Department of Informatics
School of Economics and Commercial Law
University of Gothenburg

Using an Agent-Based Recommender System to Support Competence Management –

The Case of Volvo Information Portal

Master Thesis 20 p
Spring 2001
IA7400

Abstract

There are several ways that organizations can support knowledge management (KM). Some are cognitive while others focus more on collaboration in communities. There are also a number of ways to design systems to support KM, but few of these deal with the tacit dimension of knowledge and competence. As several researchers have criticized existing KM systems for being too limited, this study focuses on a different approach, i.e. a technique that thus far had not been used for KM purposes. More specifically, we examined how an agent-based recommender system could be used as a KM system, focusing on competence in use. Based on an extensive literature research, this case study was performed at the Volvo Corporation, where our unit of analysis was designed as a portal on the corporate intranet. The study included an evaluation of the system Volvo Information Portal (VIP). This evaluation was founded on how the VIP system could support the organizational level as well as the individual co-worker. The main results of this study are: First, several fields of application, i.e. how VIP can be used to support competence management. Secondly, a number of design implications, i.e. improvements that would enhance VIP in this aspect.

Authors	Supervisor	Examinator
Annika Hanefors Christer Undemar	Rikard Lindgren	Urban Nuldén

Acknowledgements

The work with this thesis began in January 2000 and our plan was to submit it for examination in May that year. However, due to a number of circumstances beyond our control the work was delayed. We were forced to start all over and the final phase was postponed several times. When we needed a new principal, Volvo Information Technology invited us to conduct our study there, for which we are very grateful. We would especially like to thank senior information architect and researcher Dick Stenmark who, in spite of a very busy schedule, took time to explain VIP and took part in the work with the analysis of the interviews.

The sixteen test users, who graciously agreed to be interviewed, also need to be recognized, and among them especially HR manager Tiina Hyvönen, who, besides answering our questions, took care of all the practical details regarding our time at Volvo.

We had the privilege to cooperate with the Viktoria Institute and our thesis became a part of their KM project. The participants in this project shared valuable comments with us, during the two seminars that we attended. One of the Viktoria researchers, Ph. D. student Rikard Lindgren, became our supervisor. We would like to express our sincere appreciation for his commitment to help us through our prolonged project, especially after August 2000 when we both started full time employment and were forced to schedule our appointments to early mornings before work. We are deeply indebted to him for his encouragement, support, and invaluable input.

Gothenburg 2001-05-18

Annika Hanefors & Christer Undemar

Table of Contents

1 Introduction

The importance of knowledge management (KM) increases continually as organizations become more and more knowledge-based and dependent on the competence of their co-workers. However, it is difficult to find effective solutions for KM, since knowledge and competence often are hidden in the mind and subconscious of people, and when they leave the organization the loss of expertise can be severe. One common way to address this issue is to store information about knowledge and competence in databases and KM systems in order to transform the information into organizational assets (Hansen et al., 1999). A necessary prerequisite for such systems is that the information is explicit and codifiable. In addition, the information must fit predefined categories and be possible to grade. One problem with this cognitive approach is that knowledge and competence can be tacit, and therefore not easily transformed into explicit information. This awareness provides a platform for another approach, the community approach, which has a more social perspective on KM (Hansen et al., 1999). Here knowledge is considered socially constructed (Bannon & Kuutti, 1996) and tied to the individual. It is created and shared through the interaction between people. In this approach information technology (IT) is used to aid communication, not for storage of knowledge and competence, and to support the building of communities (Robertson et al., 2000; Swan et al., 2000).

Historically, KM is related to the evolution of IT-based systems (Swan et al., 1999). The closely related research area organizational memory (OM) has also tried to find IT solutions to support KM. The purpose of OM systems is to support organizations and their employees in their efforts to capture and retrieve experiences, to find and interact directly with experts, and through that collaborate more effectively (McDonald & Ackerman, 2000). Technologies used for this are for example groupware packages, hypertext systems, and intranets (Snis, 1999). Other technologies for KM include repositories of knowledge and search tools that make it possible to retrieve stored knowledge objects. There are also a number of applications that aim to capture and store competence, i.e. organizational groups of people with a certain expertise, by creating a competence structure with roles and categories (Lindgren & Wallström, 2000).

Research has shown that existing KM systems are not being fully utilized (Fahey & Prusak, 1998 cf. Lindgren & Wallström, 2000). There are few users of the systems due to lack of time and reluctance to contribute to updating the systems, which conveys that the information might not be up to date and sufficient. Another problem is that hierarchical design limits the use, i.e. they are often designed for the use of management alone (Lindgren & Wallström, 2000). Systems used within OM have been criticized for being limited when it comes to supporting human problem solving (Davenport, 1996). KM has experienced similar critique and limitations (McDermott, 1999; Scarborough, 1998; Swan et al., 1999) and it is more and more obvious within these research areas that they need to focus less on IT infrastructure only, but rely more on organizational issues such as the collective knowledge (Sarvary, 1999). Consequently, there is a need for new perspectives on KM to support the individuals who build this collective knowledge. However, both the cognitive and the community approaches focus on knowledge and competence defined by management and used by

the organization as a whole. Little research has been conducted on how to support the individual.

Our view of knowledge and competence, and consequently on KM, is influenced by Habermas' (1986) theory on knowledge-constitutive interests. He states that there is an evident relationship between interest and knowledge, and that our interests inevitably control us. Humans perceive reality based on their interests and on how they see themselves in relation to others. By paying attention to interests the more elusive tacit knowledge can also be supported, while a too heavy focus on well-defined, concrete, and graded expertise can lead to its loss. Furthermore, knowledge that is not considered core knowledge, but still is important for both the individual and the organization, can be encouraged through attention on interests. We found this relationship between knowledge and interest to be of such importance that we adopted it as one of the propositions of our thesis, especially after extending it to also include competence, i.e. knowledge put into action. Studying Habermas' theory led us to draw the conclusion that there are more areas than explicit knowledge that need to be supported by KM systems. Our second proposition is consequently that traditional KM systems are insufficient in fulfilling this purpose. Therefore, research on technologies previously not used within KM need to be performed to find out how they can serve as a complement.

Thanks to the expansion of the Internet, there is a platform for a number of new techniques supporting both individual and organizational interests, and facilitating networking over as well organizational as geographical borders. As a step in aiding the users to navigate the information domains, different tools, e.g. search engines, recommender systems (RS), and agent-based systems, have been developed. These are not traditional techniques used for KM, but lately some KM researchers have started to take an interest in them. Stenmark (2001) studied how tacit knowledge can be visualized with the help of agent technology. Other examples are research on how RS (McDonald & Ackerman, 2000) and software agents (Vivacqua, 1999) can be used to locate expertise.

In this thesis, we examine how an agent-based RS can be used as a KM system, focusing on competence in use. The study focuses on support for the organization as a whole, as well as the individual co-worker. Part of our study was performed at Volvo where we got access to a system, i.e. Volvo Information Portal (VIP), on the corporation's intranet, i.e. Violin. VIP came to serve as a platform for the empirical study, which together with the chosen theory yielded the following questions:

1. How can VIP be used to support management of competence?
2. Which changes are needed in order to improve its functions in this aspect?

The main objective of the thesis is to contribute with ideas on how an agent-based RS can support the management of competence. We propose that interest also needs to be supported since it is intimately interwoven with knowledge and competence. We draw this conclusion after merging the ideas and research of Habermas (1986) and Stenmark (2001).

This introductory section is followed by a description of our course of action and the method we have chosen to use (section 2). Next, we describe the theories we have applied (sections 3 and 4). In section 5, we account for VIP, i.e. the system on which we based our empirical study. Furthermore, we give a brief description of TP/HR, which is a KM system that was implemented at Volvo at the time for our study. Then follow the results of the interviews (section 6). Finally, we discuss our findings (section 7) and draw conclusions (section 8).

2 Method

Our method consists of two main parts, a theoretical and an empirical study. They are closely related, and both parts were conducted in parallel with each other. From the material gathered in these studies we performed the analysis and drew the conclusions accounted for in the end of this thesis.

2.1 Course of Action

We based our study on a hermeneutic worldview and used a qualitative method, i.e. case study, to perform our research. Our case study consists of six parts (fig. 1): 1) Choice of research area; 2) Literature study; 3) Evaluation of VIP; 4) Seminars on KM; 5) Interviews; and 6) Analysis of our findings.

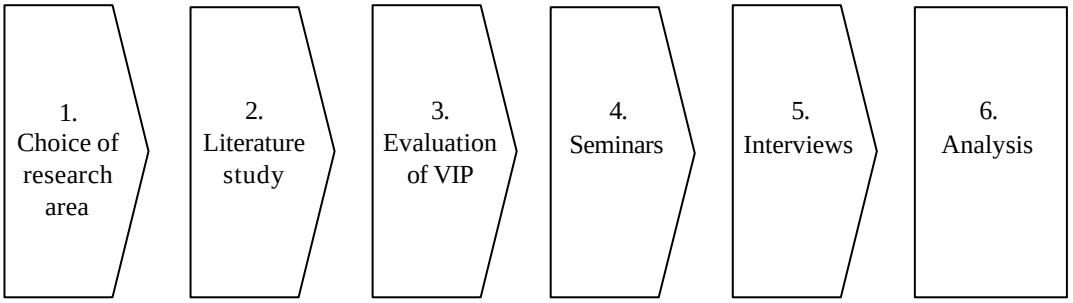

Figure 1: Course of Action.

2.2 Case Study

We chose case study from an abundance of scientific methods. Each method is a help to answer questions about a certain phenomena. These can be questions like 'what' and 'why' (Asplund, 1970). According to Easterby-Smith, Thorpe and Lowe (1991), a scientific method constitutes the overall configuration of the research and helps to recognize useful designs, but also to identify designs outside the researcher's past experience. Methods can be compared to sunglasses in different shades that give the same view different appearances. Thus, the choice of method is the choice of perspective from which to attack a problem.

According to Dahlbom and Mathiassen (1995) all scientific methods have roots in the two main historic worldviews, i.e. positivistic and heuristic. Easterby-Smith et al. (1991) also discusses these two traditions, but calls them objectivist (positivist) and subjectivist (hermeneutic). From the objectivist philosophy stems the quantitative method, which is based on the traditional assumption of a more or less objective reality, separated from mankind (Backman, 1998; Habermas, 1986). Objects, conditions and events exist independent of human beings. This naturally affects research. One observes, registers and measures a given reality. The individual is separated from the surrounding world. In order to explain this objective reality the

researcher formulates theories about it and deduces hypotheses to find out if they can be verified or falsified.

The hermeneutic philosophy, which we adopted for our study, sees the world as an individual, social, and cultural construction (Backman 1998; Easterby-Smith et al, 1991), and it is impossible to separate knowledge from knower (Alvesson & Sköldberg, 1994). The observer is part of what is observed and human interests drive science. The researcher should focus on meanings, try to understand what is happening, and develop ideas through induction of data. Here, a multiple choice of methods is preferred to establish different views, and small samples should be investigated deeply and over time. The central question in a qualitative method, e.g. case study, is how individuals experience, interpret, and structure a surrounding reality in relation to earlier knowledge and experiences, and data is considered a construction or a result of interpretation (Backman, 1998). Usually the individual is studied in real life situations. Processes rather than products and results characterize the qualitative perspective. The researcher is close to the studied subject and is sometimes part of the method. This is an inductive method, i.e. the research begins with the collection of data and continues with the formulation of hypotheses of theories.

Within the qualitative research there are courses and directions ranging from grounded theory and phenomenology to poststructuralism, postmodernism, and even feminism (Alvesson & Sköldberg, 1994). Common for all of these is the empirical study of a reality full of contradictions, and the focus on the lingual, interpreting, and selective part of research. Another significant method within the qualitative research is the method we chose for our thesis, i.e. case study:

> *"A case study is an empirical inquiry that investigates a contemporary phenomenon within its real-life context when the boundaries between phenomenon and context are not clearly evident and in which multiple sources of evidence are used."*
>
> (Yin, 1988).

Yin (1988) states that case studies are to be preferred when the research focuses on 'how' and 'why' questions, when the investigator has little control over events, and when the focus is on a contemporary phenomenon within some real-life context. This type of research is ideal when there is a need to understand complex social phenomena, but also when the relevant behaviors cannot be manipulated. Even though the case study involves techniques similar to those in other methods, such as the study of documents and artifacts, it also relies heavily on systematic interviews and direct observations. According to Yin, this ability to deal with multiple sources of evidence is the strength of the case study.

There are several applications of case studies. It can explain causal links, describe real-life contexts, illustrate, or explore situations. Our study can be described as an exploratory case study, since we are exploring how a certain technology can be used in an unconventional setting and for a new purpose. Typical for an exploratory case study is that there is no clear single set of outcome, and this is also evident in our study.

Yin (1988) also mentions five important components of research design for a case study: the study's question, propositions, units of analysis, logic linking of the data to the propositions, and criteria for interpreting the findings. The question is most commonly a 'how' or a 'why' question, which also is evident in our first question: How can VIP be used to support management of competence? In the introduction we mentioned our propositions, even though Yin states that this is not necessary in an exploratory case study. We will describe the unit of analysis, i.e. VIP, in section 5. In section 7 we link the empirical findings to the propositions. The criteria for interpreting the findings are described in subsection 2.2.6. Next, we will describe the different parts of our case study.

2.2.1 Choice of Research Area

We were invited to cooperate with the Viktoria Institute (Viktoriainstitutet, 2001) and Volvo Information Technology (Volvo IT, 2001). The latter also provided the environment for our empirical study. Our study became a small part of the KM project at the Viktoria Institute, in which Volvo IT is one of several participating organizations (KM Project Participants, 2001). We found research performed by Rikard Lindgren, Christoffer Wallström and Dick Stenmark within the KM project interesting. Lindgren and Wallström (2000) conducted a study on different KM systems for managing competence and the deficiencies found in them. Stenmark (2001) researched on how to turn tacit knowledge tangible. He means that agent-based retrieval systems can be used to capture and visualize professional interests, thus making otherwise elusive tacit knowledge tangible for others to benefit from. This created an interest in us whether there are technologies and approaches new to KM that can be used for such a purpose, and we decided to choose this as our research area.

2.2.2 Literature Study

Theory can be important to case studies in several ways (Yin, 1993). This kind of study may for example help with case selection, specification of what is being explored, and generalization of the results to other cases. Therefore we started with a comprehensive study of literature on knowledge, competence, interest, knowledge management, KM systems, intelligent agents, and RS. This theoretical study was the starting point for our research but it also continued as an ongoing process during the entire work with the thesis. The studied literature inspired us to choose our research area and helped to narrow it down, but it also gave us a structure, against which we could map the results of the empirical study.

2.2.3 Evaluation of VIP

As stated above, a case study relies on multiple sources of evidence (Yin, 1988). Another characteristic is that the researcher is not an objective outsider, but observes from within and interprets based on previous understanding. We have already described literature as one of our sources. The others are our own evaluation of VIP,

the seminars, and the interviews. The first thing we did as we got access to the Volvo corporate intranet, i.e. Violin, and VIP was to learn as much as possible about, and do our own evaluation of, the latter. We studied the help-function and read all documentation available. Then we thoroughly walked through all functions several times and discussed our findings both with each other and with Dick Stenmark, to whom we also put all questions that came to mind. We also created several channels each and tried out the different functions connected to those.

2.2.4 Seminars

In April and June 2000 we participated in two seminars arranged by the KM project. The first seminar took place at Volvo and the second at the Viktoria Institute. At both these seminars we had the opportunity to give presentations about the work with our thesis. Several members of the KM project from different organizations were present at the seminars and we received valuable input from them.

2.2.5 Interviews

In May and June 2000 we interviewed 16 employees at Volvo. The interviewees were employed at the following corporate divisions: Volvo Cars, Volvo Trucks, Volvo Penta, and Volvo IT. Their job descriptions varied from project managers to info masters and systems developers (appendix 1). They were all part of the test group and had had access to the system for different periods of time, lasting between one week and several months, prior to the interviews. Each interview lasted between 40 and 75 minutes and took place at either their or our place of work.

During a lecture, Bergqvist (1999) gave suggestions for designing interviews. One was that one constantly should proceed from the purpose of the study when choosing methods, interviewees, and delimitation of quality. The latter depends on how well the researcher succeeds in collecting, processing and presenting the material, as well as planning his time. Jones (1985) means that while preparing the interviews one will, and should, have a few broad questions in mind. One needs a framework from which to proceed, but must at the same time not be too restricted by it. In that way one can follow all interesting tracks that the interviews take. This type of interview is called semi-structured and is the technique we chose to work with.

The objective of our interviews was to get an idea of how the respondents perceived VIP, which functions they had used, which fields of applications they saw, and what features they considered missing (appendix 2). We also asked questions about the interviewees' backgrounds, how they search for information, and their experiences of the Internet and intranets. In the second part of the interview we asked the interviewees to log on to VIP and then, with the system as a background, we asked more specific questions about it. This served two purposes: as we realized it might be difficult to remember all features and functions of the system we wanted this part of the session to be a reminder of them, and secondly we believed this would provide additional nourishment to their thoughts and reasoning.

As mentioned above, we chose to use in depth interviews performed in a semi-structured way (Easterby-Smith, et al., 1991), i.e. we asked open questions and followed up with more questions to make sure that we understood what the interviewees wished to express. We did not want to limit the interviewees by asking too structured questions, but sought to encourage their free reasoning. As we prepared the interviews we tried to start with easy-to-answer, non-threatening questions, followed by broad questions about the system allowing the interviewees' thoughts to take different tracks. This also helped us to stay un-biased. The interviews were recorded on a mini-disc recorder and in addition notes were taken.

2.2.6 Analysis

According to Yin (1988), there are no fixed formulas on how to perform the analysis of a case study, but much depends on the researcher's own style of thinking. One approach could be statistic analysis by coding events into numerical form. Another is to use different analytic techniques, e.g. putting information into different arrays, putting evidence within categories in a matrix, and tabulating the frequency of different events. Such analysis must be done carefully to avoid bias, since the goal is to treat the evidence fairly, to produce compelling analytic conclusions, and to rule out alternative interpretations. Yin also stresses the importance of having a general strategy for the analysis. We chose the most preferred strategy: relying on theoretical propositions. Such propositions reflect a set of research questions, reviews of literature, and new insights. Our propositions helped us to shape the data collection and focus on certain data during the analysis.

Easterby-Smith et al. (1991) suggests a method for analyzing in-depth interviews. In this theory the researcher goes by feel and intuition, aiming to produce common or contradictory themes and patterns from the data, which can be used as a basis for interpretation. In contrast to quantitative research, the structure used for the analysis first has to be derived from the data, which means systematic analysis in order to find themes, patterns, and categories. As we analyzed and interpreted the material we went through the following seven stages mentioned by Easterby-Smith et al.:

1. *Familiarization*. In June 2000 we started our analysis by transcribing and studying the interviews. We read the transcripts several times and used brainstorming to find interesting things. During this stage we tried to stay impartial, but interesting discoveries in one transcript lead us to look for similar thoughts in other transcripts.
2. *Reflection*. As we had an extensive amount of material, we tried to categorize it to make it easier to handle. We also turned to our supervisor Rikard Lindgren and researcher Dick Stenmark for valuable input.
3. *Conceptualization*. At this stage a number of concepts emerged. In order to secure their relevance we went back to the transcripts to mark their appearances. During this process we had to redefine some of the concepts and some were disregarded altogether. In the end, we had fourteen concepts, seven fields of application and seven desired improvements. These are further described in section 6.
4. *Cataloguing concepts*. We labeled the concepts and marked the names of the interviewees, who had said something about them.

5. *Recording*. Next we went back to the transcripts and studied more carefully what was actually said. Once more we had to give some thoughts to whether or not the concepts needed to be redefined and recoded.
6. *Linking*. At this stage we began to link together all the identified variables, in order to get a more holistic perception. We mapped the results from the empirical study with our chosen theories.
7. *Re-evaluation*. Finally we gave drafts to our supervisor who commented on and criticized them. After having received his input we rewrote the drafts and gave them back for further comments.

Many of the stages mentioned above were undertaken several times. During the analysis we also finally defined our study questions, a procedure that is common for an exploratory case study (Yin, 1993). In this type of study, fieldwork and data collection are conducted first. Such research may be perceived as intuitive, but the purpose is, according to Glaser & Strauss (quoted in Yin, 1993), often to make discoveries by directly observing a social phenomenon in its raw form. Therefore, only the broad features of the study design are determined in advance.

2.2.7 Validity and Reliability

Validity and reliability was originally used in quantitative science, and in this approach there are a number of different methods to assess both (Easterby-Smith et al., 1991). These methods might not be as easy to use within qualitative research, since the hermeneutic philosophy does not view the world as absolute and objective. Hence, it can be difficult to determine whether the used instruments succeed in measuring what they are supposed to measure, i.e. validity, or gives a reliable result, i.e. reliability (Wiedersheim-Paul & Eriksson, 1997). However, Easterby-Smith et al. (1991) mean that the concepts can be applied in qualitative research, provided that the researcher is committed to providing a faithful description of others' understandings and perceptions. To determine validity in a qualitative study they suggest the question: "*Has the researcher gained full access to the knowledge and meanings of informants?*" The corresponding question for reliability is: "*Will similar observations be made by different researchers on different occasions?*"

Naturally, there is no way for us to answer these questions with complete accuracy. However, we tried, to the best of our ability, to meet the interviewees with open minds and listen to their responses without prejudice. Both of us were always present during the interviews, one of us responsible for questioning and the other taking notes. Since both of us also were free to follow up with new questions whenever something was unclear, we believe that we give a fair representation of their views. Furthermore, in this kind of research there is always a risk of bias, since it depends on the researcher's view and interpretation of reality. We were aware of this and tried to maintain objectivity during the interviews and not ask leading questions. Our belief is that we succeeded as well in this effort as any other researcher would have done. If anyone else had performed the same observations at this time we believe that they would have reached similar results.

2.2.8 Further Critical Observations

There are a number of factors that influence the outcome of our study. For instance, the interviewees had only tried the system for a limited amount of time. They were all part of the small group that had tested the system, from which a majority agreed to be interviewed. If the system had had more users we could have made a random choice of interviewees. We were also restricted to schedule the interviews in May and June, since they had to be completed before the summer vacations. Repeated interviews over a longer period might have yielded different results. We are also aware of the fact that this is a subjective study, due to its qualitative character. Therefore the results will be somewhat colored of our opinions and interpretations, no matter how objective we try to be. This is the very nature of a case study. Finally, we admit that the terms knowledge and competence may cause confusion, due to their similarities and close relationship. We have tried to remedy this predicament by the definitions made in the next section, and we also did our best to distinguish between them during the interviews.

3 Management of Knowledge, Competence, and Interest

Knowledge is an ancient concept that has been given many definitions over the centuries (Encyclopædia Britannica, 2000). Plato stated, for example, that knowledge is justified true belief, i.e. in order to be knowledge, a statement must be true, and in addition, individuals have to believe that it is true. Aristotle meant that actual knowledge is identical with its object, and Descartes distinguished two sources of knowledge, i.e. intuition and deduction, where intuition is an apprehension of something experienced and deduction depends upon thought or reason. Nowadays, various researchers still elaborate with different distinctions of knowledge (Bertels & Savage, 1998). Frequently used categorizations include explicit and tacit (e.g. Polanyi), embodied (e.g. Zuboff), encoded (e.g. Zuboff), embrained (e.g. Blackler), embedded (e.g. Berger and Luckman), and procedural knowledge (e.g. Zander and Kogut).

The different distinctions above tell us something about the complexity of knowledge and the difficulty of finding an all-embracing definition. However, for the scope of our thesis we will concentrate on the categorization made by Polanyi (1966), i.e. explicit and tacit knowledge. Explicit knowledge is formal and systematic. Therefore it can be easily communicated and shared, on product specifications or a scientific formula or a computer program. Tacit knowledge on the other hand is highly personal. It is hard to formalize and therefore, difficult to communicate to others. As Polanyi (quoted in Nonaka, 1994) says: "*we know more than we can tell*". Tacit knowledge is deeply rooted in action and in an individual's commitment to a specific context. It is partly made up of technical skills, but at the same time it has an important cognitive dimension. It consists of mental models, beliefs, and perspectives so ingrained that we take them for granted, and therefore cannot easily articulate them.

The concept of competence is closely related to knowledge. To know means to be aware of, familiar, or aquatinted with something (Encyclopædia Britannica, 2000). To be competent means to have requisite skills, necessary qualifications, capabilities, power, and eligibility (Stenmark, 2001), i.e. to be able to put knowledge into action. However, similar to knowledge, competence is also discussed in different terms (Bertels & Savage, 1998), e.g. core competencies (e.g. Prahalad and Hamel), core capabilities (e.g. Zander and Kogut) and skills (e.g. Aaker). We have deliberately chosen to ignore the term competencies (sing. compentency) since it would cause confusion to use two terms with the same meaning. Instead, we refer to competence as group related expertise found in organizational settings. An individual who is highly skilled within a certain area, and have experience of applying the skill in complicated work tasks, is viewed as an expert. Groups of such experts are valuable to organizations, which continuously look for them.

3.1 Knowledge Management

There are several definitions of knowledge management (KM) given by e.g. Nonaka and Takeuchi (1995), Marshall and Prusak (1996), Sveiby (1997), and Davenport (1996). Even if they differ they point to some common purposes of KM:

- to *create* knowledge
- to *capture* knowledge
- to s*hare* and *recycle* knowledge
- to r*educe risks* of losing valuable knowledge
- to *create value* from knowledge

It is crucial for organizations to learn how to manage not only knowledge, but also competence, and KM has come to also include this. The management of competence includes internal marketing of expertise and, from a top-down approach, strategic management and mapping of competence (Lindgren & Wallström, 2000). Management of knowledge, i.e. know-what, and competence, i.e. know-how, is closely related and sometimes difficult to separate. We view the management of competence as a part of KM, and the emphasis of our thesis lays on this part. Therefore, when we refer to KM we also include the management of competence, but sometimes we will also distinguish between them when the discussed issue refers specifically to the management of competence (fig. 2).

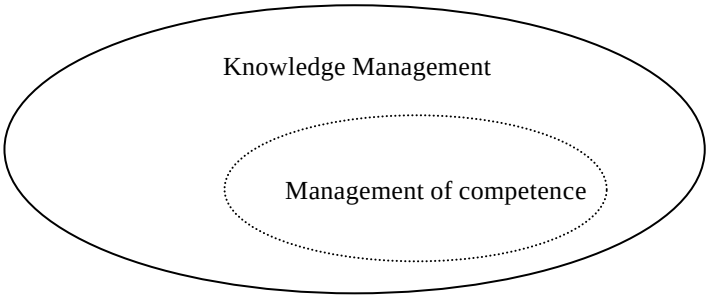

Figure 2: Our View of Knowledge Management (1).

KM is a research area that is closely related to organizational memory (OM). Both of these stem from computer supported cooperative work (CSCW), and they have slightly different views on how design and development of IT support for KM ought to be conducted. The KM community has struggled with many different scopes within KM systems, e.g. codification and personalization; and generation, codification and transfer (Lindgren & Wallström, 2000). Next we will describe two major perspectives on KM, i.e. the cognitive and the community approach. These approaches emphasize different concepts, e.g. exploitation and exploration of knowledge (Swan et al., 2000), codification and personalization (Hansen et al., 1999), and either one or the other of Polanyi's categories of knowledge mentioned above (figure 3, subsections 3.1.1 and 3.1.2). The approaches may have different ways of viewing knowledge and KM, but even so they cannot be totally separated from each other.

Cognitive approach	Community approach
Explicit knowledge	Tacit knowledge
Knowledge exploitation	Knowledge exploration
Codification	Personalization

Figure 3: The Cognitive and the Community Approach to Knowledge Management.

In our thesis, we will mainly focus on the community approach of KM, but in order to comprehend this concept, a basic understanding of what the cognitive approach represents is needed.

3.1.1 The Cognitive Approach

The cognitive approach to KM aims at capturing and transferring existing knowledge. Behind this approach lies a cognitive, information processing view of the firm where valuable knowledge located inside peoples' heads is identified, captured, and processed via the use of IT tools so that it can be applied in new contexts (Swan et al., 2000). One objective is to capture the individual's knowledge and make it the organization's asset, to avoid loosing it if employees leave. This approach mainly uses a codification strategy, i.e. centers the KM strategy on the computer (Hansen et al., 1999). This is also evident when it comes to the management of competence, where one aims at classifying and structuring employees' expertise and storing them in IT-based systems in order for the organization to be able to find the right competence at the right time (Lindgren & Wallström, 2000). However, even if this IT driven approach supports capturing and sharing of knowledge, there are also several fundamental problems connected with it (Swan et al., 2000):

- There is an underlying assumption that most relevant knowledge in an organization can be made explicit and codified. However, tacit knowledge is difficult to articulate or transfer in explicit forms because it is personal and context-specific. Therefore this approach is severely limited in terms of the contribution to innovation, since it focuses on transferring only explicit forms of knowledge.
- This approach focuses more on exploitation than on exploration. IT-based tools can support processing of existing knowledge but this is only a part of KM. Most of the emphasis is on increasing efficiency by exploitation rather than on encouraging more explorative processes.
- It is a supply driven approach, i.e. one presumes that if information is widely available it will be applied in new ways to develop innovation. However, even if knowledge is codified and stored, and individuals are invited to take part of it, there is no assurance that they will use or apply it. With a vast amount of available information the risk of overload is impending. This critique is shared by Davenport (1996) who states that sources of informal documents suffers from that such knowledge can neither be used for automatic problem solving, e.g. in expert systems, nor processed by complex query answering mechanisms, e.g. databases.

Hence, the ability to support human problem solving through informal knowledge is limited.

- A typical failure with this approach is ignoring the pre-existing organizational structures, norms and cultural values that lead different groups to have divergent, possibly even irreconcilable, interpretations of what needs to be done and how best to do it.

The research field of OM has also received similar critique, for its too cognitive approach, when trying to solve KM issues. Ackerman and Halverson (1999) state that it is not sufficiently founded on studies within an organizational field setting, i.e. within a context of everyday use. IT development should be based on empirical insights rather than analyses of prototype systems, which are largely focused on technology designed to replace human and paper-based memory. They mean that artificial memory is an artifact that holds its state, but at the same time is embedded in organizational and individual processes and thereby cannot be separated from them. Bannon and Kuutti (1996) also express a wish to head in a more community-based direction.

3.1.2 The Community Approach

While the cognitive approach focuses on the use of IT-based solutions to handle existing knowledge, the community approach has a more social angle. This approach emphasizes dialogue occurring through networks, which can, but do not have to, be IT enabled. Humans always look for a good informal place to communicate in, and this is also true in the virtual world, or as Prusak (quoted in Swan et al., 1999) puts it: *"If the water cooler was a font of useful knowledge, what constitutes a virtual one?"* The community approach mainly uses the personalization strategy, i.e. knowledge is viewed as closely tied to the person who developed it and is shared mainly through direct person-to-person contacts (Hansen et al., 1999). It originates from Japan (Cohen, 1998), where many practitioners focus on developing conditions that favor the exchange of tacit knowledge between individuals through social processes, i.e. knowledge exploration.

The chief purpose of computers in this approach is to help people communicate knowledge, not to store it, and to connect people so that they can think together and turn information into solutions through actions (McDermott, 1999). The latter is enabled when people are encouraged to form communities, which essentially are social collections of individuals who communicate with each other. Knowledge and expertise in the communities are continuously recreated through dynamic, interactive and social networking. The aim is to leverage knowledge and expertise by focusing on the community that owns it and the people who use it, rather than the knowledge itself. McDermott also states that the underlying assumption in this approach is that people learn more from each other than from themselves, and consequently, this approach highlights the importance of relationships, shared understandings, and attitudes of knowledge formation and sharing. Much of the learning process involves participation in communities, and through this process people come to embody ideas, perspectives, prejudices, language, and practices of that community. The knowledge of a community member is shared with others as the participants see the logic of each

other's thinking or knowledge creation process. All contacts, received or transmitted through our senses, can be vehicles for this sharing of knowledge.

Critique against this approach includes disappointment in existing IT systems, designed to leverage knowledge to individuals and collaboration groups, which have not fulfilled their purposes. McDermott (1999) means that people mostly send email to other people they already work with and that virtual teams need to build relationships face-to-face before they can begin collaborating. Hence the systems fall short in supporting collaborative work and knowledge creation. New solutions for the community approach include groupware programs for managing knowledge and expertise, which mainly originate from OM efforts to support KM (Snis, 1999). These programs support communication, collaboration and coordination between members of a community.

The debate on how to approach KM has shifted from a cognitive, decision-making process, to a more community-oriented focus on organizational knowledge and culture (Sarvary, 1999). Scarborough (1998) states that the emergence of the community approach weighs up some shortcomings of the cognitive approach but that it still needs to mature before we can see some real KM systems within it.

3.1.3 A Different Approach

Neither the cognitive nor the community approach ascribes any significance to interest, but we found the theory of Habermas (1986) on knowledge-constitutive interests very appropriate for our thesis. Jürgen Habermas is a well-known philosopher whose research is referred to in many sciences, even though not usually in the area of IT. He has written a retrospect of Kant and Fichte on reason and interest that reflects which types of interests that build up knowledge. He defines these as knowledge-constitutive interests, and we find this theory essential for the understanding of how and why knowledge and interest relate to each other.

Habermas tells us that interest in general is the pleasure that we connect with the idea of the existence of an object or an action. The basic conditions of life have an interest structure, and interest aims at existence because it expresses a relationship between the object of interest and our faculty of desire. The interest either presupposes a need or produces one. This has to do with the distinction between empirical and pure interest introduced by Kant. Interest of the senses in what is pleasant or useful arises from need while interest of reason in the good awakens the need. In the former case the faculty of desire is stimulated by inclination and in the latter it is determined by principles of reason.

Habermas defines the term interest as the basic foundation for work and interaction, i.e. the specific fundamental conditions of the possible reproduction and self-constitution of the human species. Work and interaction include processes of learning and arriving at a mutual understanding. These processes have to be maintained if the self-formative process of the species is not to be socially endangered. Inspired by Kant and Marx he means that the experience of the emancipatory power of reflection is essential. This experience articulates itself in the concept of a self-formative process. In self-reflection then knowledge for the sake of knowledge coincides with

the interest in autonomy and responsibility. Habermas also borrows thoughts from Fichte, saying that self-reflection leads to that the ego frees itself from dogmatism and that the moral quality of a will to emancipation is required for the ego to raise itself to intellectual intuition.

While the pursuit of reflection knows itself as a movement of emancipation, reason is subject to the interest of reason. Habermas states that reason that dictates different types of interests is not pure practical reason, but reason that combines knowledge and interest in self-reflection. Similarly, he means that the interests directed toward communicative and instrumental action necessarily include relevant categories of knowledge. The knowledge-constitutive interests cannot be established permanently unless pertinent categories of knowledge, i.e. cumulative learning processes and permanent interpretations transmitted by tradition, are secured.

Habermas argues that the interest in preservation of social life is rooted in life organized through knowledge and action. On one hand they attest the fact that the cognitive processes arise from life structures and functions within them. On the other hand, they also signify that the form of socially reproduced life cannot be characterized without recourse to knowing and acting. Interest is attached to actions that both establish the conditions of possible knowledge and depend on cognitive processes. The interest of reason is necessary, and can not corrupt reason's cognitive power, when knowing and acting are fused into a single act, i.e. when individuals use their competence. Nevertheless, interest is still a part of knowledge, even if the knowledge is not acted upon. However, the knowledge-constitutive interests demand that we have entered the dimension of self-reflection, since it is in accomplishing self-reflection that reason grasps itself as interested. From this theorizing Habermas has drawn a fundamental conclusion: That objectivism, i.e. the objectivistic self-understanding of the sciences, which suppress every contribution of subjective activity, is dissolving since knowledge and subjective interest is so closely related.

Stenmark (2001) also discusses the relationship between knowledge and interest. He means that interest is an instance of tacit knowledge. Even if interests can be difficult to define, an individual usually has no problem determining if something is interesting or not. We consider the knowledge-constitutive interests as tacit and hidden within the individual. If an organization tries to leverage these interests from individuals to groups of people, in the same way as they do with expertise and competence, we talk about interest structures. We believe that the ability to visualize these knowledge-constitutive interests and interest structures is of great concern if we want to benefit from the tacit dimension of knowledge and competence. Hence, we believe that the management of interest needs to be considered as a part of KM, in the same way as we view the management of competence (fig. 4).

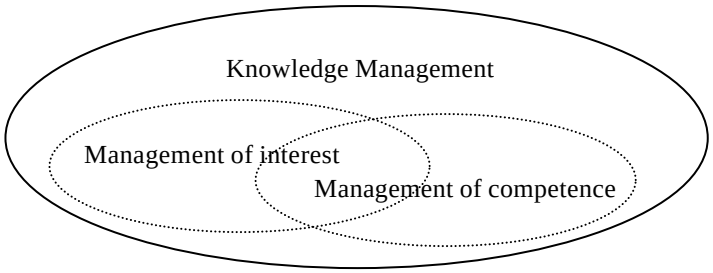

Figure 4: Our view of Knowledge Management (2).

Next, we will present different perspectives on the use of IT tools to support all these aspects of KM.

3.2 IT Support for KM

Commercial KM systems have hitherto been developed mainly with an accentuation on the cognitive approach but some have elements, which could be related to the community approach (Scarborough, 1998). Historically, KM systems for management of both knowledge and competence have been developed from a mechanical perspective. Lately a need for a softer approach has emerged. However, regardless of which approach, the systems mostly lack the ability to support the tacit dimension, why we suggest that research on new technologies is needed.

3.2.1 A Mechanical Perspective

The mechanical perspective on KM has mainly focused on careful codification and storing of knowledge in databases, where it can be easily accessed and used by anyone in the organization. This derives from the fact that a common problem in many firms is unnecessary reinvention (Swan et al., 2000). IT-based tools may increase the exploitation of existing knowledge by recording and storing experiences and thus making these available to others. In this way IT-based tools can be useful for processing information that already exists in the organization. Tools used for this purpose include document management systems, databases, data warehousing, and different groupware.

The purpose of systems for management of competence can be to enable management to see current status of expertise and needs for the future (Lindgren & Stenmark, forthcoming). Another objective may be to help an organization to categorize and visualize expertise, in order to make competence management possible. The systems are more or less designed with a top-down approach, and some of them have a hierarchical structure. This means that only management can see subordinates who in turn only can see themselves, but there are also other systems that allow every co-worker so see everyone else as well as themselves. Both kinds of systems have their advantages, but they also fail to satisfy many aspects needed.

Lindgren and Stenmark (forthcoming) state that it is difficult to choose a strategy for mapping competence within an organization no matter which perspective you choose. One strategy could be to create a specific competence structure. To do this the organization needs to find the suitable competence categories, and to fit all employees into them. If the systems contain unstructured information it might be difficult to find specific expertise, and it might also be difficult to describe the expertise in free text. Another difficulty has to do with knowledge evolution, i.e. that an individual might have the ability to perform a certain task, but has long ago moved on to new, and maybe more qualified, assignments and might not have any interest in the desired job. A third difficulty is that systems of this kind also commonly lack the possibility to connect individuals with similar interest profiles, i.e. knowledge interaction.

The limitations of these systems have revealed the need for a new perspective on KM systems, which will be discussed next.

3.2.2 A Softer Perspective

The emergence of the community approach has shown that IT systems need to help people communicate their knowledge and competence, not just to store it, and to connect people so that they can think together and turn information into solutions through actions. IT tools used for this include different types of groupware programs, and communication and co-ordination programs to support knowledge collaboration between interdependent individuals and groups that are geographically dispersed (Snis, 1999). In recent years e-mail, groupware-packages, hypertext-systems, and other systems have been further developed for this purpose. There are also further groupware efforts supporting KM, e.g. web-based applications providing for workgroups via collaborative workspaces.

The design of these systems has more of a bottom-up approach. Their purpose is to emphasize dynamic collaboration and communication between interdependent individuals or groups that form a community. Many of these systems allow participants to contact others to share knowledge and competence freely with whom they want. The structure of the systems is usually flat and each individual is responsible for communicating through them. A problem with this approach is that it might be difficult to make management aware of the employees' aims, directions, and ambitions, i.e. knowledge empowerment (Lindgren & Stenmark, forthcoming).

In both these perspectives it is evident that it is relatively easy to create IT support for explicit knowledge. However, it is more difficult to formalize and interpret tacit knowledge. Davenport (1996) states that there is a need for a hybrid solution for KM, in order to capture both. He suggests that if the effort to formalize tacit knowledge is too great it should be left informal and processed by humans. Furthermore, tacit knowledge must be interpreted in a broader context and combined with other types of information and consequently, humans are the recommended tools for this purpose. Nevertheless, during our study we have found that there are alternative techniques, i.e. non-traditional KM systems, for supporting tacit knowledge. Stenmark (2001) has studied how an agent-based RS can be used to make tacit knowledge tangible. His work involves what knowledge that governs individual activities, based on interest, and how tacit knowledge may be put to use in a community. He proposes an

interpretation on how tacit knowledge may be activated in an organizational setting. McDonald and Ackerman (2000) also have studied RS, but their research focused on recommendations of expertise. They state that RS, in contrast to many other collaborative filtering systems, rely on implicit opinions rather than explicit ratings. Therefore they can assist in finding people who have certain expertise, and who may not otherwise be identified.

As stated above, research has shown that there are certain limitations to systems supporting competence management (Lindgren & Wallström, 2000). Traditional KM systems mostly support organizations, not individuals. They store competence structures rather than visualizing tacit know-how. As far as we know, none of them support interests either. Drawing conclusions from the research conducted by Stenmark (2001) and McDonald and Ackerman (2000), we believe that agent-based RS can fit this more dynamic perspective, even though they have not yet been thoroughly studied for KM purposes. Inspired by these researchers we studied an agent-based RS, i.e. VIP, on a corporate intranet to find out how it can be used to support KM in general, and management of competence in particular, among dispersed co-workers. The study of VIP took place in an organizational setting within the Volvo Corporation.

In the next section, we will give a brief background to approaches and techniques used to retrieve information from the Internet and intranets, and through portals. We will also describe the emergence and purpose of RS and intelligent agents, and how these techniques can be used together.

4 New Techniques to Support Knowledge Management

The system on which we based our empirical study, i.e. VIP, is an agent-based RS. It is designed as a portal, and is situated on a corporate intranet. Its main purpose is to aid users in the retrieval of interesting documents, but it also has functions to support communities and to visualize interests and expertise. In the next section we will describe the system in further detail, but first we wish to give a background to the emergence of such systems. In this section we will therefore present some aspects on, and problems of, information retrieval from the Internet and intranets. We will also give some examples of tools used for this purpose. Then will we describe some features and give some examples of RS before we move on to agent technology.

4.1 Internet, Intranets and Portals

The Internet first emerged as a community of professional users with similar interests. Such a professional part of the Internet still exists, but it is largely constrained to internal networks between organizations. As the Internet develops and grows, the sharing of information over this medium becomes a less predictable, efficient, and manageable enterprise. The sheer size of information makes it impossible to preserve the informal character of the early Internet, as well as the trust the first Internet users had in the quality of the Internet-based knowledge bases (Vishik & Whinston, 1999).

The intranets of large corporations have become major channels for intra-organizational information. The notion of information overload leads to a decrease in the level of awareness of what goes on within an organization (Stenmark, 2000). Both organizations and individual employees need to deal with this vast amount of information and, according to Foltz and Dumais (1992), "*they continuously seek to find the right tools to pan out the nuggets while minimizing the need to search and filter the information*". There have been various efforts to develop IT tools that help organizations and individuals to search and filter the information. Nevertheless, Stenmark (2000) means that it, from a collaboration point of view, is unclear whether these tools cover the same ground, partly overlap, or complement each other.

It is almost impossible for a knowledge worker to determine the quality and validity of online resources. This is partly due to the fact that information related activities often occupy a fraction of the worker's time. However, the challenge of dealing with the vast amount of electronic documents in many collections, formal and informal, on the Internet and intranets remains. As Vishik and Whinston (1999) states, it is the growth of occasional audiences of Internet and intranet users that is the most important factor in improving the quality of digital products. There are a number of search tools to help escape this problem, but hitherto they have been plagued by the following four problems (Garofalakis et al., 1999):

- *The abundance problem*. The phenomenon of hundreds of irrelevant documents being returned in response to a search query.
- *Limited coverage*. There is no single tool covering all the diverse, dynamic and unstructured information.

- *Limited query interface.* Search tools, based on syntactic keyword search, limit the user's ability to formulate helpful questions.
- *Limited customization*. Search tools available to individual users lack customizable features and functions.

One way of dealing with the problems mentioned above is the effort of building portals on the Internet and intranets. Portals have been introduced to sort, filter and push the appropriate internal and external information to a user desktop (Latendre, 1999). Corporate portals present information, gathered from a variety of sources located both inside and outside an organization, on a user interface. Such interfaces are usually customizable to suit personal preferences. Portals can also support collaboration by providing users with access to knowledge and business processes, workflow tasks, and discussion groups. According to Latendre, the purpose is to deliver information in the context of everyday business operations, which can allow users to work more efficiently, and to share insights, experience, and expertise through the portal interface.

While the Internet and intranets have become more widespread in use, different techniques for retrieving information from these new 'webs' have emerged, and are becoming more and more public by the usage of search engines or similar tools (Delgado, 2000). In the following subsections we will briefly describe some of these techniques.

4.2 Information Retrieval

There are two perspectives on information retrieval (IR), which focus on how content is handled and what roles the users have. These perspectives have been labeled push and pull technologies (Stenmark, 2001), and all commercially developed IR-systems are characterized by either one of them.

The push technology focuses on how content providers can deliver added value to their customers by adapting to user behavior and learning how to recognize user preferences. This perspective is primarily used in order to help the content provider. In other words, if Amazon (2000) wants to recommend books or music it does not provide references to competitors while using its push technology. Research shows that the Internet sites using push technologies suffer from some characteristic problems, e.g. that it offers too many options, making it very difficult to decide from where to subscribe (Delgado, 2000). Moreover, the continuous flow of information even worsens the problem of choosing which information, within a certain channel, that most likely would catch the users' attention.

The pull technology focuses on the users' needs and pulls out the information requested by the users. This perspective leaves the decision making to the users, so that they freely can choose from available resources. Consequently, the users themselves must find suitable tools for retrieving, or pulling out, desirable information. However, this approach suffers from the same problems, or limitations, as described above.

Both technologies mentioned above use search engines to retrieve information, and such tools were used to retrieve documents even before the emergence of Internet. The question of how to retrieve documents from textual databases has occupied library science for many years, and lately research areas within IR have spread and now also include the context of the Internet (Delgado, 2000). There are two dominant categories of search engines, which retrieve and present information in different ways. The first is indexed search engines, where AltaVista (2000) is one commercial example. With such a search engine users are required to specify their information needs in terms of a query, i.e. as one or more simple keywords. The query is then compared with documents in a collection of potential relevance to the user. The other category is directory search engines, e.g. Yahoo (2000), which search by traversing a topic hierarchy. The hierarchy can be divided into several sub-hierarchies and the search engine is used to drill down the hierarchies until the user finds what he or she looks for.

Humans not only search for information, but also often need to make decisions, based on the retrieved information, before acting in some direction (Resnick & Varian, 1997). People often have to make choices without having sufficient personal experience of the alternatives, and accordingly have to rely on other people's recommendations, e.g. by word of mouth, recommendation letters or reviews. One type of IT tool developed to support this process is RS.

4.3 Recommender Systems

Recommender systems (RS) have hitherto been developed in various ways, depending on their backgrounds and objectives. Academic research, as well as the success of commercial products (Autonomy, 2000; Firefly, 2000; GroupLens, 2000), has shown that the different variants of RS have been successful (Stenmark, 2001). Even if there is a multitude of RS it is difficult to divide and categorize them. This partly depends on the rapid technology development, and partly on the lack of precise collaboration between the different techniques that can solve the same problem (Delgado, 2000).

4.3.1 Characteristics of Recommender Systems

Traditionally RS let people provide recommendations as input, which then is aggregated and directed to the recipients of the system. The main functionality of these systems lays either in transformation of the aggregation, or in the ability to make good matches between the recommenders and those seeking the recommendations. Delgado (2000) ascribes RS the ability to anticipate what items a user is likely to be interested in, and to recommend such items in an intelligent way. Turnbull (1999) describes RS as a new implementation of IR, and states that the difference between these two concepts mainly consists of that feedback is more important in RS. They also differ in how the aggregated feedback is measured, and in how retrieved information is used.

There are at least two incentive problems with RS (Resnick & Varian, 1997), concerning the establishment of interest profiles and content owners. The first problem is that RS have to offer some incentive to users in order for them to provide

recommendations to others. The providers could be offered money or some other stimuli as compensation for sharing recommendations. The second problem concerns content owners, who may generate a vast amount of positive recommendations for their own material and negative recommendations for their competitors'. The researchers also mention concerns about personal privacy, i.e. when people use RS it is likely that they do not want their habits or views exposed publicly.

There are different types of RS that emphasize either the recommended items or collaboration between users of the systems. These will be described next.

4.3.2 Content and Collaborative-Based Recommender Systems

Content-based RS use machine-learning algorithms for the task of classifying what is interesting to the users or not. They often train and test their algorithms on a set of previously decided web pages, before they are used in their real environment. These systems only make recommendations within a specific topic domain. For instance, the topic could be extracted from a previously decided on-line directory or a list of web pages retrieved by a search engine, e.g. Lycos (2000) or HotBot (2000). This technique does not give any advice to the users about whether the domain is interesting to them or not, and new documents that enter the domain have not been categorized. The user profiles are not automatically updated once they are built, and retraining of the systems is performed off-line. Examples of content-based RS are WebWatcher (Lieberman et al., 1999), which sorts and arranges hyperlinks on a web page, and the WebTaggerBookmarkService, which is made for organizing and sharing URL:s, but lacks some sort of recommendation engine (Delgado, 2000).

Collaborative-based RS are more rarely found in the research literature (Delgado, 2000). Delgado also states that there are more commercialized collaborative than content-based RS, and that the former usually gives recommendations in the entertainment domain. Usually, they have a centralized structure and work with indexing. Collaborative-based RS have functionality to determine similarity between users, regardless of the content of the recommended information. In these systems the users are asked to rate content, and the system then uses some predefined scale to find correlation between the users, based on the user-rated content. Examples of collaborative-based RS are Ringo (Lashkari et al., 1994), i.e. a RS for music and, based upon Ringo, Firefly (Shardanand & Maes, 1995), which is designed for recommending for example books and music. A third collaborative-based RS is GroupLens (Resnick & Varian, 1997), i.e. a system that provides personalized selections of Netnews, which is a site for on-line news organized as newsgroups. Answer Garden (Ackerman & Malone, 1990; Ackerman, 1994) is a system that is used to categorize questions and answers within an organization, forming a database comprising all of the organization. If users navigate through these questions without finding an answer, they will eventually find themselves at leaf nodes, where experts are waiting to answer their questions.

4.3.3 Hybrid Recommender Systems

There have been attempts to combine content and collaborative-based RS, and to merge the advantages found in each approach (Balabanovic & Shoham, 1997). These advantages are as follows:

- Collaborative recommendations allow users to benefit from others' experiences, which in turn can serve as a basis for more than incomplete and imprecise content analysis. The recommendations can even be beneficial to users who have not rated any of the same, only similar, items.
- Content recommendations visualize items previously unseen by others.
- Using the system's built-in profile allows recommendations of items, even if there are no other similar users available.
- Utilizing group feedback could potentially reduce the time and development cycles needed to get mature personal profiles.

Combined content and collaborative-based RS are sometimes labeled hybrid systems, or social filtering systems, when they work with both indexing and matching. Examples of such systems are Yenta (Foner, 1997) and ReferralWeb (Kautz et al., 1997), which recommend people instead of documents. Turnbull (1999) explains that these hybrid systems can have a more team-oriented purpose, and might involve some explicit ranking or rating of resources. They are more active than the systems referred to above, since they deal with constantly changing information, newly added information, and shifting relevancy. These RS often have elements of intelligent agent (IA) technology. RS combined with such technology enable a certain degree of autonomy, which further allows monitoring without intervention (Stenmark, 2000). Thus, an IA-based RS could for example detect new or updated information, without depending on whether the user is present or not.

One example of RS, which is enabled with IA technology, is RAAP (Delgado, 2000), which can filter documents and perform on-line learning of user profiles. The profiles are later used to match similarities between users, and between personal interest domains. Even if RAAP is a hybrid system, it emphasizes the collaborative approach. Another system in this category is Lets Browse (Lieberman et al., 1999), which features automatic detection of user presence, automated browsing of channels, dynamic displays of the user profiles, and explanations of recommendations. Expert Finder (Vivacqua, 1999) is another attempt to combine these technologies, and the purpose of this system is to locate expertise within predefined communities. In order to give the reader a broader understanding of how an IA-based RS may work, we will first give a description of intelligent agents.

4.4 Intelligent Agents

IA can be viewed as autonomous problem solvers, working on behalf of humans, but there are probably as many definitions of IA as there are agent systems. The notion is somewhat blurred, as both researchers and companies define different filtering

techniques as well as shopping bots[1] as IA (Feldman, 1999). Since we share Feldman's perspective, i.e. the knowledge worker's, we have also adopted his definition of IA:

> *"Intelligent Agents are software programs that can identify repetitive patterns of behavior, similarities between events or things, and changes in patterns over time, on behalf of their human user".*
>
> (Feldman, 1999).

Throughout the remainder of this thesis we will use the term agents when referring to IA.

4.4.1 Agent Characteristics

Feldman (1999) describes the behavior of an agent as gradually acquiring the patterns of its user's work practices. Once an agent is trained it can function on its own, with some occasional check with the user to sort out unclear or new situations. This is the goal of agent technology, and it is rooted in the idea that patterns of behavior can be identified and described. If the patterns can be described the user can learn something from them, which in return can help the user to train the agent to automate these patterns. According to Feldman, intelligence has other components besides creativity, which computers can manage in a better way than humans, e.g. as computers are good at following rules and patterns, we can explain the similarities, differences, and patterns that they need to learn. Even if it is more difficult to teach them how to find new patterns, this is not impossible to accomplish. Consequently, besides simply being computer programs with a knowledge base and set of rules, agents are characterized by their adaptability and learning capacity. Next we will list some basic characteristics that are agreed upon by several researchers (Feldman, 1999; Delgado, 2000):

- *Autonomy and proactivity.* Agents can learn to remember their users' interests, and can act autonomously. They can engage in cooperative processes with other agents, as well as with humans. An individual can have many agents, designed for different purposes, acting simultaneously, as alter egos, in different places at the same time. Proactiveness emphasizes that agents do not simply act in response to their environment, but can also have certain goal-oriented behaviors and take initiatives towards those goals.
- *Adaptiveness and reactivity.* Agents have the ability to learn as they react to, or interact with, their environment, resulting in a performance that improves over time. These characteristics also convey the ability to react to new circumstances, depending on how the agents have been programmed.
- *Social ability.* Agents interact with other agents via some kind of agent-language, to conduct conversations and exchange meaningful messages.

[1] A shopping bot is an intelligent agent that with predefined parameters for example searches the Internet to find the lowest price for a specific item.

- *Collaborative behavior*. This feature is based on the concept of social ability. To realize this concept users need multi-agent systems (MAS), instead of only one agent. Each agent is given a discrete task, and the agents work together to establish which agent will carry out different assignments, and how to merge the collected information so that it can be presented to the user.
- *Mobility*. Agents have the ability to migrate, in a self-directed way, from one host to another on a network, in order to perform assigned duties. The duties may include information gathering at each host, or balancing workload or traffic on the network. This can be considered as an extension of the original concept of autonomy.

4.4.2 Agent Technology in Use

Agents can be based on several technologies (Feldman, 1999). All of these use some combination of techniques, e.g. statistic operations, artificial intelligence, machine learning, inference, and neural networks. Agent-based systems are usually not plug and play, but need to be trained, and most of them require examples of correct answers or rules for appropriate behavior. An agent-based system is implemented in several stages. First, rules of training data have to be set up and secondly, users train their agents by giving them either rules or a large set of examples with the correct answers included. Once an agent-based system performs satisfactory on the training data, it is ready to work on test data to make sure that it can extend what it has learned to unknown materials. A last, but continuing, step is to evaluate the performance at several intervals. Agents should learn over time, and their performance should improve as they adapt to their users' needs.

When training an agent created for learning user profiles, the user normally works with an interface in the form of an attribute-value list and/or rules. It is desirable that the user gives as much input as possible, in order to get useful results at the first round of retrieved information. However, it is known that (Delgado, 2000):

- Users are unwilling to provide a lot of time in order to create the specification of their profiles or value lists.
- User interests may change over time, making the profiles difficult to maintain.
- Consequently, initialization and maintenance of user profiles can become difficult aspects of the design and development.

The facts above imply that incentives are needed in order for users to devote time and effort to train their agents. However, the level of incentives may vary from system to system, since the level of automation and acquisition of user profiles can range from manual input, through semi-automatic procedures, to the automatic recognition by the agents themselves, and thus demanding more or less of their users. The system on which we based our empirical study, VIP, is an agent-based RS and the characteristics of hybrid RS mentioned above became evident when we examined it closely. In the next section we will describe it more specifically.

5 The Volvo Case

During a couple of months we had the privilege to conduct our empirical study at Volvo IT, where we got access to the application Volvo Information Portal (VIP). VIP is a portal on the Volvo corporate intranet, i.e. Violin, and it is a RS based on agent technology. Its primary purpose is to facilitate the search for relevant documents on Violin, and to enable the spontaneous creation of communities. At the time of our study, VIP only existed as a test version, and there were merely approximately 20 users who had tried it.

5.1 Overall Description

VIP is built on commercially available software from Autonomy´s AgentWare platform (Autonomy, 2000), and its main functions are based on a number of intelligent agents, in VIP called channels. These channels are either predefined by someone in the organization, *General Channels*, or designed for the individual user to create and customize, *Personal Channels*. The *General Channels* are available for all users of the system, but in order to access the *Personal Channels* users have to log into the system. The *Personal Channels* can be customized to find documents based on an area of interest expressed in natural language, i.e. a user can cut and paste a large amount of text, even a whole document, into a search input area equivalent to a query field of a search engine. This larger amount of text, together with a pattern matching technique, gives a richer representation of an interest domain than a keyword-based query would do (Stenmark, 2001).

A user can retrain a *Personal Channel* to make the results better match the text, written or pasted, in the search-input area. VIP also enables users to locate colleagues with similar interests, by matching their *Personal Channels*. Each user can have up to five different *Personal Channels* at the same time. If the user wants to define additional channels, one of the existing *Personal Channels* has to be deleted. When a user defines a channel, VIP immediately returns hyperlinks to a number of matching documents. This also happens every time a user opens one of the channels. A user can also choose to be notified via email when there are new, relevant documents that match the *Personal Channel*.

5.2 Technique Behind VIP

Since our thesis focus on the usage of VIP, we will not describe in great detail how it is technically implemented. However, for the sake of understanding the basic principles we will give a brief description of the technique behind the system (fig. 5).

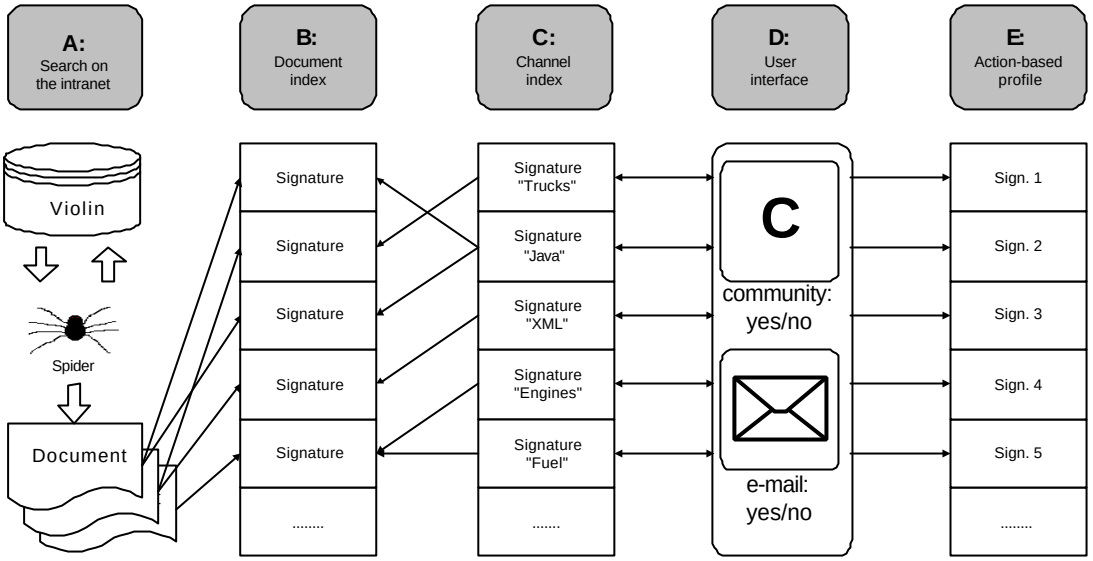

Figure 5: Illustration of VIP's Sequence of Work.

VIP is based on neural networks and pattern matching techniques. A spider, also called crawler or robot, searches Violin for documents once every twenty-four hours (A). The spider is initially given one or more URL:s when the search starts. The spider brings back all found documents within its scope. The system scans every collected document and gives each a digital signature, or fingerprint, based on its content. The signatures are then placed in a document index (B).

Every time a user creates, or edits, a *Personal Channel*, the channel is given a digital signature and is placed in a channel index (C). This signature is then immediately compared with the signatures in the document index to find out if there are any matches. Every time the user clicks on a hyperlink to a document, retrieved by his or her *Personal Channel*, the channel signature is updated. This update, and each time a channel is trained, contribute to the improvement of the network and initiate a new matching.

From the user interface (D) the user can choose to be notified via e-mail when there are updates or new documents on Violin. The user can also choose to be available for communities. This means that the user is visible for other users with similar *Personal Channels*. The matching of similar *Personal Channels* takes place in the channel index, through the creation of a temporary channel that seeks matching channels, by comparing signatures.

As the user searches Violin through VIP, i.e. clicks on document hyperlinks in *Personal* or *General Channels*, up to five different invisible action-based profiles are created. When the user clicks on one of the hyperlinks, the related action-based profile is updated (E). The purpose of these profiles is currently only to automatically recommend new documents, based on earlier searches performed by the individual user. There are no functions for comparing action-based profiles of different users or to configure these profiles manually. Consequently, they do not serve any community function.

5.3 Specific Functions

There are some specific functions that need to be mentioned, in order to give the reader a better understanding of the system. These functions are displayed as hyperlinks, which are located in a navigation bar, and appear in the main frame when activated. We will not describe all functions available in VIP, but concentrate on those relevant for our study.

5.3.1 General Channels

On the start page (fig. 6) there are a number of predefined *General Channels*, which are divided into different *Information Categories*, e.g. Automotive. Each category has a channel and a designated editor, who is responsible for training and modification of the channel. The retrieved documents of these *General Channels* are public, i.e. all users can benefit from their results.

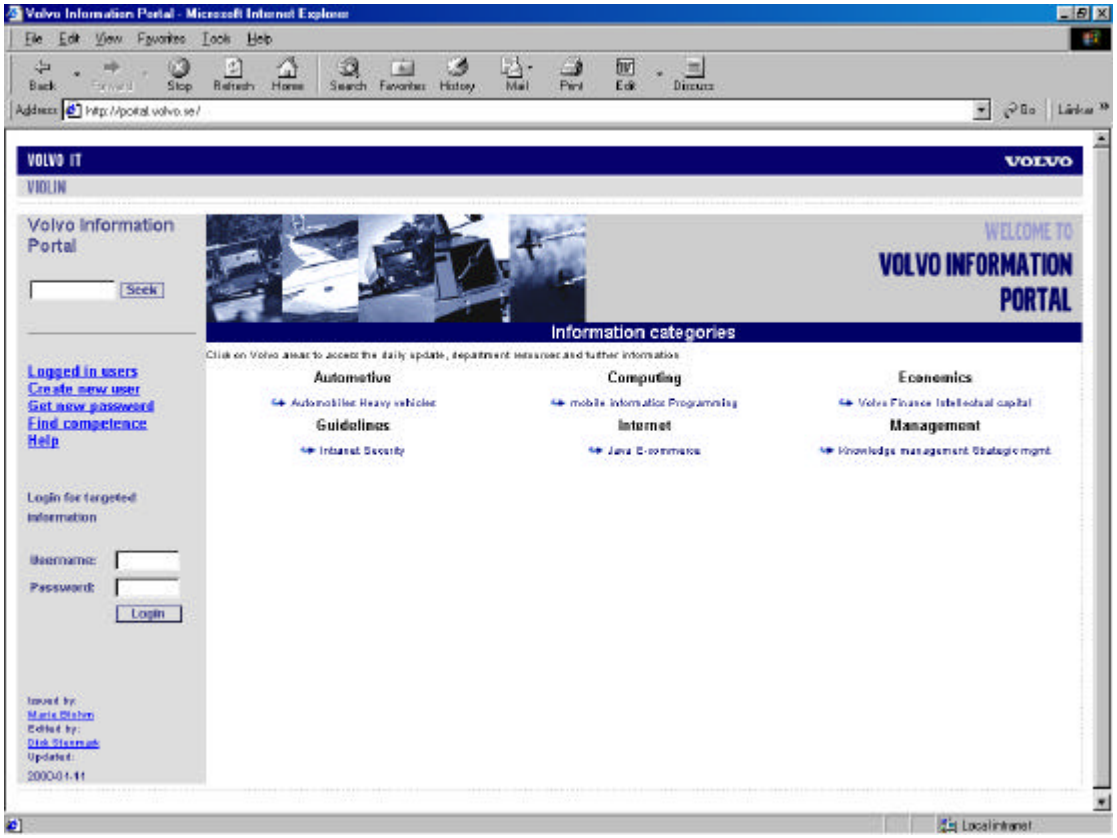

Figure 6: Overview of General Channels.

5.3.2 Personal Channels

This page, labeled *Your Channels*, displays an overview of the names of the user's own *Personal Channels* (fig. 7). A user can open one of his or her *Personal Channels* by clicking on the hyperlink with the channel name. These hyperlinks lead to pages,

which list the documents that best match each *Personal Channel*. Each displayed list contains document headers, weight, i.e. ratings of retrieved documents made by the system, and further hyperlinks to document summaries and related documents. To the left of every document header there is a selection box. If the user wants to make a *Personal Channel* more specific and narrow down the search, this can be accomplished by checking the boxes of the most interesting documents, and then clicking on the *Retrain*-button.

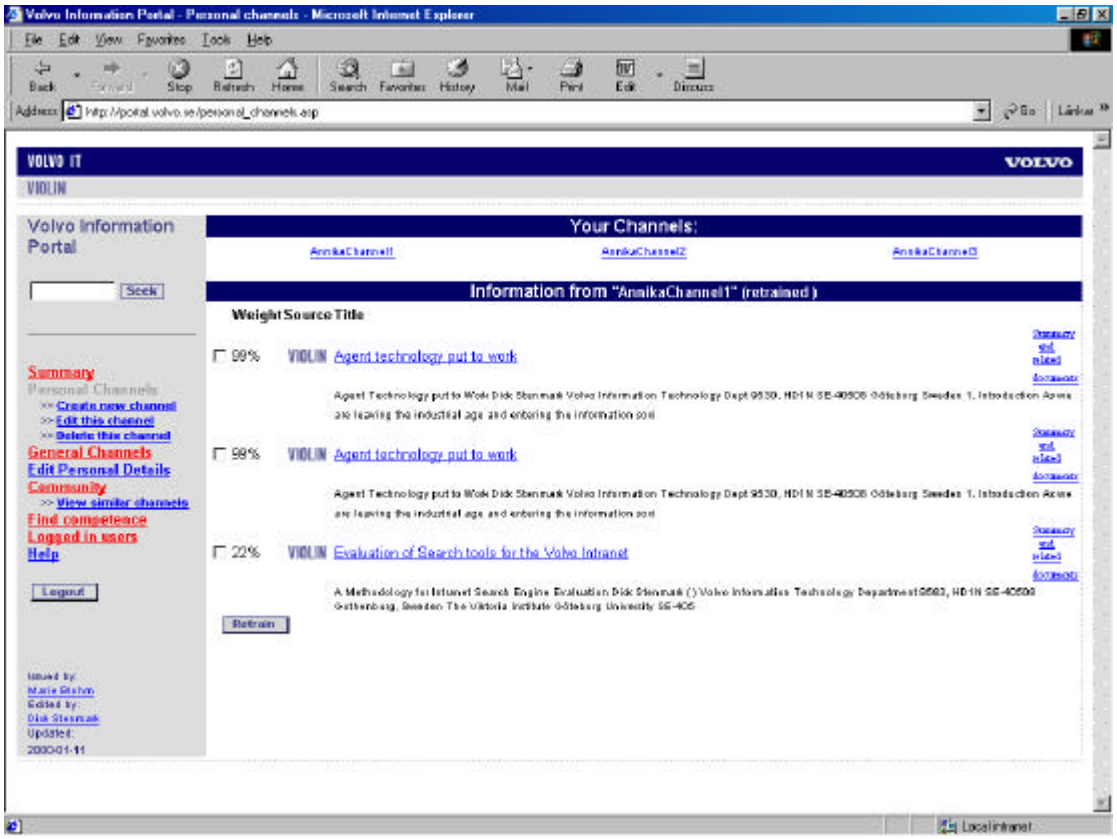

Figure 7: Overview of Personal Channels.

5.3.3 Editing a Personal Channel

From a selected *Personal Channel* a user can click on a hyperlink to edit the channel. The edit page consists of configurable parameters, which enable the user to customize the channel according to the user's own preferences (fig. 8). The text in the search input area, *Channel Training*, is editable. *Number of results* shows the number of documents displayed on each page, and *Minimum quality* sets a degree of quality of a match. The checkbox *Keep retrain* enables this channel to be retrained continuously, and *Add to community* enables the user to be visible in the community of the VIP users. The user can leave an e-mail address in the field for *Email address* if he or she wants to be notified of document changes matching the *Personal Channels*.

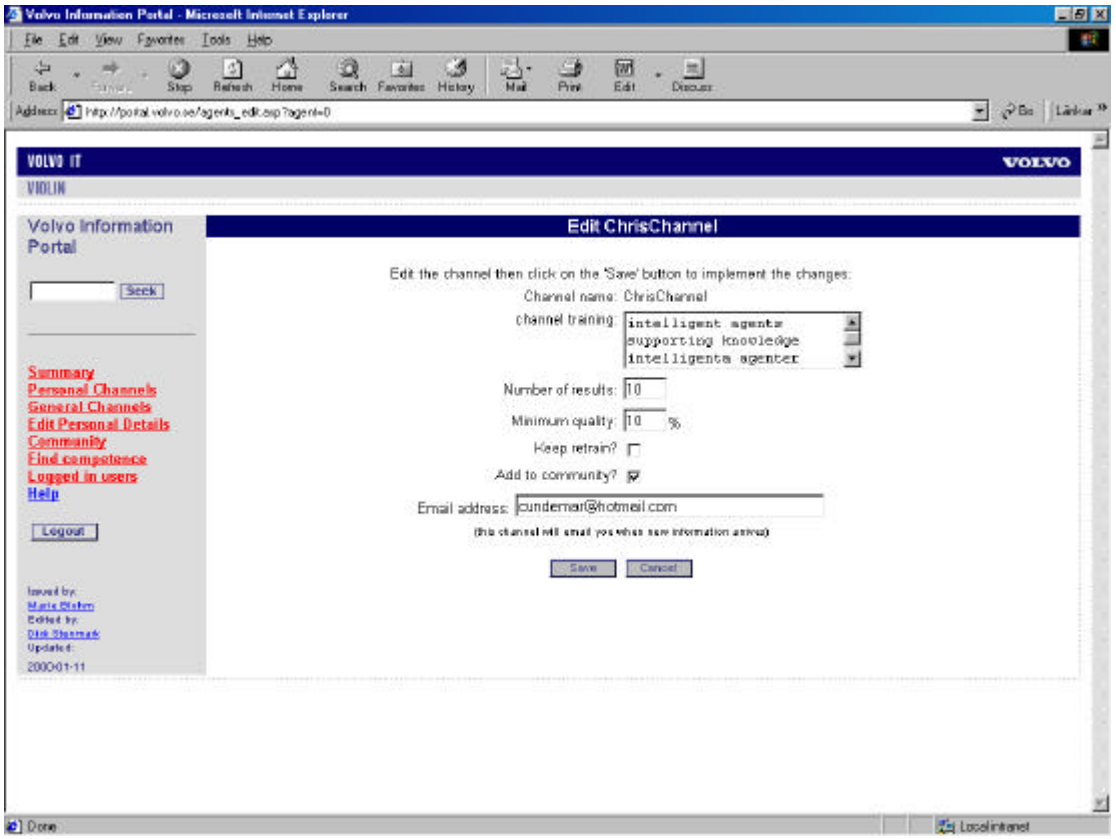

Figure 8: Overview of Editing a Personal Channel.

5.3.4 Find Competence

When a user clicks on *Find Competence* (fig. 9), a temporarily created channel is matched with the *Personal Channels*, a proceeding similar to the comparison of similar channels mentioned in subsection 5.3.5 below. The *Find Competence* page consists of a search input area, *Enter text to search for certain competence*, where the user writes or pastes texts describing which competence area he or she is searching for. If there is a match with other users a list of those other users are shown. This lists displays *User* with a hyperlink to the user's e-mail, *Fullname*, *Company*, *Department*, *Site*, *City*, *Country* and *Phone*. If there is no match, the page displays that there are no users with similar competence.

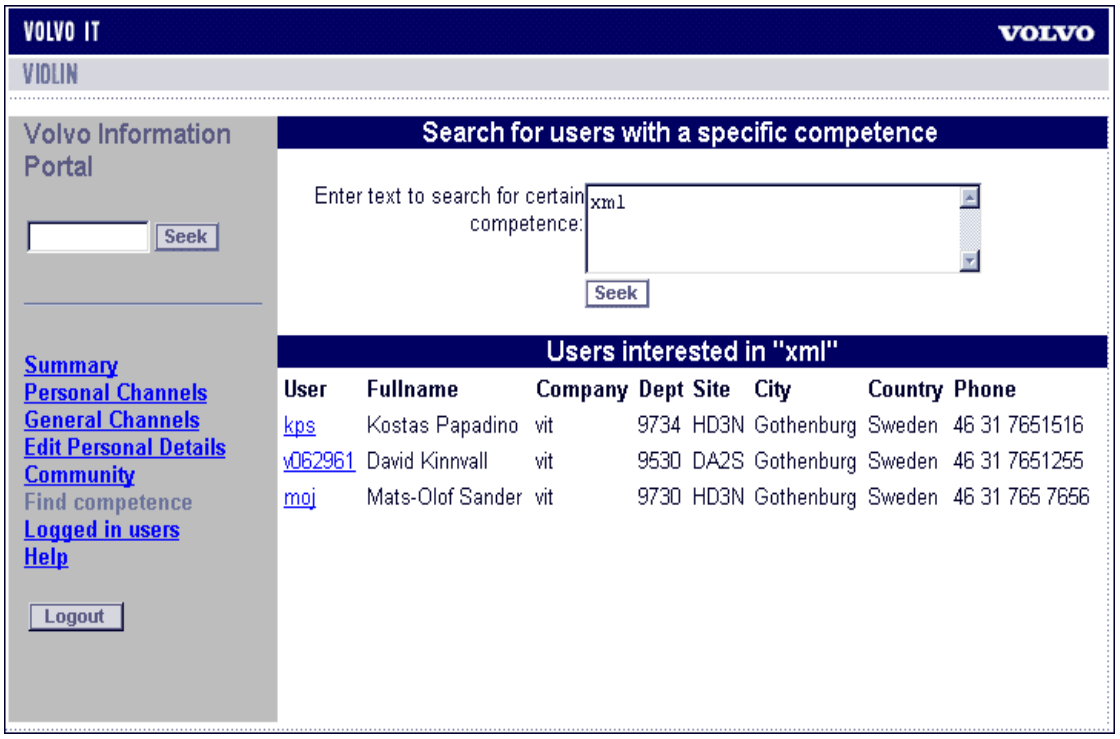

Figure 9: Overview of Find Competence.

5.3.5 Community

The *Community* feature shows all other users who have one or more channels that match one or more of the user's own. If there are users who have defined similar channels the page displays the same information as *Find Competence* (fig. 9). No further information is given about which specific channels that match or what they contain, only who own the matching channels. If there is no match the page displays that there are no users with similar competence.

5.3.6 Logged in Users

This page contains information about all users who are logged into VIP at the moment. It displays *User* with a hyperlink to the user's e-mail, *Fullname*, *Company*, *Department*, *Site*, *City*, *Country* and *Phone*. However, this function is not a very good indicator of active users of the system. The only time a user has to log into the system is when he or she has to create, edit, or retrain a *Personal Channel*. Another reason for logging in is to view the results of the *Personal Channels*. However, these results can also be retrieved by e-mail, which means that entering the system is not required in order to be an active user.

5.4 Related KM Efforts at Volvo

Above we have described the system on which we base our empirical study. During our time at Volvo, another KM project took place. Although this project, the TP/HR project, was outside the scope of our thesis, and the TP/HR application was not part of our study, it was referred to by several of our interviewees. Therefore, we have chosen to give a brief description of the project and the application here.

5.4.1 The TP/HR Project

The main purpose of the TP/HR project was to identify competence structures for Volvo IT, and to enable the employees to be mapped according to the competence structure in the system. The defined competence structures were to be used in several areas, e.g. resource planning and gap analyses (Lindgren & Stenmark, 2001). Other project objectives were to simplify internal recruitment, help the organization to steer towards its goals, and to get a common and structured language. Additionally, there was a wish to make the competence searchable. TP/HR was also intentioned to be a support for goal and planning conversations, and personal development discussions.

5.4.2 The TP/HR Application

The TP/HR (TietoPersona/Human Resources) application is a commercial 'off-the-shelf' product from Tieto Datema AB. It is a client/server module-based system, which supports a hierarchical competence structure. TP/HR is based on a number of competence windows where competence data are grouped in a hierarchical tree structure. The top levels in the tree are categorized as functional and technical skills, which in turn has sub-levels, consisting of different skills. Functional skills refer to the work you perform and technical skills refer to the methods or techniques used to perform tasks. These competence windows are customizable to contain special types of information, e.g. background, education, courses, and certificates achieved. TP/HR allows users to compose CV pages that contain expertise.

The defined functional and technical skills have to be graded on a five-graded scale and there is another five-graded scale for competence grading. These scales build up a matrix used for competence mapping, and for further group and gap analyses. Here the employees are able to express their wanted skill level and/or new areas of interest relevant to a specific competence area. One of the TP/HR functions enables searches for people holding a certain skill on a particular level. Another function allows two categories of group analysis. The first analysis, i.e. a gap analysis, can show how well the employees' expertise match given demands for each work task during different times. The second analysis indicates how critical competence is distributed within a group and for each respective work task.

You must be an authorized TP/HR to be able to log into the system. The employees are only authorized to see their own profiles. Group managers are authorized to see the profiles of all employees in their own departments or groups. The authorization roles are set manually and new users have no authority whatsoever, until a manager have set up a user role and given access to each employee.

5.4.3 Our Interpretation of Traditional KM Systems

In the thesis, we refer to the term traditional KM systems and TP/HR is one example of such systems. Therefore we feel the need to give a brief presentation of our interpretation of the term. With a traditional KM system we mean a tool in which the knowledge and competence categories are centrally structured. This means that the information is static, explicit, and easy to retrieve. However, it is difficult to fit and find tacit information in the system. It is designed for usage from a top-down approach, i.e. for management control of the knowledge and competence within the organization. It is used and updated in one of two ways. Either all users are allowed to update their own information profile or a small number of users, e.g. within the HR function, have access to the system for updating and maintenance. In the latter case, the information has to be gathered from the employees and thereafter fed into the system. In both of these cases there is a need for some kind of incentive to secure the cooperation of the individual employee.

6 Empirical Findings

Our study comprised a literature study, an evaluation of VIP, two seminars, and interviews with sixteen users. In this section we will account for the results of the latter. The users that we interviewed had tested VIP during periods of different length, lasting from one week to two months. Their varying backgrounds and profiles contributed to the rich picture we got of the application. While conducting the interviews two major tracks emerged. One track covered areas where VIP could be useful and the other track highlighted areas where VIP shows some deficiencies, but where it could be further developed to fill a function. In the next subsection we will describe the first track, i.e. fields of application, which thus includes what VIP can be used for today.

6.1 Fields of Application

The VIP-specific part of our study took place in May and June 2000. As described in the previous section, i.e. the Volvo case, another KM project was carried out in the organization at the same time. This project concerned the implementation of a tool for competence mapping, i.e. TP/HR. Competence mapping involves getting an overall picture of the competence within an organization. The information received can be used to study the competence history and to make predictions about future demands for competence. In TP/HR information about the employees' competence and experiences were to be registered according to predefined parameters and levels. Someone at the HR department would then perform updates, in consultation with the individual employee. As a HR manager said, it is a system for controlling what you have.

> *"TP/HR is a lot more about order... order and to be in control of the situation... and to know what we have and the level of education of our co-workers, how many of these and how many of those. Then this portal is something else. It is what people do on an everyday basis. It is what they use their skills for. It is sort of in the next step. To me they don't even come close to each other. "*
>
> (HR Manager)

While a traditional system for competence mapping supports the historic view of competence in an organization, VIP gives an idea of competence used on an everyday basis. It reflects the activities of the employees, as their personal channels continually search for information about subjects they find interesting. For example, an employee working with development of engines can have several channels. One of these might search for documents on petrol engines, and thus reflects what the employee works with. Another channel might search for documents on non-polluting fuels and together, these channels might tell that the employee has an interest in development of environmentally friendly engines. In other words, VIP can function as a complement to competence mapping, since it may add everyday information to the competence information found in a traditional system.

Furthermore, VIP can help to identify the different roles that employees have within the organization. Roles can depend on assignments, position within a project group, knowledge and skills etc. A traditional KM tool allows co-workers to see which roles an employee has been assigned by the organization. VIP on the other hand shows other roles that individuals might, subconsciously or deliberately, embrace in different situations. In a project organization, the roles of employees often vary from one project to another. While a tool like TP/HR can give a more historic role description, VIP mirrors the role an employee has at the actual moment of search. Another difference is that the roles are not explicitly defined in VIP, as in TP/HR. Thus, VIP can be used to find unstructured information about roles that cannot be found in a more traditional KM system.

> *"I have not indicated a role… it can be an area of activity or a task or something like that. For example it could be that if you work with CAD in this company then there are many units developed that could form a network… I just see this function as very useful since you have a database with individuals' profiles anyway. One could easily access that information, so to speak."*
>
> (Systems Developer)

The possibility of finding roles in VIP conveys that it can contribute to the creation of communities of practice, i.e. spontaneously created networks nourished by interaction between the participants. Through the search after users with similar interests employees can find others, whose existence they had no previous knowledge about. If and when these contact each other, a network will be created.

> *"It is very interesting. I see it as a very good function, a possibility to build networks. It is interesting to be able to find colleagues with the same… who for some reason are interested in the same things especially if you have chosen a narrow search area. For our main problem here is that there are people who work with similar things everywhere. And you don't really find them. We thought when we started the group for project managers that networks have been built since the dawn of time here and surely there are a bunch of people who have worked with this and there are surely some really good network builders who have… but it turned out that there are none. However, people have worked with it off and on, but no one have… no one kind of boasts with I have built networks, and I know how to. Strange. So it turned out to be an aha experience."*
>
> (Project Manager)

Another area that became clear during the interviews was the possibility to identify driving forces, i.e. motivational factors that produce a strong effect on employees' work situations. Driving forces can be a number of factors, e.g. interest. An employee motivated by commitment and interest might be better suited for a certain task than another individual, who lacks such motivational factors, even though the formal competence might be lower.

> *"When you take initiatives, beyond assigned tasks… Then there is a commitment to and an interest in participating in changing things. Commitment really is worth more, says more, because I don't really have to do it, no one is forcing me to and I am not measured by it. You can perform miracles in ten minutes if you have enough motivation. It doesn't have to take days. Therefore I have found it exciting to look for those with an interest and not those who are designated to do it, because they are not always best suited to."*
>
> (HR Manager)

If one can form an understanding of what individuals are interested in and committed to, this will give an advantage when assigning tasks and appointing project members. Project configuration means searching for and finding the right participants for a project. Even if TP/HR may be a better tool for finding precise, measurable competence, VIP can still contribute with significant parts, as mentioned above. As several interviewees expressed, commitment and interest are important building blocks of a project and should thus be seriously considered when configurating a project group. The interviewee cited below implied that not only competence, but also hunger and motivation, are important in a project. Therefore it is crucial to consider the interest that nourishes the hunger.

> *"The basic level for competence is to find out things, not just to sit down and wait. Basic competence and attitude, competence and interest, it is sufficient for me to put together a project around that. That people are interested in it. Of course it also depends on how large or how acute the project is. Then you will have a hungry person and that is more important than having a competent person."*
>
> (Technology Watcher)

VIP is an open system, i.e. it is available for all co-workers and allows them to freely search the mutual domain, i.e. Violin. This conveys a great possibility, when using VIP to search within areas interesting to a project, to find co-workers with interest in and knowledge about these areas in different organizational units.

VIP can also be used to identify competence areas, i.e. delimited domains such as Java programming or project management. Identifying these areas is important since they give a picture of the competence within an organization. Through a search for users with a certain competence or interest, employees with competence mapping a certain area can be found. Also, gaps in desired competence areas can be visualized. However, in order to find both existing and lacking competence the personal channels need to be named in a way that reflects their contents.

> *"Yes of course it gives a hint of that there is no one else than I who is interested in these areas. Yes it could be… would be able to show lacks. Missing competence for example… that there is a*

> *shortage in a certain area. Yes, then you could find areas that are neglected or where you are weak or, that is what it could be."*
>
> (Systems Developer)

Interest management was a further area where VIP was found to be useful. The technology watcher cited below implied that an analysis of the personal channels could visualize trends and wishes of the employees as well as competence gaps towards the organization. One way to amend the latter could be to edit the information categories in the general channels and add areas that hopefully would create an interest in new competence areas. If an interesting personal channel was found during the analysis, a general channel covering the same topic might be created. Another example of interest management in VIP could be delimitation or expansion of the search domains.

> *"A personal agent speaks about a momentum that people want... then you realize that maybe… thanks to an analysis of personal agents you might realize that there are competence gaps towards the company. Then you can create new areas that can make people see that there are more possibilities…"*
>
> (Technology Watcher)

VIP can also be a tool for competence analysis, i.e. to find out where employees' competence is today compared to where it has been and where it is heading. Within every organization there is a need for quick overall pictures of present situations, trends and directions. By studying existing channels, their owners, and for how long the channels have existed, trends can be visualized. However, this also demands that the users name their channels in such a way that the names match the contents.

> *"The most powerful thing I see is a possibility to map. Properly used this tool can provide a possibility to map in order to get a quick feeling for where people have been and where they are heading."*
>
> (Technology Watcher)

A traditional KM system might be better suited for visualizing an employee's competence at specific times. VIP does not have the same possibility to give historic information, but it can still indicate trends in competence development through snapshots, i.e. frozen pictures of personal channels including search criteria and retrieved documents. The comparison of snapshots taken at different times will indicate a pattern, enabling mapping of previously not visualized competence and interest areas. An analysis of the channels can also show nuances of expertise among different individuals, who in a more static system seem to have the same competence. An example of this might be the different interests interface design and database programming within the area of Java. Thus, VIP can contribute to a richer and timelier picture of individuals' competence.

6.1.1 Summary

During the interviews a number of different areas of usage, i.e. fields of application, were highlighted. These areas are closely related and sometimes partly overlap. The fields of application are as follows:

- Competence mapping complement
- Role identification
- Communities of practice
- Project configuration
- Competence area identification
- Interest management
- Competence analysis

6.2 Design Implications

During the interviews we also discovered a number of areas where enhancements of existing, and development of new, functions could give VIP a wider platform, from which to function as a competence management system. One of these suggested improvements was something that several expressed a wish for, i.e. a function that would create awareness of other employees' activities. Such indicators are important since individuals are influenced by other people's actions, whether they are aware of it or not. An individual's actions may affect other individuals in either a positive or negative direction. The influence of these actions becomes stronger when interests are shared or in situations with insecurity, caused by too many options.

> *"So, first I have to find… this channel looks interesting or I don't completely trust the search agent in all aspects, but I will do a manual evaluation as well… then I look at… yes, it looks interesting. Then I would go on, dig out more information… about the person and of course a subject before I talk to them. At a web conference that I attended a long time ago there was a registration system where you checked and decided which lectures to attend… and you could see who else attended certain lectures. That is rather stressful… for example what lectures to attend. Then I used this type of information… he attended the same lecture as me yesterday… let's see what he will attend today. It can very well be the same lectures that I am interested in. In that situation I felt I benefited from knowing what other people did. "*
>
> (Systems Developer)

To some extent, VIP can be said to support awareness of others, but only as far as visualizing users with similar interests and logged in users. A public list of all users and their channels would be an improvement. Such a function could to a higher degree influence users to search within the visualized topics, in a way they would not

have come up with on their own. Further functions could be a counter in the general channels showing most visited channels or a date field, indicating when new documents have arrived in an information category.

Another suggested design implication concerned training of the personal channels. Training largely depends on how experienced the user is and therefore some users with little or no experience of the procedure found it insufficient and in need of improvement. The problem here is that the first result from a channel totally depends on the quality of the text used as a basis for the search. This text might not be as representative as expected. Even though further retraining is possible, it will not yield satisfactory results if the initial text was not sufficiently representative. One way to deal with this could be to visualize the content and search criteria of other users' personal channels. Several interviewees expressed a wish for this. If the results of these channels would be found interesting, the users would like to have access to their search criteria to examine and maybe even copy these, in order to improve their own channels. The possibility to clone other users' channels was a feature that many interviewees thought would be useful. Such features could improve the quality of personal channels and spread the knowledge of how training can be performed.

> *"Then I have probably been too extreme in my wishes. Been to narrow… I believe I would kill the channel, create a new and make it wider to get a hit and then gradually narrow it down. It depends a little on how long you have trained the system. For I know that during the agent test you finally and with deep regrets had to destroy an agent. You realized that now there was too few and too little information for this to result in anything and then there was certain inertia inherent to the system. Subsequently you can't retrain in a right angle, but it is a slow process… if at all possible. Or I could… maybe I could try to copy another if the system made it possible. That I could get access to another that was placed here more openly and which… that looks good… I could have that as a starting point. Then a large part of the job would be done."*
>
> (Process Developer)

There was also a wish to see further information about other users. A couple of functions in VIP offer the possibility to find out more about users. These functions include links to information about name, organizational belonging, position, telephone, and workplace. The purpose of this information is to make it easy for users to contact each other for information exchange and building of communities. Nevertheless, several of the interviewees found this information insufficient, and expressed a desire for more, e.g. a photograph, a link to a home page, or information about current or previous assignments. The following quote describes this:

> *"Then you could imagine that if you click on the name… then maybe you could find… see a photo and some information about that I work with this. For there are many departments that have… but maybe it differs a little between different places, but some information about what you do… what you are involved in, so that you can see the person before you call."*
>
> (Financial Analyst)

In a traditional KM system competence is described according to predefined categories. The advantage of this course of action is that it makes it easy to search for explicit competence, but if the competence in demand does not fit within any of the predefined categories you have to look elsewhere. VIP, on the other hand, handles unstructured information, and in order to find the desired information the user has to be able to express the search criteria in a way that covers all or parts of the interesting domain. This procedure puts new demands on the user, as some of the user profiles found in the system might not have explicitly defined their competence, but instead have expressed it as an interest within a certain area.

> *"Maybe my interest is in areas in which I am not that competent. I have competence concerning recruiting and people are interested in finding me for that reason, but I do not express that competence in this system… so, there can be a gap. You can find people that are merely interested and not competent. Often interest indicates competence though, but not always."*
>
> (HR Manager)

The advantage of this way of searching for competence is the possibility to find competence that cannot be found in a traditional KM system. However, as the HR manager above implied, the drawback is that the users must be able to read between the lines and draw their own conclusions about competence from the information in the system. This will suit some users but not all, why we believe that VIP could benefit from a link to a formal competence description or a CV-page. This would eliminate eventual misunderstandings and enhance the understanding of an individual's competence when this is needed.

One of the functions in VIP helps users to find other users with similar interests. However, it can sometimes be difficult to decide how well another user's interests agree with one's own. Is the interest professional or does the user just want to be up to date with what is happening within a certain area? To be able to decide on the level of agreement, one usually needs a richer and deeper picture of the specific interest. VIP can give an indication of a person's interest, but in the end it is the users who must do a subjective evaluation of their impressions and determine the relevancy of the interest. This process could be simplified if a date- and timestamp was added each time that a personal channel was trained. The technology watcher below expressed that management also can have an interest in studying the employees' interests:

> *"Of course it is interesting to know if there is an interest in this area, whether we are heading in that direction or not. It's perfectly all right to add your name to a list if you have an academic interest or if you are really interested. Of course this in interesting, the question is how you… information that could be interesting to find here is for example if… some kind of analysis of the person's agents within the area. If you have two agents with different search criteria connected to the same area of usage, then you're interested. How many, and how you… even things like how long time you have used to build your agent, how often it is changed (trained) and that kind of information is really very interesting. However, then you're approaching an integrity limit, and where do you draw that line… For if you work a lot with it then you're interested, it's one dimension anyway."*
>
> (Technology Watcher)

VIP can be used to take bearings on the direction of interest groups within the organization. This can be accomplished by taking snapshots from log files. The difficulty with this procedure is that the information has to be compiled and aggregated manually. If this information could be retrieved in a more automatic manner, it would be easier to compare results from VIP with the overall direction of the organization, i.e. strategic and tactic goals. Such a measurement in VIP would give management a quick and flexible overview of the organizational status. The results of the measurement can be compared to results received at, for example, a development dialogue, taking place once or twice a year. Certain issues, which are not mentioned by the employee at such a meeting, can still be visualized through analysis of information in VIP.

If the organization could get an aggregated picture of the number of personal channels, searching a certain area, and how actively they are updated it would show how groups of individuals use their competence in practice. Thus, the organization would get an overview of what the employees take an interest in right now. The purpose for this could be a need for confirmation that the organization is on the right track, or a need for signals of discrepancies between the goals of management and different interest groups. As we mentioned earlier VIP can display lists of channels, but does not display their contents. However, such a public visualization would enable easier analyses and evaluations of the employees' interests.

> *"Well, I would… I would like to have… maybe it works but I don't know how to do… that is to find out what kind, these agents so to speak. Of course, I can watch other people's agents and find out things, but I would like to have a picture of… the amount and what search criteria people use. To get a map over whom, not who look for what but how many looks for what."*
>
> (Technology Watcher)

VIP functions well in the dynamic retrieval of information and snapshots, but lacks ways of aggregating the information and saving it for later use. However, freezing the

results of the channels at certain moments of time and saving the snapshots at different intervals could easily achieve this. If many snapshots were taken during different time periods it would result in a picture of the channels' movements. Further examples of functions could be statistic and graphical analyses and presentations of how interest groups move over time.

> *"The drawback of these agents is that they loose their historic information since they change all the time. There you would like to take snapshots of how the competence development looks like to be able to find a direction… historic information is always interesting for the direction. So preferably one would like to have snapshots… for the fact is that a certain interest group that have a certain appearance… and here seems to be an element right now. And then four years later it has change and then one would like to see that all right, they have developed."*
>
> (Technology Watcher)

Both the personal and the general channels continually search for documents within certain predefined domains. The content of the domains is dynamic, i.e. new documents are added while untimely ones are removed. This conveys that the results of the channels change in a similar manner, which can be both positive and negative for the users of the system. The negative part was mentioned before when discussing the difficulty in saving historic information, but there can also be benefits if the two types of channels are allowed to interact and exchange information.

> *"I believe that the general groups and the personal agents must interact for the reason that the personal agents tell that there is an element that people want. Then maybe you realize that there is a… thanks to an analysis of the personal agents you can discover a competence gap towards what the company desires. Then you can create new areas here… that can make people see more possibilities … so a reciprocal action here in between, I believe."*
>
> (Technology Watcher)

Today, there is no interaction between the two types of channels, except when an employee sees the result of a general channel, gets inspired, and creates a similar personal one. In a further developed VIP, the personal channels could give useful information to the general so that these could be improved and vice versa. Then the editors of the general channels could receive input from the personal channels in a way similar to when an employee gives a contribution to a suggestion box. VIP could regularly generate a list over personal channels, which then would be sent to the editors for analysis. This would give the editors a possibility to change the general channels, in order for them to serve the same interests as the personal. On the other hand, the personal channels could not only contribute to, but also benefit from, the general channels, if these were made open and public. Then the users could copy all or part of the search criteria of the general channels to improve their own channels.

6.2.1 Summary

The interviews did not only result in some fields of application. A number of areas where VIP today displays some deficiencies were also visualized. However, these are areas where the system could be used in the future, provided that further development and improvements take place. These design implications are as follows:

- VIP can be improved to indicate activities of other users.
- VIP needs better functions for training a channel, and for the possibility to clone agents.
- VIP can facilitate contacts between employees if further information about the persons displayed in the system is provided.
- VIP needs to include links to formal competence descriptions.
- VIP can better assist in the evaluation of interests if links to further information are added.
- VIP does not support historic views of information but could easily be improved to do so.
- VIP would benefit from a further developed relationship and interaction between personal and general channels.

7 Discussion

There are many different definitions of knowledge to be found within the KM context, and therefore, it is problematic to refer to one common basis for the KM research. Earlier, we stated that interest is a necessary building block of knowledge, i.e. know-what, and competence, i.e. know-how. This conviction is founded on the theories of Jürgen Habermas (1986) who means that interest is one of the cornerstones, maybe even the very foundation, of knowledge. Interest, expressed in work and interaction, is also one of the basic prerequisites for human life. Furthermore, interest constitutes the very foundation for interaction between individuals, which Habermas means is a prerequisite for the preservation of social life. Thus individuals have great power, perhaps greater than normally acknowledged, over their own knowledge and competence, i.e. how it should be developed and what direction it should take.

7.1 Aspects of Knowledge Management

It becomes more and more important for organizations to find strategies for KM in order to benefit from their employees' knowledge and competence. One strategy is to let management control and direct. Another approach is to let the employees themselves take responsibility for their own knowledge development, through the organization's support of their interests. Humans view the world in either an objective or a subjective way. These approaches reappear in the two dominating aspects of KM, i.e. the cognitive and community approaches. The dominating approach, in research as well as in commercial development, is the cognitive. This is the more objective approach, viewing knowledge as explicit and codifiable. The community approach on the other hand, focuses more on the subjective, as the elusive tacit knowledge, which is closely related to reflections and interpretations. It also has a strong focus on interaction and cooperation between individuals. As far as we know, this approach gains more and more territory as new forms of organizations, e.g. project-based and knowledge intensive, become more common. Even if interest is not emphasized in the cognitive approach its impact cannot be totally ignored. Knowledge may be concrete and objective, but we believe that subjective interest always will be one of its building blocks.

7.2 IT Solutions for Knowledge Management

In the cognitive approach, KM systems are mostly viewed as storage for knowledge and competence, and as tools for management. With this kind of systems the knowledge is already acquired when it is stored and therefore it might be difficult to determine its topicality. Traditional KM systems are often found in this category. However, one problem with these, which we find severe, is that tacit knowledge seldom is found within them. This has to do with the fact that tacit knowledge is difficult to structure and categorize. Interests are also in many cases hard to categorize, estimate and evaluate. They often overlap and therefore it is difficult to draw distinct borders between them. As stated earlier, Habermas (1986) means that interest constitutes the foundation for knowledge, but interest is not given much scope

within the cognitive approach. In the community approach, IT-based systems are regarded as means for cooperation and sharing of knowledge and competence. The spreading of knowledge is considered to take place through systems, used for e.g. advisory service, sharing of experience, or reference to information. As far as we know, there are no commercial KM systems designed to fit this approach, but there are several other techniques and systems that could be used for the same purpose, e.g. OM systems for cooperation and RS. Our study has shown that, properly adapted, these techniques could very well be used as KM systems. We especially believe RS to be a suitable platform on which to build, preferably in combination with agent technology.

VIP features the basic characteristics of agent-based RS. Its basic function is to help users stay up-to-date within one or more interest areas and it also allows cooperation between users of the system. One benefit of the latter is that the users at all levels are free to share their results with each other. This can take place in different ways, but regardless of which, an exchange of knowledge is made possible, which in turn can lead to the creation and growth of communities. The qualitative growth of the community consists of the members sharing knowledge, and quantitative growth means that new members are added. Since communication within the community often is informal, there is a great possibility to share tacit knowledge. The community members can also recommend hyperlinks and other sources of information that can be useful for the daily work. When these yield good results they will be used more often, and this can in turn lead to that the sources are further recommended to other members, and to the expansion of the community, both qualitative and quantitative. This is one of the many benefits of hybrid RS.

7.3 Fields of Application

Using VIP as a KM system may reveal new or hidden competence within the organization. Not only is structured or categorized information of interest. Information about day-to-day work and how it is done is seldom documented, but if it were known tacit knowledge would easier be passed on. The apprentice who is taught a craft by a master exemplifies this kind of knowledge and competence acquisition. Here VIP has an advantage over to traditional KM systems, since it can visualize the daily work. One beneficial feature of VIP is the possibility to see which roles different users are associated with in varying situations. However, since the information in the system is unstructured, it demands that the users can read between the lines and draw their own conclusions from the information. The possibility of using VIP to find roles and users with similar interests can contribute to the creation of communities of practice. These networks can be of varying duration, temporary as well as long term. After the community is created, it is up to the participants to decide to what extent VIP shall continue to be used and what place VIP as a medium will receive in the network. The building of networks is often considered important, but is also often taken for granted in many organizations, since they are a prerequisite for cooperation. People are often attracted to like-minded and gladly build networks with these to exchange thoughts and ideas. If only their existence was known there would be many more who would be interesting to network with. One of VIPs strengths is the ability to help find these co-workers.

One reason why we consider interest important is that it can be an essential driving force. Besides competence, driving forces are needed to produce a positive effect on the outcome of different work situations. An individual driven by commitment, motivation and interest might be better suited to accomplish a certain task, even though someone else has a higher formal competence. In VIP, driving forces can be found through identifying content and/or search parameters of the channels, the frequency of usage, number of channels with similar search parameters, etc. By identifying driving forces they can be enforced and through that better results can be achieved. Driving forces are also important to consider during configuration of project groups. Therefore VIP can be well worth consulting in such a situation.

VIP also supports interest management, i.e. how to deal with and use interests and encourage the development of them. Even if management cannot control their co-worker's interests they can pay attention to them, take an interest in them, increase awareness of them, and take advantage of them in different ways. Provided that there are no organizational restraints, the users are free to create channels to search any interest areas. This flexibility allows the channels to nourish interests without the need for more than an initial engagement from the users, since the channels continually return information matching the search criteria. Not only will this benefit the individual, but also the entire organization, since its knowledge assets will grow and mature. One way to look after organizational interests is to combine push and pull technologies. In VIP, the organization can push out messages through general channels and allow users to pull in whatever they choose to search after through personal channels. Additionally, the collective knowledge will be visualized if an aggregated picture of the channels is made possible. Through analysis of this picture and the development of interests in the organization, a picture of the overall knowledge and competence emerges. This picture is different from pictures retrieved through traditional KM systems, since VIP not only visualizes existing knowledge and competence, but also how it is being used on an everyday basis. The general channels can also be used to manage interests. Every category of general channels has an editor, who decides which information the channels should search for. Here the organization has an opportunity to create interest in its employees and push out information that they ought to see. Another way to manage interests is to print log files of all personal channels, which speak of the employees' interests. If the organization finds any of these interesting for others, the channels can be transformed from personal to general. This can contribute to the visualization, widening, and faster development of a certain interest. Through interest management individuals are encouraged to develop their interests. If a co-worker's interest can rub off on others, a higher level of motivation can be achieved within the organization. However, even if interest management is facilitated through VIP, this feature needs to be improved through further collaboration and interaction between the personal and general channels.

7.4 Distinctions from Traditional KM Systems

VIP distinguishes itself from traditional KM systems in many ways. It is primarily designed for the users of the system, i.e. has a bottom-up approach. Today, it is the individual users who benefit most from the system. Updating is not a separate activity in VIP, but takes place automatically, as soon as the system is being used. This gives

VIP an advantage, since it does not need more than a minimum of maintenance once a channel has been properly trained and consequently, there is no need for specific incentives for updating either, as long as the interests do not take a different direction. In that case, the search criteria have to be changed. Several of the interviewees compared VIP to TP/HR, which can be viewed as a traditional KM system. The two systems contain more or less the same information. Although, VIP is a more dynamic system managing unstructured information, which means that it retrieves information in a different manner. TP/HR stores the information as long as it agrees with the specification and if the information needs to be changed the system must be updated. One advantage with TP/HR is that it is easy to find historic information and perform analyses. This is not possible in VIP today, since it only shows instant information. The dynamic aspect of VIP conveys that old information may fall away if the information on Violin changes. In TP/HR roles, positions, and competence are explicitly categorized and defined. This means that the information is easily found, but at the same time, certain information can be lost if it does not fit in the predefined structure. However, such information can be found in VIP. As we mentioned before, in VIP the roles and positions are not explicitly expressed, but the daily work can be visualized. Both VIP and TP/HR have advantages and drawbacks and it is difficult to decide which serves best as a KM system. We have constantly been reminded of that the two systems can complement each other. Therefore we believe that Volvo would benefit from their collaboration.

The overall purpose of KM systems is that the organization should have constant access to timely information about the employees' competence. This demands that the information is updated continually, something that can be a difficult task. The individual employee has to give information about the own competence, either directly in the system or indirectly through another person. However, there is always a risk of the individual hiding certain information (cf. Lindgren & Wallström, 2000). This might be due to a shift in interest and that the employee no longer has a wish to work with certain tasks, even though he or she has a skill for it. Another problem with updating is the time aspect when other assignments have to be prioritized. This in turn conveys further problems, since the system has to be updated regularly so that the information stays timely.

Traditional systems for competence management are usually designed for usage from a top-down approach. Their purpose is to serve management, e.g. to support gap-analyses, control, and development. The common problem with this kind of systems is that they don't offer any natural incentives to the individual employee for updating his or her profile. VIP, on the other hand, which is an agent-based RS, is updated automatically as soon as it is being used. RS often depend on the number of users. If the number of users increases the system will hopefully yield more and better recommendations. On the other hand, if combined with agent technology, the users will not be completely depending on each other's recommendations since agents autonomously and continually collect and pass on recommendations. The users can to a certain degree control the recommendations by training of the agents, i.e. the more frequent an agent is trained the better results it will return. However, this type of system needs a critical mass, i.e. a sufficient number of individuals who will use, update and handle the information, to make the system attractive for more users. To address and overcome this problem, systems design must emphasize incentives for usage and bring forward individual as well as collective benefits (Grudin, 1994). VIP

depends on the number of active users. If they are too few, the system will not provide satisfactory results conveying that present users will lose interest in the system and recruitment of new users will be difficult. The system requires that people pay attention to it, but the paradox is that people are unwilling to pay attention unless others are using it. In order to achieve this, there are areas in VIP, which need to be improved.

7.5 Design Implications

One way to overcome the critical mass problem is to enable the system with social activity indicators (SAI), which can create awareness of what other users do (Ackerman & Starr, 1995). SAI are especially useful in collaboration situations, since individuals often want to know more about what people around them are doing before choosing with whom to collaborate. A SAI could for example be a list of rated software. The list shows who have rated certain software highly and this can, in turn, trigger other people's choice of software. Other SAI could be awareness of a colleague's participation in a seminar, the amount of people sharing in an activity or using a chat function or a project room. A further example could be a general indication of how many users who currently are logged into a system. This function exists in VIP, but does not serve any particular purpose since the system is active for all users, whether they are logged in or not. A more useful SAI in VIP could be an indication of how often a personal channel is trained or how popular a general channel is. Another way that VIP can contribute to minimizing the threshold before contacting new people is to provide links to further information about the users. These links could lead to home pages, formal CV pages, or more structured information in another system, e.g. TP/HR. This would make it easier to get in touch with new people, since the more that is known about a person the easier it is to decide whether a contact is desirable or not, which in turn will give a feeling of control over the situation.

Another improvement concerns training of the channels. After creating a channel in VIP and having received a first search result there is a possibility to refine the search, i.e. training. Marking the most interesting returned documents and performing another search can do this. Training largely depends on how experienced the user is and therefore some users with little or no experience of the procedure found it insufficient and in need of improvement. If the first search criterion is insufficient, the first search result will also be unsatisfactory, and if the user does not understand the reason for the poor outcome he or she will soon loose interest in the system. Therefore, we believe that better instructions are needed and also the possibility to use other users' channels as a basis for retraining. This could be done by visualizing contents and search criteria of the personal channels, and then enabling copying of these.

Furthermore, VIP provides instant information, but lacks the ability to save and aggregate it. Today this must be accomplished manually. The first problem is relatively easy to deal with by saving snapshots at regular intervals. However, to aggregate useful information from these snapshots is a more difficult task. First, topics and categories need to be found and then common parameters for interpretation of the data need to be decided. However, this is not easy since some of the information is tacit and unstructured, but it still needs to be solved if VIP shall be able to aggregate in an automated manner.

7.6 Prerequisites for Implementing VIP

There are a number of other factors that need to be considered before implementing VIP. Since a critical mass of users is needed, the organization must make sure that the employees are motivated and committed to use the system. One way of doing this is to guarantee the quality and validity of the on-line resources, i.e. make sure that the documents within the search domains are timely and relevant. Securing the relevance of the general channels will build trust in the system and in the organization behind it. Understanding of the system will also increase through good education and on-line help. Contrary to most traditional KM systems, VIP is designed for the users, i.e. with a bottom-up approach. This opens up new possibilities, but also demands more of the users. Since the information is unstructured, it requires that the users can learn to evaluate and draw conclusions from the retrieved material. Then it is their responsibility to act on the tacit information. Furthermore, all activities that include drawing conclusions from the personal channels have an essential prerequisite, i.e. the channels have to be named in a way that reflects their contents. If so it will be enough to scan the names to get a quick overview of the channels. Further analysis of their contents will not be necessary unless the search criteria and retrieved documents are of interest.

Furthermore, the organization must secure protection of the employees' integrity if it is to aggregate and analyze personal channels. This is especially important since the co-workers can have a feeling of discomfort when their personal search behavior is visualized. However, since the search domain, i.e. Violin, is work related this is not as delicate as the visualization of a private search behavior on the Internet, but it still needs to be carefully dealt with. There are also an additional number of organizational aspects, e.g. culture, leadership, and internal marketing, that needs to be considered, but we will not further discuss these in this thesis. However, they are important areas that could be interesting for further research.

During our study we have come across both the cognitive and the community approaches to KM. They represent two sides of the same coin and are mutually dependant. VIP supports both of these approaches. Two other opposites in the scope of our thesis are static and dynamic systems. VIP is dynamic and can be used to retrieve unstructured information, while traditional KM systems mostly focus on structured and categorized information. Therefore, we are convinced that VIP very well can be used as a KM system, especially to support the tacit dimension of knowledge and competence.

8 Conclusions

Based on our case study, we conclude that VIP can support the management of competence by the following fields of application:

- VIP can function as a complement to competence mapping by adding everyday information to the more structured competence information found elsewhere.
- VIP can support visualization and identification of roles, both formal and informal.
- Through the identification of roles mentioned above, VIP can support the creation of communities of practice.
- It is possible to identify driving forces via VIP, which in turn can aid project configuration.
- Through a search for users with certain expertise or interest, VIP can help to identify competence areas.
- VIP can support the management of interest.
- Finally, VIP can be used as a tool for competence analysis.

VIP can be improved in supporting the management of competence by the following design implications:

- By adding functions, e.g. public lists of all users and their channels, VIP can be improved to indicate activities of other users.
- VIP needs better functions for training a channel, e.g. visualization of content and search criteria of channels in order to inspire individual employees. It would also benefit from the possibility to clone channels.
- If VIP included further information about the persons displayed in the system, it would, to a higher degree than before, facilitate contacts between employees and the building of communities.
- Since VIP handles unstructured information, which demands interpretation abilities from the users, it would benefit from hyperlinks to formal competence descriptions.
- VIP could be improved to better assist in the evaluation of interests, e.g. by facilitating aggregation of information.
- VIP does not support historic views of information but could easily be improved to save instant pictures. The aggregation mentioned above could also help in this aspect.
- VIP would benefit from a further developed relationship and interaction between personal and general channels.

9 References

Literature References

Ackerman, M.S. & Malone, T. (1990). Answer Garden: A Tool for Growing Organizational Memory. In *Proceedings of ACM Conference on Office Information Systems.* (pp. 31-39).

Ackerman, M.S. (1994). Augmenting the Organizational Memory: A Field Study of Answer Garden. In *Proceedings of Conference on Computer-Supported Cooperative Work (CSCW '94)*. Chapel Hill, NC: ACM Press, (pp. 243-252).

Ackerman, M.S. & Halverson, C. (1999). Organizational Memory: Processes Boundary Objects and Trajectories. In *Proceedings of HICSS*.

Ackerman, M.S. & Starr, B. (1995). Social Activity Indicators: Interface Components for CSCW Systems. In *Proceedings of UIST-95*.

Alvesson, M. & Sköldberg, K. (1994). Tolkning och reflektion: Vetenskapsfilosofi och kvalitativ metod. Lund: Studentlitteratur.

Asplund, J. (1970). *Om undran inför samhället*. Lund: Argos.

Backman, J. (1998). *Rapporter och uppsatser*. Lund: Studentlitteratur.

Balabanovic, M. & Shoham, Y. (1997). Fab: Content-Based Collaborative Recommendation. *Communications of ACM.* Vol. 40, No. 3, (pp. 66-72).

Bannon, L. J. & Kuutti, K. (1996). Shifting Perspectives on Organizational Memory: From Storage to Active Remembering. In *Proceedings of the 29th Hawaii Conference on System Sciences (HICSS-29)*. Los Alamitos: IEEE Computer Society Press, Vol. 4, (pp. 156-157).

Bertels, T. & Savage, C.M. (1998). Tough Questions on Knowledge Management. In von Krogh, G., Roos, J. & Kleine, D. (Eds.), *Knowing in Firms* (pp. 7-25). London: Savage Publications.

Cohen, D. (1998). Toward a Knowledge Context: Report on the First Annual U.C. Berkley Forum on Knowledge and the Firm. *California Management Review*. Vol. 40, No. 3, (pp. 22-39).

Dahlbom, B. & Mathiassen, L. (1995). Computers in Context: The Philosophy and Practice of Systems Design. Oxford: NCC Blackwell.

Davenport, T.O. (1996). Information Ecology: Mastering the Information and Knowledge Environment. Oxford: Oxford University Press.

Delgado, J. A. (2000). *Agent-Based Information Filtering and Recommender Systems on the Internet*. (Dissertation Thesis). Nagoya: Nagoya Institute of Technology, March.

Easterby-Smith, M., Thorpe, R., & Lowe, A. (1991). *Management Research: An Introduction*. London: Sage Publications.

Fahey, L. & Prusak, L. (1998). The Eleven Deadliest Sins of Knowledge Management. *California Management Review*. Vol. 40, No. 3, Spring.

Feldman, S. (1999). Intelligent Agents: A Primer. *The Searcher*, Vol. 7, Issue 9, (pp. 42-56).

Foltz, P.W. & Dumais, S.T. (1992). Personalized Information Delivery: An Analysis of Information Filtering Methods. *Communication of the ACM*, Vol. 35, No. 12, (pp. 51-60).

Foner, L.N. (1997). Yenta: A Multi-Agent, Referral-Based Matchmaking System. *First International Conference on Autonomous Agents, (AGENTS 97)*.

Garofalakis, M.N., Ramaswami, S., Restogi, R. & Shim, K. (1999). Of Crawlers, Portals, Mice and Men: Is There More to Mining the Web? *SIGMOD '99*, (p. 504). ACM Digital Library.

Grudin, J. (1994). Groupware and Social Dynamics: Eight Challenges for Developers. *Communications of ACM*. Vol. 37, No. 1, (pp.92-105).

Habermas, J. (1986). *Knowledge and Human Interests.* Boston: Beacon Press.

Hansen, M. T., Nohira, N. & Thierney, T. (1999). What's Your Strategy for Managing Knowledge? *Harvard Business Review*, March-April, (pp. 106-116).

Jones, S. (1985). Depth Interviews. In R. Walker *Applied Qualitative Research*. Aldershot: Glover.

Kautz, H., Selman, B. & Shah, M. (1997). Referral Web: Combining Social Networks and Collaborative Filtering. *Communications of ACM*, Vol.40, No. 3, (pp. 63-65).

Latendre, D. (1999). Collaborative Portals Add Value to Your Intranet. *Computing Canada*, Vol. 25, No. 41, (p. 33).

Lashkari, Y., Metral, M. & Maes, P. (1994). Collaborative Interface Agents. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*. Cambridge, MA: MIT Press.

Lieberman, H., Van Dyke, G. & Vivaqua, A. (1999). Let´s Browse: A Collaborative Web Browsing Agent. In *Proceedings of IUI-99*. ACM Press.

Lindgren, R. & Stenmark, D. (forthcoming). *Information Agents as Competence Indicators*. Submitted for publication.

Lindgren, R. & Wallström, C. (2000). Features Missing in Action: Knowledge Management Systems in Practice. In *Proceedings of ECIS 2000*. (pp. 701-708).

Marshall, C. & Prusak, L. (1996). Financial Risk and The Need for Superior Knowledge Management. *California Management Review*, Vol. 38, No. 3, (pp. 77-101).

McDermott, R. (1999). Why Information Technology Inspired But Cannot Deliver Knowledge Management. *California Management Review*, Vol. 41, No. 4, (pp. 103-117).

McDonald, D. & Ackerman, M. (1998). Just Talk To Me: A Field Study of Expertise Location. In *Proceedings of CSCW '98*. Seattle, WA: ACM Press (pp. 315-324).

McDonald, D. & Ackerman, M. (2000). Expertise Recommender: A Flexible Recommender System and Architecture. In *Proceedings of CSCW 2000*. New York: ACM Press, (pp.231-240).

Nonaka, I. (1994). A Dynamic Theory of Knowledge Creation. *Organization Science*, No. 5, (pp. 14-37).

Nonaka, I. & Takeuchi, H. (1995). The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation. Oxford: Oxford University Press (Trade).

Polanyi, M. (1966). *The Tacit Dimension*. London: Routledge and Kegan Paul.

Robertson, M., Sörensen, C. & Swan, J. (2000). Facilitating Knowledge Creation with GroupWare: A Case Study of a Knowledge Intensive Firm. In *Proceedings of HICSS 33*. IEEE Press.

Resnick, P. & Varian, H.R. (Eds.) (1997). Recommender Systems. *Communications of ACM*, Vol. 40, No. 3, (pp.56-89).

Sarvary, M. (1999). Knowledge Management and Competition in the Consulting Industry. *California Management Review*. Vol. 41, No. 2.

Scarborough, H. (1998). Path(ological) Dependency? Core Competencies from an Organizational Perspective. *British Journal of Management*, Vol. 9, (pp. 219-232).

Shardanand, U. & Maes, P. (1995). Social Information Filtering: Algorithms for Automating 'Word of Mouth'. In *Proceedings of CHI 95*. Denver: ACM Press. (pp. 210-217).

Snis, U. (1999). IT Support in a Knowledge Management Process: A Field Study of a Quality Support Group in a Pharmaceutical Company. In *Proceedings of IRIS 22*, Keurru.

Stenmark, D. (2001). Leveraging Tacit Organizational Knowledge. *Journal of Management Information Systems.* Vol. 17, No. 3, (pp. 9-24).

Stenmark, D. (2000). Collaborative Aspects of Information Retrieval Tools: Summarizing Three Action Case Studies. In *Proceedings of IRIS 23*, Uddevalla.

Sveiby, K. E. (1997). The New Organizational Wealth: Managing & Measuring Knowledge-Based Assets. San Francisco: Berrett-Koehler.

Swan, J., Newell, S. & Robertson, M. (2000). Limits of IT Driven Knowledge Management Initiatives for Interactive Innovation Processes: Towards a Community-Based Approach. In *Thirty-Third Hawaii International Conference on Systems Sciences (HICSS-33)*. Maui, HI: ed IEEE.

Swan, J., Scarborough, H. &. Preston, J. (1999). Knowledge Management: The Next Fad to Forget People? In *Proceedings of the Seventh European Conference on Information Systems (ECIS '99)*, (pp. 668-678).

Turnbull, D. (1999). Interacting with Recommender Systems. *CHI*, April.

Vishik, C. & Whinston, A.B. (1999). Knowledge Sharing, Quality and Intermediation. *WACC-99*, No. 2, ACM Digital Library.

Vivacqua, A. (1999). Agents for Expertise Location. In *Proceedings of AAAI-99, Spring Symposium on Intelligent Agents in Cyberspace*, Stanford, CA. (Tech. rep. SS-99-03).

Wiedersheim-Paul, F. & Eriksson, L.T. (1997). *Att utreda, forska och rapportera*. Malmö: Liber.

Yin, R. (1993). *Applications of Case Study Research: Design and Methods*. Newbury Park, CA: Sage Publications.

Yin, R. (1988). *Case Study Research: Design and Methods* (Rev. ed.). Newbury Park, CA: Sage Publications.

Internet references

AltaVista. (2000). URL: http://www.altavista.com

Amazon. (2000). URL: http://www.amazon.com

AUTONOMY. (2000). URL: http://www.autonomy.com

Encyclopædia Britannica. (2000). URL: http://www.eb.com:180/

Firefly. (2000). URL: http://www.firefly.com

GroupLens. (2000). URL: http://www.grouplens.com

Lycos. (2000). URL: http://www.lycos.com

HotBot. (2000). URL: http://www.hotbot.com

KM Project Participants. (2001). URL: http://www.viktoria.se/groups/Knowledge Management/participants.html

Viktoriainstitutet. (2001). URL: http://www.viktoria.se

Volvo IT. (2001). URL: http://www1.volvo.com/frameset.html?url=/group/ other_volvo_companies/volvo_information_technology

Yahoo. (2001). URL: http://www.yahoo.com

Personal references

Bergqvist, M. (1999). Lecture in the class *Informatik som vetenskap*. Department of Informatics, School of Economics and Commercial Law, University of Gothenburg. Nov 18.

10 Appendices

Appendix 1 – Titles of interviewees

Appendix 2 – Interview questions

Appendix 3 – Abbreviations

Appenix 1 – Titles of Interviewees

Title	Number of interviewees
Department manager	1
Financial analyst	1
Group manager for systems architects	1
HR manager	2
Info master	2
Methods developer	1
Process developer	1
Project manager	2
Systems developer	4
Technology watcher	1

Appendix 2 – Interview questions

Introduction and background

1. Describe your education.
2. Describe your previous job descriptions, including those before your time at Volvo.
3. For how long have you been employed at Volvo?
4. What role/job description do you have in the organization?
5. For how long have you held this position?
6. How do you proceed when you look for information within an area that is new to you?
7. Why do you choose this line of action?
8. Are there alternative ways?
9. To what extent do you work in teams? (The purpose of this question is to find out if the need for a Community-function in VIP is greater when working in teams than in individual work.)
10. Do you develop more or less when working in a team than when working on your own?
11. Is there a difference in the way you look for information during working hours from how you go about it in your spare time?

Information retrieval

12. How do you act when you look for information on the Internet?
13. For what purposes do you use the Internet? Are you familiar with other areas of usage?
14. What experience do you have of search engines (good/missing functions etc.)?
15. Do you know anything about portals? (Describe, give examples, give URL:s, describe how they work.)

Experience of the Internet/intranets

16. What experience do you have of the Volvo corporate intranet, Violin?
17. How often do you use Violin?
18. How do you use Violin?
19. What distinguishes Violin from the Internet?
20. Which advantages and disadvantages do you find with Violin compared to the Internet?
21. How do you like the functions of Violin?

Experience of VIP

22. Approximately how many times have you used VIP?
23. What did you do then?
24. Why did you do it?
25. What more could you do?
26. How do perceive the purpose of VIP?

Ask the user to log in to VIP

27. Which of these functions have you used?
28. The search window in the top left corner – what is it used for?
29. *Logged in users* – in what way is this relevant to you?
30. How have you worked with your personal channels?
31. The *Personal channels* page – Do you find the recommended documents relevant? In what way?
32. The *Personal channels* page – Is the function *Summary and related documents* relevant? In what way?
33. *General channels* – are the topics relevant to you?
34. *Summary* – Can you describe what you see on this page?
35. The *Create new channel* page – How do you go about it when you fill in the field *Training*?
36. The *Create new channel* page – Do you use the *Community* function? If no – Do you know what it involves? If not – explain the purpose of it and continue: Would you consider using it?
37. The *Create new channel* page – Do you use the *Email address* function? If no – Do you know what it involves? If not – explain the purpose of it and continue: Would you consider using it? How often would you like to be notified via email?
38. The *Community* page – What do you think about the information you receive from this page?
39. The *Community* page – Has this page made you discover new people that you did not previously know had a similar interest like you? If yes – how did you act on that information? If no – how do you think you would have acted if you had discovered new people?
40. The *Community* page – Have you received the answer: ”There are no similar agents in community”? What does that mean? What did you do then?
41. The *Find competence* page – In what way is this function relevant to you?
42. The *Find competence* page – Do any of the functions here need to be altered or are any functions missing?
43. If *Find competence* does not show any users, what does that mean? How will you act then?
44. If one of your channels does not retrieve any documents, what does that mean? What will you do then?

General questions

45. Has VIP contributed to your work in a concrete manner, e.g. to get in touch with other people?
46. Do you believe that VIP has the potential to do that? If yes – in what way?
47. What consequences can VIP have for an individual?
48. What consequences can VIP have for a group?
49. Did you, at any point, experience that VIP behaved in a different manner than what you expected?

Appendix 3 - Abbreviations

CSCW	Computer supported collaborative work
CV	Curriculum vitae
HR	Human resources
IR	Information retrieval
IT	Information technology
KM	Knowledge management
OM	Organizational memory
RS	Recommender systems
SAI	Social activity indicator
TP/HR	Tieto Persona/Human Resources
Violin	The Volvo corporate intranet
VIP	Volvo Information Portal
Volvo IT	Volvo Information Technology