

## Turning Tacit Knowledge Tangible

Dick Stenmark

*Volvo Information Technology, Dept. 9530 HD1N, SE-40508 Göteborg, Sweden  
and  
IT & Organisation, Viktoria Institute, P.O.Box 620, SE-40530 Göteborg, Sweden  
it.dixi@memo.volvo.se*

### Abstract

*People are able to determine whether or not a given document is interesting just by glancing through it. However, when asked to make explicit the rules upon which such a decision is based, they are unable to do so. This, I argue, is because tacit knowledge is involved in this process. Tacit knowledge constituted the major part of the body of knowledge and it is therefore important for organisations to sustain and exploit this asset. While studying how an intranet recommender system prototype was used, I discovered how tacit knowledge in form of professional interests could be shared among the organisational members in an unobtrusive way. Based on these empirical findings, and informed by Michael Polanyi's theory of tacit knowledge, I claim that agent-based retrieval systems can be used to capture and visualise our professional interests, thus making otherwise elusive tacit knowledge tangible for others to benefit from.*

### 1. Introduction

Knowledge Management (KM) in the meaning of caring for and sharing of knowledge is really nothing new; it has been around since man first shared the knowledge of how to make fire, and has since been employed by parents teaching their children and by masters training their apprentices. In recent years, however, the importance of knowledge has risen dramatically, and gone from being one resource amongst many others to become the primary resource. Being able to manage this resource has thus received the attention of chief executives and KM as a concept has become vividly debated.

Without going too deep in the philosophical debate of what exactly knowledge is we may notice that on an epistemological level knowledge may be split along several dimensions. One way is, as suggested by Polanyi [17], to distinguish between tacit and explicit knowledge. Though others have since developed this separation further, (e.g. Blackler [3] speaks of embodied, embedded,

embrained, encultured, and encoded knowledge), I shall stay with Polanyi's definition for the scope of this paper. Thus, with explicit knowledge, I will mean knowledge that has been captured and codified into manuals, procedures, and rules, and is easy to disseminate. When using the phrase tacit knowledge, I refer to knowledge that cannot be easily articulated and thus only exists in people's hands and minds, and manifest itself through their actions.

#### 1.1. The duality of tacit knowledge

An interesting but also troublesome property of tacit knowledge is the inherent tension between its value on the one hand and its elusiveness on the other hand. The high value stems from the fact that most of the body of knowledge is made up of things we know but are unable to express. With Polanyi's words: "*We can know more than we can tell*" [18, p.136]. Leonard and Sensiper go even further by stating that "*we can often know more than we realize*" [13, p.114]. Many KM initiatives have thus focused on trying to explicate our tacit knowledge. For instance, Nonaka and Takeuchi [15] suggest that tacit knowledge becomes explicit through the process of externalisation, i.e. by sharing metaphors and analogies during social interaction. Once the knowledge has become explicit we may store it in databases and manuals.

However, such a process is both difficult and costly, and the fact that the tacit knowledge must be externalised before it can be exploited limits its usefulness. It may even be questioned whether it is at all desirable to explicate knowledge [11, 12].

The elusiveness of tacit knowledge can be derived from at least two reasons; we are ourselves not fully aware of our tacit knowledge, and there is a lack of incentive on the individual level to make it explicit. Firstly, Davenport and Prusak observe that tacit knowledge "*incorporates so much accrued and embedded learning that its rules may be impossible to separate from how an individual acts*" [8, p.70]. In our daily activities, our tacit knowledge informs our activities without us thinking of it as knowledge. We know how to ride a bike, or what

cinnamon smells like, but we cannot document it in a manual, nor explain it to others.

Secondly, our knowledge is something that resides within us, and manifest itself through our actions, and we therefore do not *need* to document it for our own sake. We just use it. Should we have to express our tacit knowledge in words, it would not be for our own sake but for the benefit of someone else in our organisation or community. Grudin [10] has argued convincingly that situations where one is forced to do the work and someone else gets the benefit very often result in failure. Leonard and Sensiper [13] further observe that our tacit knowledge may be considered a valuable competitive advantage that we would not want to share with others without getting something in return. Stenmark's [21] findings, that lack of proper reward mechanisms on the individual level will effectively hinder sharing of ideas despite potential organisational benefits, shows that this is true not only for tacit knowledge. Extensive knowledge sharing may create a situation where an organisational member has "automated away" the reasons for his or her existence in the organisation.

## 1.2. Interests as tacit knowledge

In his ethnographic study of how work is conducted, Orr [16] describes how an organisation's view of how work is carried out can contrast sharply to what it really takes to get a job done. Though we have our formal job descriptions, these are seldom enough to account for the actions we perform during a working day. Our interests as professionals often make us elaborate within, and often even outside, our role definitions. Much of our daily office activities are thus governed by professional interests that dictate what reports we read, what documents we write, what discussions we engage in, and what we search for on the web.

Suchman [22] observes that tacit knowledge enables us to take actions that are situated in particular social and physical circumstances, and that tacit knowledge thus is contextual bounded. In an office setting as the one described in this paper, our interests, and the actions they provoke, are limited to a professional context. If we could capture some of those activities and derive our underlying interests, we might be able to externalise part of our tacit knowledge, and thus make it - if not explicit - at least tangible. Such a possibility would be useful to an organisation, as it would enable the sharing of this value resource, without having to go through the costly process of fully externalising it.

I suggest that our (professional) interests are instances of tacit knowledge. Though we may be unable to produce an exhaustive definition of our interests, we usually have no problem in determining whether or not any given document is interesting. Therefore, we intuitively know

what we are interested in when we see it but we are unable to make our interests explicit for others to learn. I will later show how we may apply Polanyi's theories to explain this phenomenon.

## 1.3. Retrieval systems from a new perspective

While studying the usage of an agent-based web retrieval prototype I observed unexpected user behaviour which led me to do further investigations. Based on the studies I claim that such systems can provide the mechanism that allows us to solve the two problems mentioned above; it helps articulate tacit knowledge, and; it creates a natural incentive to try to do so. By identifying certain documents as interesting, an agent-based retrieval system could maintain a dynamic profile that represents some of my tacit knowledge without requiring explicitly defined keywords or manually updated records. Since this profile is used to provide me with information that is more accurate and search results that are more precise, a natural incentive exists for me to give feedback and thus cultivate the profile. The resulting profile, I will try to show, represents part of my tacit knowledge, which in a sense thus becomes tangible. The profiles enable the organisation to find people with a certain interest.

Cohen *et al.* [6] take a similar but reversed approach in their Expert Browser, when they note that experts read web documents and that this is an indication of that the document in question is relevant within a certain field. Others may follow the path of the expert to find useful information. A prerequisite is that the expert (or a group of possible experts) is known. Our approach is instead to follow the *interest* that the documents represent to find the otherwise unknown expert.

Research concerning agent-based retrieval systems has focused mainly on user-to-object or user-to-information objectives, but has sometimes also addressed the user-to-user considerations. (See e.g. [19] or [9] for references to recommender systems and their implementations.) No one, however, has approached agent-based retrieval systems from a knowledge management perspective; i.e. discussed what knowledge governs the individual activities and how tacit knowledge may be put to use in the community. My work contributes by proposing an interpretation that explains how tacit knowledge is activated, and how it may be made tangible in an organisational setting. The research described herein is thus not about recommender systems *per se*. I have studied people using technology rather than the technology itself. However, the way in which the recommender system prototype was implemented help explain the findings and I will thus briefly describe some of the main features in the tool used. The choice of tool

was however not significant for the research - any equivalent product would have served the purpose.

#### 1.4. How this paper is organised

In the next section I will explain my research methodology and describe the domain in which I performed my empirical study. I will then refer to some text about knowledge in general and tacit knowledge in particular before briefly describing the fundamental features of the prototype. Thereafter I continue with a report of the users' experiences and comments, followed by a discussion of these findings. I conclude the paper with some conclusions.

## 2. Research settings and method

The empirical fieldwork took place at Volvo Information Technology, an IT service company within the Volvo Group. Volvos intranet consisted of some 450 web servers and had approximately 400,000 documents. Most of the content was official or semi-official information, such as department presentations, project reports, Frequently-Asked-Questions (FAQ's), and online help material.

In my research, I subscribe to a grounded theory-inspired approach, meaning that I let my empirical findings form the hypothesis on which my analysis is built. This is an iterative process during which the discovered theory is gradually refined until it comprises all observed cases.

I invited approximately 80 users to participate in the study, which ran from August to November 1998. The incentive to participate came from the assumption that the prototype being tested would be able to provide them with more targeted information at a lower user effort.

The interested users were invited to a 2-hour introduction meeting, where I explained the purpose of the research, the concept of agent-based systems, the design of the application and how to operate it, how to register and login, and how to set up and run individual agents. I also asked the participants to keep informal records of particular incidents that they considered worth noting, and informed them that I was going to contact them during or after the test to collect their viewpoints. The seven users that were unable to attend either of the three introduction meetings received the above information via email. Most, but not all, of the 48 users who actually registered and participated in the test were Volvo IT employees and their job descriptions varied from technicians and system developers to content providers and administrators. All were experienced computer users.

User experiences as well as hard data have been collected in several ways, including both interviewing,

questionnaires, and log file analysis. First, all users were invited to a group interview but only eight showed up. Certain emerging patterns could however be noticed and a first tentative theory was formed. The remaining 40 users were then sent an email questionnaire, which again only some (12) answered. After re-looping the analytic phase, based on the so far received answers and the application log files, I conducted seven semi-structured qualitative interviews which shaped the final conclusions reported herein. The interviews were open-ended and lasted between 28 and 66 minutes.

## 3. A look at knowledge epistemology

I will argue that Polanyi's theories of tacit knowledge can explain the mechanisms that enable recommender systems to help visualise or interests. However, tacit knowledge has in Western history not received as much attention as has the - at least in some respects - more obvious explicit knowledge, so before going into tacit knowledge, we shall first position ourselves by making a small cultural and historical journey.

### 3.1 Knowledge culture and history

Nonaka and Takeuchi [15] point out that Western philosophy has, unlike that of Japan, a long tradition of knowledge epistemology. They describe how the philosophical debate, departing from the work of Plato and Aristotle, goes on via the likes of Descartes and Locke, and later Kant, Hegel, and Marx until today. What has characterised the Western tradition is the separation of the knowing subject from the known object. Descartes gave this separation its methodological basis when he declared the "Cartesian split" between the mind and the body. This Cartesian split has had a profound influence on Western social sciences, including economics, management, and organisational theory. According to Nonaka and Takeuchi, this split explains, at least partly, why Western view largely has focused on *explicit* knowledge and looked upon learning as something performed with the mind.

### 3.2 Tacit knowledge

When developing their hypothesis regarding knowledge creation, Nonaka and Takeuchi [15] draw heavily on the distinctions between explicit and tacit knowledge. They base much of their discussion on the work of Polanyi, though, as pointed out by Leonard and Sensiper [13], Polanyi assumed knowledge to exist on a spectrum where all knowledge has a tacit dimension. In this article, however, I shall concentrate on the tacit end of

this spectrum and elaborate on what constitutes the major part of the body of knowledge.

Choo [5] observes that Schön's [20] definition of knowing in action describes the characteristics of tacit knowledge:

- There are actions and judgements, which we spontaneously know how to carry out without having to think prior to or during their performance.
- We often find ourselves doing these things without being aware of having learned them.
- The knowledge revealed by our actions can usually not be described, and we may never have been really aware of it.

Further, Choo concludes that tacit knowledge is "*distributed in the totality of the individual's action experience*" and that tacit knowledge is "*relying on tactile cues registered by the human body interacting with its environment*" [5, p.117]. To better understand these statements we may turn to Polanyi.

Polanyi claims from his research that tacit knowledge has two distinct properties, which he names its *proximal* and *distal* terms [17]. The proximal term is the part that is closer to us, while the distal part is further away. Polanyi exemplifies by describing how the police help a witness who is unable to describe a perpetrator to create a phantom picture by selecting pictures from a large selection of human features such as eyes, noses, and hair. By attending *from* the first, closer image that resides within, *to* the second, more distant picture collection, the witness is able to communicate her awareness of the face. Similarly, Polanyi refers to an experiment where a person was presented with a large number of nonsense syllables and after certain syllables, the person was given an electric shock. The person was able to anticipate the shock at the sight of the shock syllables but on questioning remained unable to identify them. Again, by attending to the distal term - the shock - the test person became aware of the proximal term - the shock association.

Tacit knowledge is, argues Polanyi, the understanding of the unity that this proximal/distal pair together constitutes. We become aware of the proximal term only in the presence of the distal term - the tactile cue - but remain unable to communicate the former. Should we indeed focus our attention on the proximal term we may become temporarily paralysed and unable to perform the task at hand. For instance, a pianist may know a complicated musical piece by heart and be able to play it without concentrating on where to put his hands. However, should he start to think about where each finger should go - i.e. should he attend to the proximal term - he would soon find himself unable to continue playing. Hence, we should not attend to the particulars in themselves but be aware of them as the entity they

constitute. Only by dwelling in them, we can understand their joint meaning, according to Polanyi.

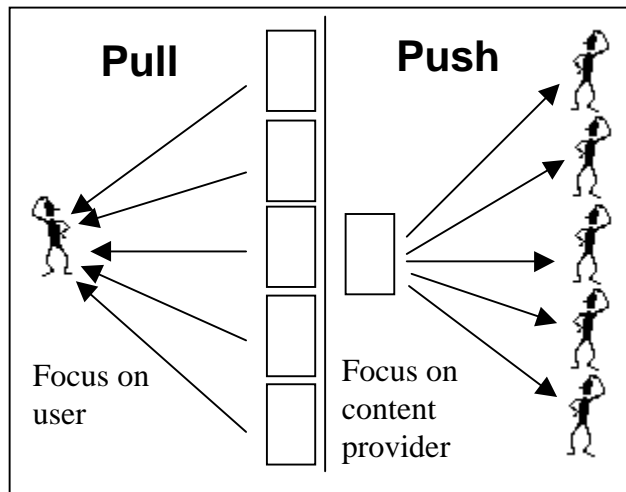
Despite not being easily expressible in words, tacit knowledge may be both shared and taught. Choo [5] explains that through rich modes of discourse including analogies, stories, and metaphors, tacit knowledge may be revealed. Hansson [12] agrees and claims that tacit knowledge is not at all "tacit" - it just expresses itself in another form than spoken language. Applying Polanyi's theories here, we see that when attending *from* our interests - the proximal term - and attending *to* the document - the distal term - we are able to recognise and express our interests. Through documents, the tacit knowledge of our interests may thus be communicated. All we need now is a mechanism that allows us to both identify documents and share this insight with others. I suggest that a web-based recommender system provides such a mechanism.

#### 4. Prototype system features

The aim of this research was not to develop and study agent-based retrieval systems per se, but rather to examine how such a technology could be used in a new and innovative way. To speed up the development process I therefore wanted to build on existing software tools. While examining the commercial tools available at the time, I came to realise that there were two different perspectives on how content was handled and what role the user or customer had. These two views may be labelled *Push* and *Pull* respectively.

Push-oriented products focus on the content providers and how the site owners can best deliver added value to the customers. Though being able to adapt to user behaviour and learn to recognise user preferences, this is done in order to help the content provider primarily. Since every server wanting to have this feature must have the appropriate software installed, this solution works best when a single web server is used. For example, when Amazon uses push-based technology to recommend books or music, it only recommends books and music from the Amazon site. It does not provide references to competitors. However, for an intranet, this approach is less useful.

A pull-oriented product, on the other hand, takes departure in the users' needs and pulls whatever information it can find that matches the users' interests, from any web server in the net, and delivers it to the users' browsers. No modifications to or restructuring of existing data is needed and no additional software has to be installed on the web servers. Given the objectives of my research I considered pull-oriented technology more suitable. Figure 1 below illustrates the difference between these two perspectives.



**Figure 1.** Pull sets the user in focus and collects information from a variety of sources making it suitable for an Intranet. Push focuses on the content provider, sending information from a single server to many different (Internet) users.

#### 4.1. Basic recommender system features

Recommender systems are able to anticipate what items a user is likely to be interested in and can thus, in a hopefully intelligent way, recommend such items. The details on how this “anticipating intelligence” is implemented vary from product to product and are not relevant to this paper. Academic research as well as the success of commercial products has shown that such systems do work and we may safely assume this to be true also in this particular case. The pull-oriented tool used to implement my prototype could reduce web documents to 0.5K digital signatures or “fingerprints”. By using proprietary neural net and pattern matching algorithms, the product could identify documents similar to each other. It also provided features for creating agents that could find documents matching the users’ interest, and re-training, e.g. adapting the agent to a set of documents marked as relevant by the users. See [2] for details.

#### 4.2. Design decisions

The prototype was designed and implemented to support the following:

1. Offer individual agents that could be set to find intranet documents based on a user profile, e.g. a richer representation of an interest than merely a keyword-based query.
2. Provide mechanisms to enable refinement of the agents based on positive user feedback on retrieved documents.

3. Enable users to locate colleagues with similar assignments and organisational roles by matching user job descriptions and thus facilitating the creating of online communities.
4. Display users with similar interests by matching their agents.

Below, these features are described in more detail.

**4.2.1. Creating agents.** Since each user was to be offered personalised agents, they had to identify themselves by logging in. Once inside, the users could create an agent, give it a name, and assign it a task. The task corresponded to a search engine query, but was not limited to simple keywords and Boolean connectives since many users found such syntax difficult. Instead, the task was to be expressed in natural language and the best results were achieved when the users cut and pasted (a large chunk of text from) a relevant document and asked the agents to find more similar documents.

After the agent was created, it was represented on the screen as an icon. For each agent the users have four options; delete it, edit it, find similar agents, or check the result.

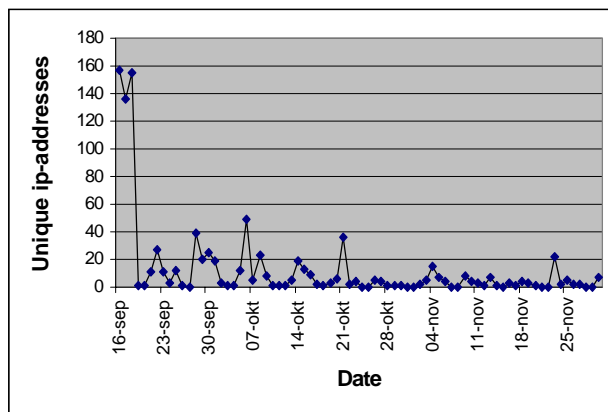
**4.2.2. Agent refinement.** The search results from the agents were displayed in a simple list, similar to those generated by most search engines, and by clicking on the associated hyperlinks the documents were retrieved. When the user had read and verified that one or more of the returned documents were indeed relevant, the user could provide the agent with positive feedback by marking the document(s) and clicking the retrain button. The digital signature of the agent was then merged with the signature(s) of the selected document(s) and the result became the new agent signature, replacing the previous one.

**4.2.3. Profile-based communities.** New users were also suppose to create a user profiles in which they were to describe their professional interests, job roles, or work responsibilities in a free text fashion. If a user already had a CV stored elsewhere, copying it into this field would have been a good idea. The profile, once saved and stored, was then converted to a digital signature, and when clicking on the Community button, the user profile signature was matched with that of other users and the resulting users were listed on the screen. The user could now display the email address or the profile of any found user by clicking the corresponding hyperlink, and had the opportunity to contact him or her. My intention with this feature was to make the users aware of each other’s presence and thus facilitate the emergence of online communities.

**4.2.4. Finding similar agents.** The Similar Agents feature was a rather late idea added more or less because it was rather unproblematic to implement. I actually planned to let the users be able to search for and find similar agents to have them cloned by copying them to their own private area. This, I believed, would help new and inexperienced users to get their agents to a decent quality level quicker. However, this extra functionality was not implemented due to lack of resources and the only feature offered to the users during the test were the option to find other users with similar agents. The result of such a search was a list of users; identical to the list produced by the Community feature.

## 5. Results and user experiences

The week immediately following the initial introduction, where all participating users were provided with user-id, password, and the URL of the prototype, the usage was high. During the following two weeks, usage declined slightly before settling on a stable level. This level was then maintained throughout the rest of the test.



**Figure 2:** Number unique ip-addresses accessing the prototype per day during the test period.

Since the test site used Dynamic Host Configuration Protocol (DHCP), by which each user receives a dynamically generated ip-address at each logon, the number of unique addresses is higher than the number of actually registered users. Figure 2 still shows that the users were most active in the beginning of the test period.

Another pattern observed from analysing the log files was that the users typically used the application often, sometimes heavily, during a day or two. They have then stayed away for some while before returning for the next session.

### 5.1. Mixed results and feelings

Overall, the user reactions were positive. All 27 responding users claimed the prototype to be useful or at least potentially useful. The respondents said they believed in this technology and considered it to be “*an extremely important asset*” with a “*great potential*”. The most frequently reported reasons for these beliefs were that it was “*easier to construct queries*” and that it “*saved time not having to search*”. One user put it this way; “*In the future we’re gonna be bombarded with even more info and this may be the only way to stay ahead*”. Eight users did however not consider the prototype useful in its current state, but they believed that a future version would probably be able to deliver.

More specifically, seven users explicitly expressed their appreciation of not having to come up with descriptive keywords, since “*they do never fully contain the meaning you have in mind any way*”, to use the words of one interviewee.

Despite the general claims that these sorts of retrieval agents were welcomed and appreciated, many users had experienced mainly negative actual results. A majority of the users (15 of 27) reported what they referred to as “*strange*” or “*unexpected*” document matches. “[It is] *hard to get something useful out of it. After retraining it with relevant documents it comes up with nothing*” as one user put it. However, the users tended to blame these bad results on their own inability rather than on the application. One user having received very little useful information said, “*The rather shallow results may depend on me not using the right words. Otherwise, I like the idea. Keep improving!*”

### 5.2. Little interest for static job descriptions

Not many of the test users exploited the Community feature, i.e. the feature that locates user with similar work-description profiles. The reasons given for not using the function were that the users already knew enough people doing similar jobs or that most users with similar profiles worked at the same department as the respondents. The respondents were not too interested in finding like-minded colleagues. As one user put it “*What’s the use of hooking up with people doing the same stuff I do? [...] It would probably be better to team up with those who knows stuff I don’t know.*” This last opinion was shared by several users, who suggested that an opportunity to search for users with complementary profiles would have been more useful. Those who actually did try the Community feature used it only once or, in one case, twice. All interviewees but one considered the Community feature to be working, or to use their words; it delivered what it was supposed to do. One user, however, claimed to have been connected to people with whom he had nothing in common. This was

not what he had expected and his reaction to it was rather negative. “*This was clearly a bug*” were his words.

### 5.3. Actual similarity more appreciated

The Similar Agents feature, which based on a given agent located users who had defined similar agents, was however much more frequently used, though it generated much the same sort of knowledge as did the Community feature. “*It’s really interesting to see who else is searching for this sort of things*”, one of the users commented. Many respondents reported that they were surprised to find certain people sharing their interests (6 users), or that the Similar Agents feature returned users whom they had not expected to be interested in a particular topic (4 users). In contrast to the Community feature remark, these comments were not uttered in a negative way. On the contrary, the users regarded these results as useful new insights and no one questioned the correctness of the results.

## 6. Discussion

The users clearly preferred exemplifying their interests by pointing to relevant documents rather than having to invent clever keywords. This, I argue, is because the former involves tacit knowledge while the latter requires a translation to explicit knowledge. However, to many organisational members the use of keywords is the established way of searching and they have difficulties trying to re-think. Despite the instructions to use entire documents as query input they continue to type in (a few) keywords. This suggests that the system should more actively encourage and facilitate the use of documents rather than keywords - possibly by letting the user enter a URL instead of text. This would prevent the user from entering keywords only.

In analogy with the face description and the shock association described by Polanyi, I suggest that our interests constitute the proximal term of our tacit knowledge. In the presence of the document - i.e. the distal term as previously the picture cards and the syllables - we are able to attend from the first term to the second. I further suggest that this distinction between tacit and explicit knowledge also explains the different ways in which the Community feature and the Similar Agents feature were used.

### 6.1. Espoused theory in formal job descriptions

The Community feature, which is based on explicit knowledge, has not been used much at all. What I believe to be the most feasible explanation for this is the fact that the Community feature was built on static profiles

provided by the users themselves to mirror the official responsibilities placed upon them by the organisation. People are often viewed as performing their jobs according to their formal job descriptions though everyday practice provides evidence of the opposite ([4], [16], and [1]). Argyris and Schön [1] refer to the worldview and values that people believe their behaviour is based on as “*espoused theory*” as opposed to “*theory-in-use*”. The organisational structure and the department descriptions that are not only already known to the members but also experienced as fictitious depict the espoused theory of work. The users rightly or wrongly assume that they know what the Community feature will return and they dismiss it as of little interest.

### 6.2. Practice considered more relevant

The Similar Agents feature is different from the above in that it does not rely on static profiles provided to describe an official role. Instead, Similar Agents relies on the tacit knowledge of our interests, made tangible through dynamically retrained agents created with a totally different purpose than the static profiles. If the prompt “*Enter your profile*” connotes a question equivalent to “*what is your official job description?*” the agents are instead created for personal benefit only and no official considerations are taken into account. True and real interests govern the choice of topics, which makes these search profiles more “*believable*” than the previous job describing ones. The most notable observations from the interviews are that when matching job profiles built on explicit knowledge and espoused theory of work, the user being linked to unexpected colleagues referred to the result as “*strange*” in a negative meaning. At the same time, the users matching agents built on tacit knowledge and practice commented similar results as “*interesting*” in a positive meaning. The tacit theory-in-use is obviously regarded as more trustworthy.

In a future version of this prototype system, an added feature would be the possibility to explicitly search for a specific competence by entering a description and match it against both agents and user profiles. This will enable the organisational members to find competence based on people’s tacit knowledge.

### 6.3 Avoiding explication of tacit knowledge

A weakness with Nonaka and Takeuchi’s [15] rather well-referenced model of knowledge creation and sharing is that it largely ignores the fact that knowledge is a competitive resource not only on the organisational level but also on individual level. People do not share knowledge without a strong personal motivation, and they would certainly not give it away without concern for what

they may gain or lose in the process. This problematic circumstance may be avoided by an approach where knowledge does not have to be externalised.

Instead, I suggest that the *whereabouts* of the knowledge should be identified and explicated. In such an approach, the users' value will increase, both for themselves and for the organisation. When the risk of being tapped of their knowledge and replaced by a database is gone, the users' reluctance to contribute is reduced. With this suggestion, I adopt the viewpoints of Davenport *et al.* [7], who suggest the introduction of information politics - that collective knowledge of the organisation is worth managing, but not necessarily worth capturing. A similar stand is also expressed by Hansen *et al.* [11] when they conclude that the management strategies for knowledge should be informed by the nature of the business. When mainly tacit knowledge is used to solve problems, the face-to-face approach to communicate knowledge should be facilitated, rather than an attempt to store it. Trying to externalise tacit knowledge can then lead to serious problems since the nuances and details that are exchanged in physical interactions are lost.

### 6.3. Problems yet to be addressed

However, the approach suggested in this article has certain shortcomings that remain to be solved. For instance, McDonald and Ackerman [14] point out that many recommender systems do not distinguish between different *levels* of knowledge. This is a weakness also in this approach. There is no way of telling whether a user with an interest is an experienced expert or just a curious novice.

Further, interests are in themselves rather elusive while knowledge remains much longer. Interests may shift over time but that does not imply that the knowledge is gone. A senior C++ programmer with a corresponding interest may develop an interest in Java programming, and eventually focus entirely on this new field. As the interest profile shifts it would then not be possible to identify this user as a C++ expert.

One must also remember that being able to find this sort of knowledge is only a first step. It helps identifying people within the organisation but it does not prevent these people from leaving the organisation nor guaranteeing that they will have time to share their knowledge on request. Davenport and Prusak observe that "*mapping who knows what in an organization creates an essential knowledge inventory, but does not guarantee the ongoing availability of knowledge*" [8, p.81].

## 7. Conclusions

I have argued that agent-based retrieval system technology could act as a facilitator in the knowledge managing process of capturing tacit knowledge on an intra-organisational web and making it tangible. The benefits of such an approach are mainly two. Firstly, the otherwise hard to solve problem of being able to produce an exhaustive definition of one's interests is replaced with the much simpler task of determining whether a given document is interesting or not. This I claim is because the latter is based on tacit knowledge. This is also illustrated by the fact that users prefer pointing to documents rather than inventing keywords when searching the web. From these findings, we can derive the following proposition:

- *P1: Retrieval systems should support a richer means of query representation than just keywords.*

Secondly, since a good profile results in more accurate information, a natural incentive to set up and maintain the profile exists. It may be claimed that the explication of tacit knowledge is both difficult and costly, and not always desired. Knowledge does not move without motivating force, and people will not give away valuable possessions such as knowledge without concern for what they may gain or lose in the process. My approach does not require the users to "automate away" the reason for their existence in the organisation. Rather, it helps identify the user as a resource. This leads up to the following propositions:

- *P2: Recommender systems may function as a visualiser of tacit knowledge without having the users to explicate away their competitive advantage.*

Previous research on agent-based retrieval system has studied how to connect users with information or users with other users. While this study confirms earlier findings that both these goals may be achieved simultaneously, it also introduces a third, and until now unnoticed, aspect of agent-based retrieval systems. My contribution is the suggestion that profiles based on the tacit knowledge of our interests and identified by practice are conveyed as trustworthier than the espoused theory-based job descriptions. The former profiles can be used to facilitate the externalisation of tacit knowledge in form of user interests to searchable knowledge in a low-intrusive way. This allows us to formulate a third and final proposition:

- *P3: Profiles based on tacit knowledge that are identified by practice are considered trustworthier than the espoused theory-based job descriptions.*

## 8. Acknowledgements

Thanks are due to Volvo Information Technology for their support during this research and to Autonomy Nordic

for letting me use their software. I am also grateful to Jan Ljungberg and Magnus Bergquist at the Viktoria Institute, and to anonymous HICSS reviewers for useful comments on earlier versions of this paper.

## 9. References

- [1] Argyris, C., and Schön, D., *Theory in practice: Increasing professional effectiveness*, Jossey Bass, San Francisco, CA, 1974.
- [2] Autonomy Technology White Paper, available from their web site at <http://www.autonomy.com/tech/whitepaper.pdf>, 1998.
- [3] Blackler, F., "Knowledge, Knowledge Work and Organizations: An Overview and Interpretation", *Organization Studies*, vol. 16, no. 6, 1995, pp 1021-1046.
- [4] Brown, J. S., "Internet technology in support of the concept of 'communities-of-practice': the case of Xerox", *Accounting, Management and Information Technologies*, no. 8, 1998, pp. 227-236.
- [5] Choo, C. W., *The Knowing Organization*, Oxford University Press, Oxford and New York, 1998.
- [6] Cohen, A. L., Maglio, P. P., and Barret, R., "The Expert Browser: How to Leverage Distributed Organizational Knowledge", paper presented at the Workshop on Collaborative Information Seeking at CSCW'98, Seattle, WA, 1998.
- [7] Davenport, T. H., Eccles, R., and Prusak, L., "Information Politics", *Sloan Management Review*, Fall 1992.
- [8] Davenport, T. H. and Prusak, L., *Working Knowledge: How Organizations Manage What They Know*, Harvard Business School Press, Boston, MA, 1997.
- [9] Fagrell, H. and Ljungstrand, P., "Make an Agent and you shall find: An Intranet Recommender System", In N. J. Buch *et al.* (Eds.), *Proceedings of IRIS21*, Volume 1, Department of Computer Science, Aalborg University, Denmark, 1998, pp. 197-206.
- [10] Grudin, J., "Social Evaluation of the User Interface: Who does the Work and Who gets the Benefit?", In Bullinger, H.-J. and Shackel, B. (Eds.), *Proceedings of INTERACT '87*, Elsevier Science Publishers, North-Holland, 1987.
- [11] Hansen, M. T., Nohria, N., and Tierney, T., "What's Your Strategy for Managing Knowledge?", *Harvard Business Review*, March-April 1999, pp. 106-116.
- [12] Hansson, H., *Kollektiv kompetens* (Collective Competence, in Swedish), PhD Thesis, School of Economics, Göteborg University, Sweden, 1998.
- [13] Leonard, D. and Sensiper, S., "The Role of Tacit Knowledge in Group Innovation", *California Management Review*, vol. 40, no. 3, 1998, pp. 112-132.
- [14] McDonald, D. W. and Ackerman, M. S., "Just Talk to Me: A Field Study of Expertise Location", *Proceedings of CSCW'98*, ACM Press, Seattle, WA, 1998.
- [15] Nonaka, I. and Takeuchi, H., *The Knowledge-Creating Company. How Japanese Companies Create the Dynamics of Innovation*, Oxford University Press, Oxford and New York, 1995.
- [16] Orr, J., *Talking About Machines: An Ethnography of a Modern Job*, Cornell University Press, USA, 1996.
- [17] Polanyi, M., *The Tacit Dimension*, Routledge and Kegan Paul, London, UK, 1966.
- [18] Polanyi, M., "The Tacit Dimension", In Prusak, L. (Ed.) *Knowledge in Organization*, Butterworth-Heinemann, Boston, MA, 1998.
- [19] Resnick, P. and Varian, H. R. (Eds.), "Recommender Systems", *Communications of the ACM*, vol. 40, no. 3, 1997, (entire issue).
- [20] Schön, D. A., *The Reflective Practitioner: How Professionals Think in Action*, Basic Books, New York, NY, 1983.
- [21] Stenmark, D., "Asynchronous Brainstorm: An Intranet Application for Creativity", Shortpaper in *Proceedings of WebNet '99*, AACE Press: Honolulu, Hawaii, 1999.
- [22] Suchman, L., *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge University, 1987