

SCALE: Supporting Community Awareness, Learning, and Evolvement in an Organizational Learning Environment

Amy Soller¹, Renata Guizzardi³, Alessandra Molani² & Anna Perini¹

¹ITC-IRST
Via Sommarive 18
38050 Povo, Trento Italy
+39 0461 314 358
soller@itc.it

²University of Trento
Dept of Information and Communication Technology
Via Sommarive 18, 38050 Povo, Trento Italy
+39 0461 314 359
molani@dit.unitn.it

³University of Twente
7500 AE Enschede
The Netherlands
+31 53 489 3786
souza@cs.utwente.nl

Abstract: The SCALE objective is to develop a toolbox of intelligent solutions designed to enhance the current state-of-the-art distributed organizational learning and knowledge management technology. This toolbox will attempt to mediate and facilitate the sharing of knowledge between peers, and promote community awareness, learning, and evolvement. In this paper, we describe the first two intelligent tools that we have begun to develop toward this effort. The first is a social network based visualization tool, intended to raise users' awareness of existing communities of practice, and the second is a socially aware recommendation agent.

Introduction

As an organization develops, its knowledge and expertise becomes increasingly distributed. While this process promotes the growth of specialized knowledge communities, it also makes discovering relevant knowledge from relevant communities more difficult, and understanding the specialized nature and functions of the newly discovered knowledge even more challenging (Hoadley & Pea, 2002). When knowledge is shared effectively across a large distributed organization with pools of specialized expertise, the organization may experience improved learning and development, and increased productivity and growth (Brown & Duguid, 1991; Soller & Lesgold, in press). The SCALE research project aims to develop tools for supporting knowledge sharing and community awareness in distributed knowledge organizations. In this paper, we describe two intelligent SCALE tools that are currently being developed to enhance the KEEx peer-to-peer organizational learning environment (Bonifacio, Cruel, Mameli, & Nori, 2002): a social network based visualization tool, and a socially aware recommendation agent.

Organizational learning environments electronically aid employees in the discovery, sharing, learning, and application of knowledge (Ayala, 2001). These systems support the natural processes within an organization that promote knowledge exchange and construction. Researchers that study organizational learning explain that organizations naturally share knowledge by forming small groups based on similar interests, personal affinity, and trust. These groups are termed *Communities of Practice (CoPs)* (Brown & Duguid, 1991; Lave & Wenger, 1991) because they function within the organization as cohesive communities that share a common sense of purpose and interest. CoPs facilitate the sharing and creation of new knowledge, and are therefore important to the stability and growth of an organization. Members of CoPs are motivated to discover, join, and create new CoPs because of their interest, desire to learn, and inherent information seeking behaviors. In an organizational learning environment, the knowledge that users might discover was created, used, and transformed by a CoP, and hence derives its meaning in part from its roots in that community. This means that when learners seek and share knowledge, they are also sharing social and cultural references that carry different meanings in different social communities (Mantovani, 1996). The knowledge sharing process is hence intertwined with the processes that form the fundamental social practices of CoPs such that when users share and create new knowledge in a distributed learning environment, they naturally contribute toward the evolution of the communities grounding that knowledge (Scardamalia & Bereiter, 1994). The role of the computer environment, then, is to support, mediate, and guide these processes of knowledge seeking and sharing, and community learning and evolvement (Jermann, Soller, & Muehlenbrock, 2001).

The SCALE initiative centers around supporting the learning, development, and growth of communities of practice. In particular, we are developing a toolbox of solutions, grounded in real business process analyses,

designed to enhance the current state-of-the-art organizational learning technology. This toolbox will attempt to mediate the exchange of knowledge between learners by (a) helping users become aware of each other and their communities, (b) promoting interaction, knowledge sharing, and unintentional, organizational learning, and (c) facilitating the evolution of community practices.

In this paper, we briefly present an example analysis of an actual organizational learning case driving our development of new technological solutions. We then describe the first two tools that we have begun to develop to address the issues uncovered in our organizational analyses. The first is IVisTo, an intelligent social network based visualization tool, intended to raise users' awareness of existing CoPs and available knowledge sources. IVisTo displays a weighted combination of social networks, showing the learner and his peers, and their relation to the existing knowledge communities and artifacts. The second tool is a distributed team of socially aware intelligent recommendation agents that might, for example, put the learner in contact with an online expert or instructor from a strategically selected community of practice. We first briefly describe our open-ended organizational learning environment, and then our business analysis method, and socially-aware technological solutions under development.

An Open-Ended Distributed Organizational Learning Environment

The SCALE initiative builds upon a foundation of traditional knowledge management technology to develop a uniquely distributed peer-to-peer environment for supporting organizational learning practices. Our environment is self-sustaining, and operates without centralized servers or databases. The distributed nature of this peer-to-peer solution is intended to reflect the naturally distributed nature of knowledge and expertise in an organization. Peer-to-peer technology supports the horizontal relationship between people, seeing them as both learners and designers of knowledge. Each learner (or peer) controls his own personal knowledge ontology containing artifacts (e.g. documents, videos, web pages), and can decide to exchange knowledge with other peers based on, for example, their common interests, roles, expertise, or trust. Because employees in an organization interact similarly to peers in a peer-to-peer network, naturally grouping themselves in communities, the distributed peer-to-peer knowledge platform seems to be a natural choice for supporting the creation and sharing of knowledge.

The KEEEx (Knowledge Enhancement and Exchange) open-ended peer-to-peer organizational learning environment (Bonifacio et al., 2002) (Figure 6) serves as the foundation for the intelligent mediation tools described in the following sections. KEEEx allows users to share and search for artifacts, without requiring them to agree on a common representation language or ontology. Each user "publishes" his knowledge in the form of a concept hierarchy, or *context*, containing relevant artifacts. Then, when a learner seeks information, he simply runs a keyword or context query, prompting the system to search the other learners' contexts available on the network and their associated artifacts, and return a list of users, contexts, and discovered knowledge items. The system uses a sophisticated context matching algorithm to match the query's focus with the available, published contexts (Bouquet, Serafini, & Zanobini, 2003). The learning environment is open-ended in the sense that it does not include a curriculum or set of predetermined exercises; instead it is designed for workers to use in their everyday learning and work activities. In the following section, we describe an actual business case derived from the data we collected during the deployment of this system in our client company. We then describe our research efforts to enhance the KEEEx environment with socially aware plug-in software components for facilitating the awareness, exchange and creation of knowledge within and across Communities of Practice (CoPs).

User and Community Modeling

Although a distributed peer-to-peer system might reflect the natural environment in which members interact and share knowledge, it also maintains the distributed nature of the knowledge that would otherwise help to explicitly identify relevant knowledge and communities of practice. The expertise in a distributed organization exists sparsely in the (internally understood and shared) members' understanding of each other's knowledge, and in the (hidden) behavioral and cognitive similarities among individual users. A distributed environment, therefore, must make special provisions to help uncover and make salient some of these hidden characteristics so that learners might become aware of the existing organizational knowledge and its corresponding CoPs. In this section, we describe our efforts towards identifying, quantifying, and modeling these "hidden" learner characteristics – those that make the difference between retrieving relevant knowledge from key communities of practice, and retrieving long lists of documents resulting from keyword matches.

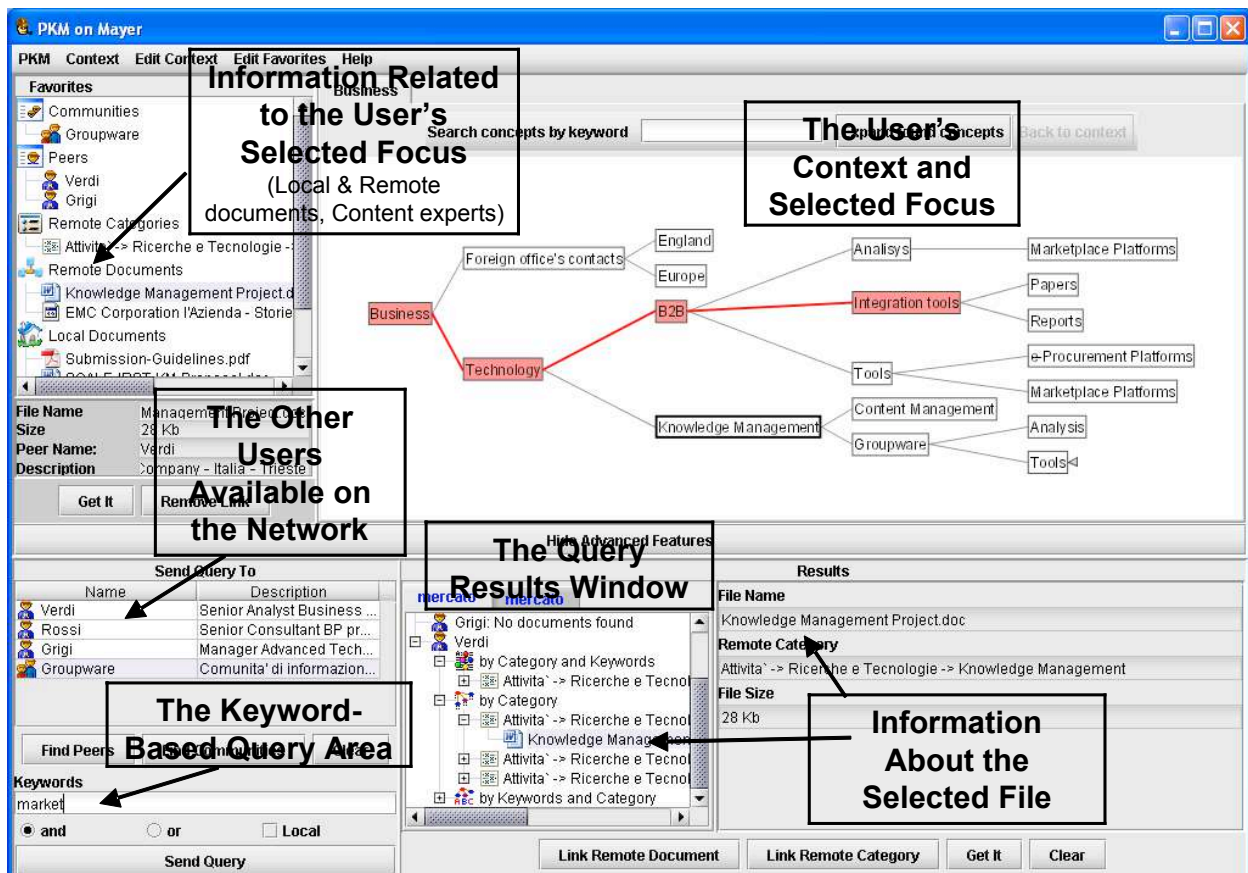


Figure 6. The KEEx Interface, showing the documents in the user's selected focus (top left), the user's context and selected focus (top right), the user's query (bottom left window), and the query results (bottom right).

Identifying Key User Model Variables through Intentional Analysis

The learner characteristics that shape a community of practice will tend to change slightly, depending on each organizational situation. For example, a newcomer to an organization, who needs to learn the fundamentals of the project to which he has just been assigned, might be interested in characteristics such as his peers' willingness to help and trustworthiness. A middle manager however, learning how to select team members for an interdisciplinary project, might be more interested in factors such as role, and teamwork or collaboration history. In the first stage of the project, we performed an *i** intentional analysis in an actual company (Molani, Perini, Yu, & Bresciani, 2003). The *i** methodology is designed to elicit the technological requirements for supporting organizational learning practices by modeling the relations between actors, goals, tasks, and resources. We used the methodology to identify the critical factors that influence the network of relationships between individuals, teams, and their learning tasks, allowing us to identify a set of critical user model variables. Figure 7 illustrates the *i** analysis that was performed for our case in which a newcomer has just joined a workgroup. The Middle Manager (shown as the circle in the upper left corner) plays the role of the supervisor for the project to which the Newcomer has just been assigned. She is responsible for guiding the Newcomer in learning the key concepts of the project by providing him with references (e.g. books, URLs), and relevant expert personnel. The Newcomer's role is to determine which actions will help him learn these concepts quickly and efficiently. The actors' goals, subgoals, and resources (for achieving these goals) are represented within the large shaded actor balloons. For instance, the Newcomer balloon shows that he has the general goal of *working well* in the new workgroup (depicted as a cloud, or "softgoal" in *i** notation), which means that he must learn about the project to which he has been assigned (*seek information* goal). This goal is refined into three subgoals: *asking an expert*, *consulting literature*, and *talking to colleagues*. The first two were actual suggestions from the Middle Manager, while the third was the action taken by the Newcomer. This was not surprising, because in the company we analyzed, knowledge is commonly shared, and knowledge sharing often occurs through informal interaction between peers or members of a CoP. One reason why the Newcomer may have chosen to talk to his peer instead of contacting the expert who was suggested by his supervisor, may have been that

the Newcomer had a softgoal of *avoiding being prematurely judged*. This is also not surprising, because the Newcomer, having just started work in the organization, wanted to maintain a good image.

In the analysis framework, we say that the softgoal, *avoiding being prematurely judged*, contributes positively toward the goal *talking to colleague*, but negatively toward the goal *asking an expert*. This is because, from the Newcomer’s point of view, the Colleague is someone who has the resource, *time* (depicted as a square), and is willing to collaborate (depicted as the resource, *good collaborative level*, and the goal, *being helpful*). The Colleague also has a similar knowledge base (*level of expertise*), and may be able to speak a similar language as the Newcomer while sharing his knowledge and organizational experience. The Expert’s bubble shows that, although he is highly competent (*high level of expertise*), he has more power and evaluative potential than the Colleague (*power role resource*). In designing user models to support such organizational practices, these are the factors that appear to impact the Newcomer’s decision making process the most.

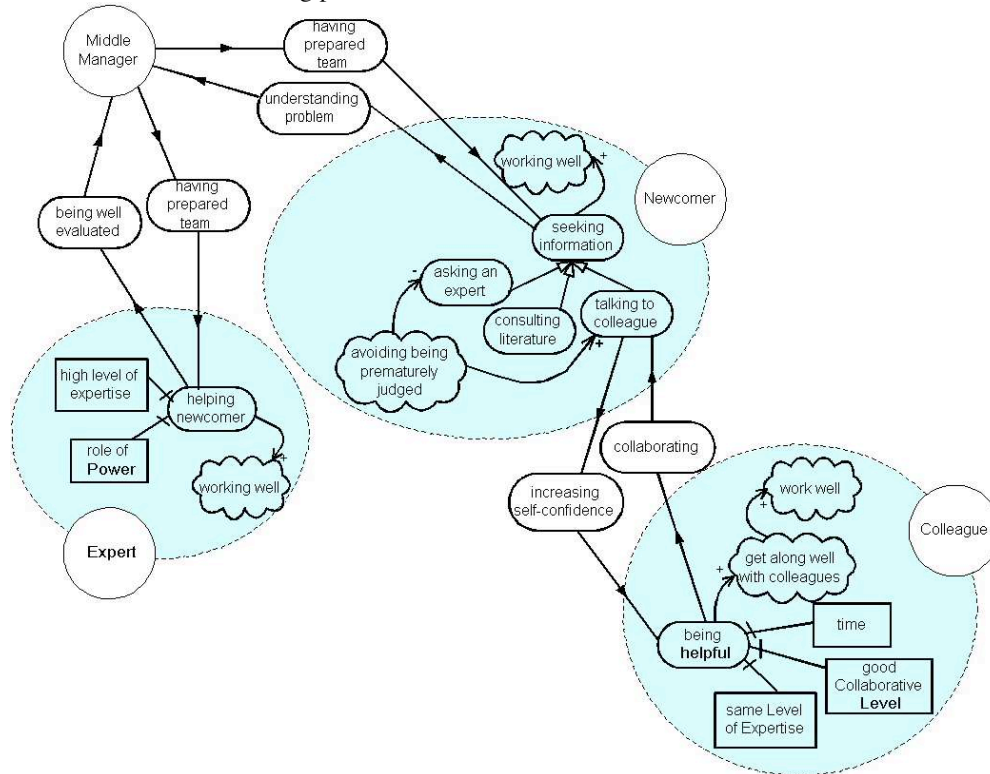


Figure 7. Intentional Analysis based on the i* framework, analyzing the Newcomer’s perspective

User Modeling in KEEx

The analysis described in the previous section uncovered several key characteristics influencing the sharing and learning of knowledge across an organization, during practices in which newcomers join workgroups. These characteristics, defined as resources in the i* analysis, were used to develop our user model variables (see Table 3). Then, two techniques were designed to elicit the initial values for the identified variables. The first involves training the organizational learning system (KEEx) to observe and model the users’ behavior, and the second involves explicitly asking the users their opinions through questionnaires and interviews. In the first case, the system *infers* the information, while in the second case the information is *elicited* from the user. In the next section we will see how the SCALE software tools will enable KEEx to intelligently update the learner models by observing his activities and interaction with peers (Gaudioso & Boticario, 2003; Vassileva, McCalla, & Greer, 2003).

Because we are working within the framework of a peer-to-peer system, we need to respect the privacy of learners. Each user should be able to decide which information, in his own user model, he would like to keep private. For example, users may choose not to share all of the information about whom they do and do not trust. The information in each user’s model that he chooses not to share will only be visible in his personal “perspective-based” model of the online community, and will be missing from other users’ assessment of the knowledge available

in the environment. As in (Vassileva et al., 2003), we address privacy issues by assigning a personal assistant agent to each learner in the peer-to-peer network. This agent interacts with its user to understand his preferences, and negotiates what user model information to disclose to other agents during the knowledge sharing process.

Table 3. The Intentional Analysis resources identified and corresponding User Model variables

Intentional Analysis Resource	User Model Variable
Good collaborative level	Trust, History of Collaboration
Power role	Organizational Role
Time	Availability
Same level of expertise	Level of Expertise, Problem Solving Strategy
High level of expertise	Level of Expertise

Intelligent Tools for Supporting Communities of Practice

The IVisTo Interactive Visualization Tool

The user models described in the previous section are designed so that a distributed organizational learning system can help users develop their virtual identity online, while increasing their awareness of existing communities of practice, facilitating their access to knowledge in different communities, and supporting the creation and evolution of new communities of practice. In working toward the first two goals, we have begun to transform the user models into dynamic social network oriented visualizations of users' behavior. These visualizations are designed to raise users' awareness of the types of knowledge communities that exist in the distributed network. Visualizations that reflect users' behaviors have been shown to increase their awareness of these behaviors, influencing their emotional attitudes, and enhancing collaboration, reflection, and motivation (Gutwin, Stark, & Greenberg, 1995; Jermann, Soller, & Muehlenbrock, 2001; Reimann, 2003). For example, the Comtella peer-to-peer system (Bretzke & Vassileva, 2003) enables users to visualize their significance in the community through a dynamic night sky background that displays each user as a star. The size and brightness of the stars denote the number of relationships or resources the users hold, respectively, and groups of peers with similar interests are represented as galaxies. Similarly, our social network based visualizations are designed to increase users' awareness of how their behavior relates their peers' behavior, as defined by our business practice analysis.

The IVisTo tool, shown in Figure 8, presents a novel way to visually rank the relevance of a learner's query results, taking into consideration both the social community-oriented information, and the more traditional and accessible lexical and semantic similarity information provided by the query-response engine. IVisTo displays a weighted combination of social networks, where each social network addresses a different user model variable, and the weights are given by the learner's social and semantic preferences. The bottom half of the interface contains a set of slider bars representing the social variables in the user model (e.g. Role, Trust), and the lexical and semantic attributes given by the KEEx lexical and context matching algorithms. Using these slider bars, the user can indicate the importance, or *weight*, of each variable. Behind the scenes, the system, generates a social network for each of the user model variables, and then computes one single network by calculating a weighted sum of the individual networks. For example, Sally's visualizations show her as an object in the center of the screen, and her peers as objects in the periphery, while the length of the links between her and her peers suggests their degree of similarity according to each user model variable. In the case of "Trust", the length of the links would suggest the degree to which Sally trusts (or would trust, based on the system's assessment) each of her peers. In the case of "Role" or "Expertise", the length of the links suggests the degree to which Sally holds a similar role or level of expertise as each of her peers. In this way, IVisTo provides each learner across an organization with a personalized set of visualizations from his perspective, weighted according to his interests. In the case of the Newcomer (from the example described previously), factors such as whether or not a peer is trusted, available, and willing to collaborate may be important, whereas the Content Expert may be more interested in obtaining knowledge having a high degree of lexical and semantic accuracy. The folder icons attached to the user nodes show the information available from each user; clicking on a folder icon enables a user to download a knowledge item directly from the provider.

As the learner carries on her day-to-day learning and collaborative work, the lengths of the links in IVisTo are re-calculated as the elicited and inferred information in her user model is updated. For instance, if Sally downloads a knowledge item (e.g. document or video) from Harry, adds it to her personal knowledge ontology, and accesses it more than once during the course of the day, the system would increase her user model values for

collaboration and expertise similarity with regard to Harry. These types of activities help the system intelligently infer and visualize different types of knowledge sharing communities (e.g. communities of trust, or communities using similar problem solving strategies), and identify potential future members and items of interest. We are applying the Advogato trust metric (<http://www.advogato.org/trust-metric.html>) for dynamically propagating the values of our user model variables through our social networks. Over the next six months, we plan to conduct a formative evaluation in our client organization to assess these social networks' ability to raise users' awareness of the social factors that define their CoPs, and facilitate their access to relevant artifacts and communities of practice.

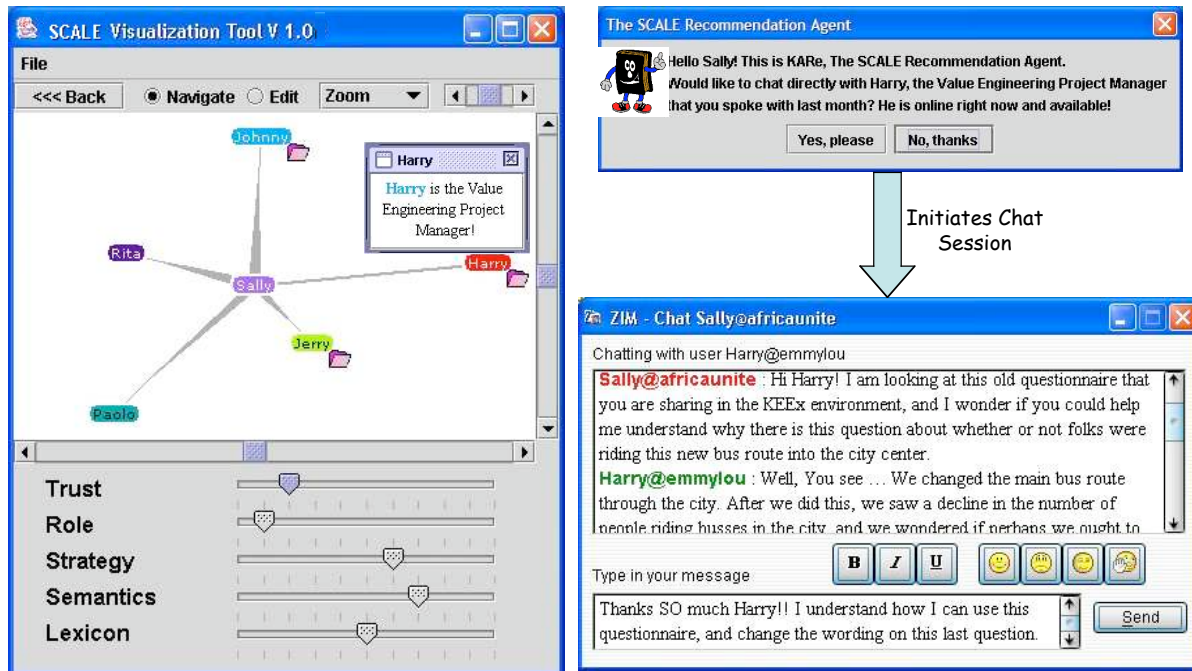


Figure 8. The IVisTo Visualization Tool⁽¹⁾ and KARE Recommender Agent

KARE: A Knowledgeable Agent for Recommendations

The KARE multi-agent system (Figure 8) is being designed, within the peer-to-peer knowledge management framework, to recommend both *documents* that match the user's knowledge and information needs, and relevant *peers* or *communities* that the user might query for information and knowledge (Guizzardi, Aroyo, & Wagner, 2003). KARE generates recommendations by reasoning over the values stored in the user models (such as the user's organizational role, problem solving strategy, level of trust, and availability), while considering the user's personal knowledge ontology in KEEEx, and current query. As in the I-Help system, (Vassileva et al., 2003), KARE guarantees the users' privacy by ensuring access to only information that is authorized by the user. The KARE privacy system is modeled using a multiagent approach, comprising Personal Peer Assistant Agents, and cooperating Broker Agents. Each user has exactly one Peer Assistant, stored on the user's machine, in accordance with the distributed peer-to-peer approach. The Peer Assistant is responsible for maintaining the user model, and performing rule-based inferencing operations (e.g. Linton, Goodman, Gaimari, Zarrella, & Ross, 2003). Recommendations are generated by reasoning over the values stored in the user models (e.g. organizational role, problem solving strategy, and availability), while considering the user's privacy settings, personal knowledge ontology in KEEEx, and current query. These potential recommendations are then send to the Broker Agents.

The multiple, cooperative KARE Broker agents are responsible for contacting suitable peers to answer knowledge requests, and recognizing patterns in users' models that might suggest the formation of communities of practice. Each Broker agent will provide a different type of community service, for example collecting the available user model information, analyzing interaction patterns, or assessing users' knowledge needs (Guizzardi, Aroyo, & Wagner, 2003; Linton et al., 2003). The Broker agents are intended to facilitate the natural processes of knowledge sharing across communities by providing users with referrals to other users holding relevant knowledge and residing in similar communities (Gaudioso & Boticario, 2003; Vassileva et al., 2003). We expect such services to be

especially useful for newcomers who do not yet have an established network of peers or an understanding of the knowledge and expertise that exists in the distributed environment.

The Integration Model

The SCALE integration model (Figure 9), defines the interactions between the KEEEx environment, the IVisTo visualization tool, and the KARE multi-agent system. The results of a user's query in KEEEx, including the lexical and ontological similarity data, are sent to the SCALE environment as an XML file. This file is then parsed by the User Model Engine (XML and UM Update arrows), combined with the available social and behavioral knowledge, and sent to both IVisTo and KARE. The User Model Engine plays a critical role in the integration model because it is responsible for guaranteeing the consistency of the user model information while both IVisTo and KARE update the user model based on interactions with the user.⁽²⁾

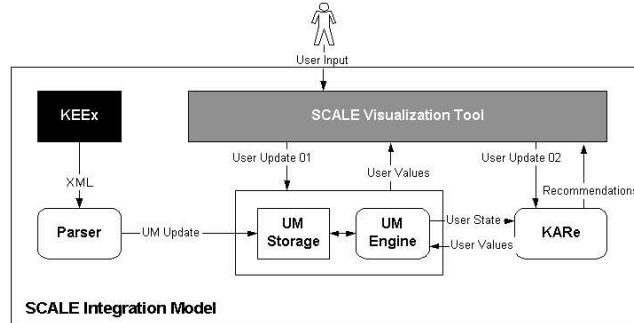


Figure 9. The SCALE Integration Model

Many of the technologies that comprise the SCALE environment have been applied independently in other contexts. For example, the aLF environment (Gaudioso & Boticario, 2003) performs user modeling to evaluate the users' level of activity, analyze their forum postings, and recommend sub-groups within existing communities. The system also provides navigation support to tailor the course content and recommend information sources. The success of Vassileva et al.'s I-HELP system (2003) provided the research motivation for the development of the KARE multi-agent system. I-HELP's personal and broker agents recommend content experts and artifacts to users, based on the characteristics of users distributed throughout the system. Ogata, Matsuura, and Yano's (2000) Knowledge Awareness Map is similar to IVisTo, in that it depicts a specialized social network based on the "knowledge pieces" held by each participant, and graphically shows users who else is discussing or manipulating their knowledge pieces. In this case, the distance between users and knowledge elements on the map indicates the degree to which users have similar knowledge. The SCALE approach is distinct from other approaches (with the exception of Bretzke & Vassileva, 2003) in that the environment is entirely distributed in nature, without any central databases or servers. It is designed to help users identify relevant communities of practice by not only modeling the learners' knowledge and behaviors, but also discovering similarities in their personal knowledge ontologies, which may differ in structure and language, and constructing user-controlled dynamic social network visualizations (IVisTo) built upon the foundation of real organizational practice.

Summary and Future Work

In this paper, we presented an example i* intentional analysis designed to identify the critical factors that influence knowledge sharing and learning across an organization, especially during practices in which newcomers join workgroups. These factors were realized as user model variables, and provided the foundation for the SCALE IVisTo interactive visualization tool, and KARE multi-agent recommendation system. These two tools are designed to promote community awareness, social interaction, knowledge sharing, and learning within the KEEEx open-ended peer-to-peer organizational learning environment. The KEEEx environment presents the results of the user's query as an ordered list, ranked by the accuracy of the keyword query results and contextual match. The IVisTo tool enhances this interface by providing the learner with an interactive social network that displays his query results in terms of social and behavioral factors, promoting the awareness of online knowledge communities. KARE further enhances this environment by suggesting community building actions to encourage the creation of new knowledge.

In the future, we plan to replace the XML-based user model with a description logic-based model (e.g. OWL) that affords more advanced reasoning capabilities. We also intend to continue to evaluate the business practices within our two target organizations so that we might better understand which user model variables play

major roles in their learning activities. Our developed system will be deployed within these organizations so that we might evaluate the SCALE system's ability to facilitate the learning process in which newcomers join workgroups, and in the future, other key organizational learning practices. This phase should result in a feedback loop, suggesting the foci of the next Intentional Analysis phase, and the refinement of the user models.

Endnotes

- (1) The IVisTo prototype was developed using Touchgraph_{LLM} [Available: <http://www.touchgraph.com>].
- (2) IVisTo (under development at ITC-IRST, Italy) and KARE (under development at the University of Twente, the Netherlands) are designed to also function as separate plug-in components.

References

- Ayala, G. (2001). Intelligent agents supporting the social construction of knowledge in a learning environment. In Q. Chen (Ed.) *Human Computer Interaction: Issues and Challenges* (pp. 44-63). Idea Group Publishing.
- Bonifacio, M., Cruel, R., Mameli, G., Nori, M. (2002) A peer-to-peer architecture for distributed knowledge management. *Proceedings of the 3rd International Symposium on Multi-Agent Systems, Large Complex Systems, and E-Businesses (MALCEB'2002)*.
- Brown, J. S., & Duguid, P. (1991). Organizational learning and communities of practice: Toward a unified view of working, learning, and innovation, *Organization Science*, 2(1), 40-57.
- Bouquet, P., Serafini, L., & Zanobini, S. (2003). Semantic coordination: A new approach and an application. *Proceedings of the 2nd International Semantic Web Conference*, Sanibel Island, FL, 130-145.
- Bretzke, H., & Vassileva, J. (2003). Motivating cooperation on peer to peer networks. *Proceedings of the 9th International Conference on User Modeling*, Johnstown, PA, 218-227.
- Gaudio E., & Boticario, J. (2003). Towards web-based adaptive learning communities. *Proceedings of Artificial Intelligence in Education 2003*, Sydney, Australia, 237-244
- Guizzardi, R., Aroyo, L., & Wagner, G. (2003). Agent-oriented knowledge management in learning environments: A peer-to-peer helpdesk case study. *AAAI Spring Symposium on Agent Mediated Knowledge Management*.
- Gutwin, C., Stark, G., & Greenberg, S. (1995). Support for Workspace Awareness in Educational Groupware. *Proceedings of CSCL'95*, Bloomington, IN, 147-156.
- Hoadley, C., & Pea, R. (2002). Finding the ties that bind: Tools in support of a knowledge-building community. In K. A. Renninger & W. Shumar (Eds.), *Building Virtual Communities: Learning and Change in Cyberspace* (pp. 321-354). New York: Cambridge University Press.
- Jermann, P., Soller, A., & Muehlenbrock, M. (2001). From mirroring to guiding: A review of state of the art technology for supporting collaborative learning. *Proceedings of the First European Conference on Computer-Supported Collaborative Learning*, Maastricht, The Netherlands, 324-331.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.
- Linton, F. Goodman, B., Gaimari, R., Zarella, J., & Ross, H. (2003). Student modeling for an intelligent agent in a collaborative learning environment. *Proceedings of User Modeling 2003*, Johnstown, PA, 342-351.
- Mantovani, G. (1996). Social context in HCI: A new framework for mental models, cooperation, and communication. *Cognitive Science*, 20, 237-269.
- Molani, A., Perini, A., Yu, E., Bresciani, P. (2003) Analyzing the requirements for knowledge management using intentional analysis. *Proceedings of the AAAI Spring Symposium on Agent Mediated Knowledge Management*, Stanford University, California, USA.
- Ogata, H., Matsuura, K., & Yano, Y. (2000). Active knowledge awareness map: Visualizing learners activities in a web based CSCL environment. *Proceedings of the International Workshop on New Technologies for Collaborative Learning (NTCL 2000)*, Awaji, Japan, 89-98.
- Reimann, P. (2003). How to support groups in learning: More than problem solving. Invited keynote, *Supplementary Proceedings of the 11th International Conference on Artificial Intelligence in Education*, 3-16.
- Scardamalia, M., & Bereiter, C. (1994). Computer support for knowledge-building communities. *Journal of the Learning Sciences*, 3(3), 265-283.
- Soller, A., & Lesgold, A. (in press). Modeling the process of knowledge sharing. In U. Hoppe, M. Ikeda, & H. Ogata (Eds.) *New Technologies for Collaborative Learning*. Kluwer Academic Publishers.
- Vassileva J., McCalla G., Greer J. (2003). Multi-agent multi-user modeling. *User Modeling and User-Adapted Interaction*, 13(1), 179-210.

Acknowledgments

Very special thanks to the members of the SCALE team, Matteo Bonifacio, Lora Aroyo, Peter Mika, and Paolo Massa, for their collaboration and suggestions, and also to Julita Vassileva for her comments and continual support.