* Paper presented at the WWW 2010 Workshop Linked Data on the Web (LDOW2010), Raleigh, NC.

Strohmaier, M., Yu, E., Horkoff, J., Aranda, J., & Easterbrook, S. (2007). *Analyzing Knowledge Transfer Effectiveness – An Agent-Oriented Modeling Approach*. Paper presented at the 40th Annual Hawaii International Conference on System Sciences (HICSS'07), HI, USA.

Veit, D., Müller, J. P., Schneider, M., & Fiehn, B. (2001). *Matchmaking for autonomous agents in electronic marketplaces*. Paper presented at the 5th International Conference on Autonomous Agents, Montreal, Quebec, Canada.

Yimam-Seid, D., & Kobsa, A. (2003). Expert-finding systems for organizations: Problem and domain analysis and the DEMOIR approach. [Article]. *Journal of Organizational Computing and Electronic Commerce, 13*(1), 1-24.

Yu, E. (2009). Social Modeling and i*. In A. T. Borgida, V. Chaudhri, P. Giorgini & E. S. Yu (Eds.), *Conceptual Modeling: Foundations and Applications - Essays in Honor of John Mylopoulos* (Vol. 5600, pp. 99-121): LNCS, Springer.

Yu, E., Giorgini, P., Maiden, N., & Mylopoulos, J. (2011). *Social Modeling for Requirements Engineering* Cambridge, MA: MIT Press.

---

[i] *Make* is a contribution strong enough to satisfy a *softgoal*, *Help* is a partial positive contribution which is not by itself sufficient enough to satisfy a *softgoal*, and *Some+* is a positive contribution whose strength is unknown.

[ii] The *means-ends* relationship will take the maximum value of its children, with satisfied being the highest value and denied being the lowest value

[iii] Fluents are properties of the real world that can change over time. There are two types of fluents: *relational fluents* refer to relations that have true values of "true" or "false"; and *functional fluents* refer to the functions as defined in mathematics (Huang, 2008).