Using Intranet Agents to Capture Tacit Knowledge

Dick Stenmark

Volvo Information Technology, Department 9530 HD1N, SE-40508 Göteborg, SWEDEN,
phone: +46 31 765 7037, email: it.dixi@memo.volvo.se

IT & Organisation, The Viktoria Institute, Göteborg University
P. O. Box 620, SE-40530 Göteborg, SWEDEN, phone: +46 31 773 5566

Abstract: People can easily determine whether or not a given document is interesting just by glancing through it. However, when asked to explicitly list the rules upon which such a decision is based, they are unable to do so. This is because our personal interests are examples of tacit knowledge. I have implemented a prototype agent-based retrieval system, capable of finding web documents matching a user's interests and connecting users with similar profiles. Based on empirical findings from this prototype I claim that agent-based retrieval systems may be used to capture our professional interests, thus making otherwise elusive tacit knowledge tangible. By using agent-based retrieval systems two otherwise troublesome obstacles are avoided: interests may be defined by examples instead of specific keywords, and there now exists a natural incentive for creating and maintaining a search profile.

1. Introduction

Epistemologically, knowledge can be split along several dimensions. One way, as suggested by [Polanyi 1966], is to distinguish between tacit knowledge and explicit knowledge. Though others have since explored this separation in greater detail, I shall stay with Polanyi's definition for the scope of this paper. By explicit knowledge, I mean knowledge that has been captured and codified into manuals, procedures, and rules, and when using the phrase tacit knowledge, I refer to knowledge that cannot be easily articulated and thus only exists in people's bodies and minds.

An interesting but also troublesome property of tacit knowledge is the inherent tension between its value and its elusiveness. 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. Michael Polanyi explains: “We can know more than we can tell” [Polanyi (1966) 1998]. One of the objectives of Knowledge Management is to bridge the gap between tacit and explicit knowledge, and any technological solution that could assist in this process would be highly appreciated.

The elusiveness of tacit knowledge can be derived from at least two reasons; we are ourselves not fully aware of it, and there is a lack of inventive on the individual level to make it explicit. First, tacit knowledge “incorporates so much accrued and embedded learning that its rules may be impossible to separate from how an individual acts” [Davenport & Prusak 1997, p.70]. In our daily activities, our tacit knowledge governs 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.

Second, our knowledge is something that resides within us, and manifests itself through our actions, and we therefore have no reason to document it. 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 1987] has argued convincingly that situations where one is forced to do the work and someone else gets the benefit very often result in failure.

Our interests are an instance 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. Yet our professional interests govern much of our daily office activities. It dictates what reports we read, which documents we write, which discussions we engage in, and what we search for on the web. If we could somehow capture some of those activities and derive our underlying interests, we might be able to externalise some of our tacit knowledge, and thus make it - if not explicit - at least “touchable”. Such a

possibility would be useful to an organisation, as it would enable people to locate knowledge instead of only information.
I argue that agent-based web retrieval systems can be used to solve the two problems mentioned above; they help articulate tacit knowledge, and they create an 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, as I will show, represents part of my tacit knowledge, which in a sense becomes tangible.
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. [Resnick & Varian 1997] or [Fagrell & Ljungstrand 1998] for references to various 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 or how tacit knowledge may be put to use in the community. My work contributes to the field of KM research by proposing an interpretation that explains how tacit knowledge is activated, and how it may be made tangible. Since my work is focused on the usage of information technology (IT), I have studied how people interact with IT rather than the technology itself. To study whether agent-based retrieval systems could assist us in defining our true interests, I needed an intranet retrieval application prototype. Therefore I used a commercially available tool to implement the prototype, but the choice of tool was not significant for the research - any agent-based product would have worked.
In the next section I will describe the domain in which I performed our empirical study and explain my research methodology. I will then present the fundamental features of the recommender system prototype before reporting the results. In the last two sections I discuss the results and suggest more general interpretations and conclude with a summary.

2. Domain and Research 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.
I invited approximately 80 users to participate in the study, which ran from August to November 1998. 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. The 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.
User experiences have been collected in several ways. All users were invited to a group interview but only eight showed up. The remaining users were then sent an email questionnaire, which again only some answered. After a first analysis, based on the so far received answers and the application log files, I conducted seven semi-structured qualitative interviews, each lasting between 28 and 66 minutes. This data was again analysed, using the grounded theorising approach suggested by [Hammersley & Atkinson 1995], where the body of data is used to generate concepts which in turn helps develop typologies and theories.

3. Choice of Tool and Design Decisions

Recommender systems are able to anticipate which items a user is likely to be interested in and can thus, in a hopefully intelligent way, recommend such items. How this “anticipating intelligence” is implemented varies from product to product and is 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 in this

particular case. The aim of my research was not to develop and study new agent technology *per se*, but rather to examine *how* such a technology could be used in an innovative way. To speed up the development process I therefore wanted to build on existing software tools, if such were available. The product used in the research used neural networks and advanced pattern matching techniques to identify text patterns in profiles and look for similar patterns in other profiles or web documents. Each web document was synthesised into a 0.5K digital representation, a "fingerprint", and the characteristics that give the text meaning were determined. Once the fingerprint signature was created, the reasoning engine could perform concept matching (e.g. finding documents relevant to each other), agent creation (e.g. setting up agents that can find relevant documents), and agent retraining (e.g. adapt the agent to a set of relevant documents).
The prototype was designed and implemented to support the following features:

- Offer agents that could be set to find documents based on user profiles, e.g. a richer representation of an interest than just a keyword-based query.
- Provide mechanisms to enable retraining of the agents based on positive user feedback on retrieved documents.
- Enable users to locate colleagues with similar job descriptions or organisational roles by matching user profiles.
- Display users with similar interests by matching their agents.

Since each user was to be offered personalised agents they had to identify themselves by logging in. In addition to personal agents we also decided to offer *general agents*, i.e. publicly available agents pre-trained by us. [Figure 1a] depicts the log-in screen with the three general agents to the left and the private section to the right.

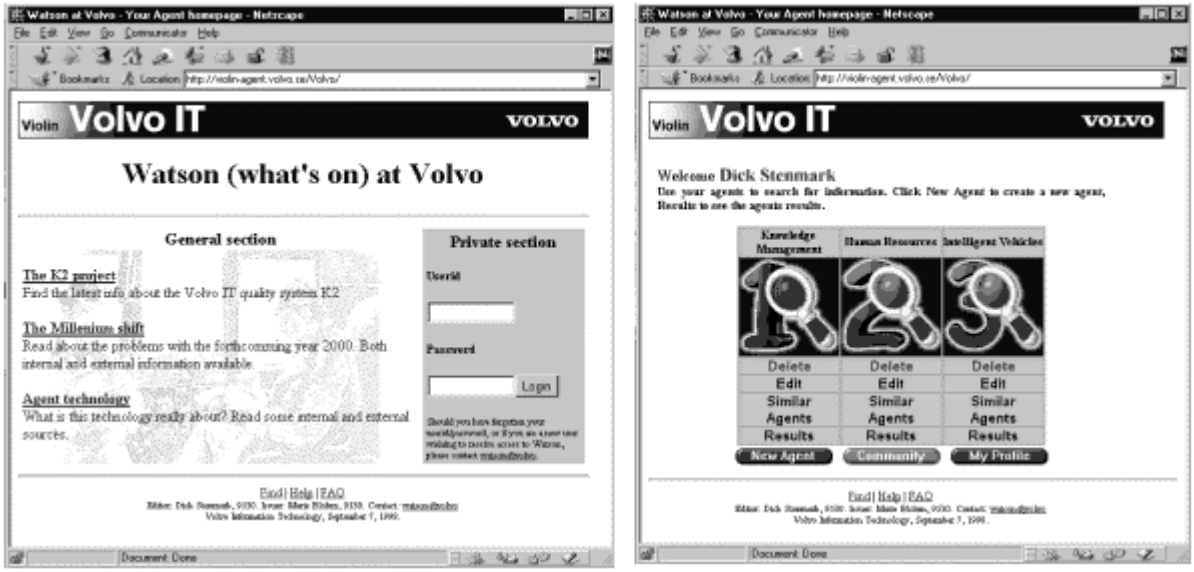

Figure 1: a) Watson's start page with the three general agents (the K2 project, the Millennium shift, and Agent technology) and the private login section. b) The individual agents available for inspection. The user may also create new agents, browse the community, or update the profile.

Once logged in, the users can create a *New Agent*, try to locate other employees with similar profile by invoking *Community*, or create/update their profiles by clicking *My profile*. When creating an agent the user gives the agent a name and assigns it a task. The task, equivalent to a search engine query, is not limited to keywords and Boolean connectives but may, and should, be expressed in natural language. In fact, the more text entered the better. The best result is achieved when the user cuts and pastes a large chunk of text from a relevant document and asks the agent to find more similar documents. After the agent is created, it is displayed on the screen. The example in [Figure 1b] shows a user with three agents; Knowledge Management, Human Resources, and Intelligent Vehicles. For each agent the user can delete it, edit it, find similar agents, or check the result. The number of individual agents was limited to five per user.
One of the first things a new user should do is to create a user profile. In the profile, the user can describe his or her professional interests in a free text fashion. If the user already has a CV stored elsewhere, it may be pasted

into this field. The profile, once saved and stored, will be converted to a fingerprint and used when the Community function is invoked.
When clicking on *Community* or *Similar Agents* the user profile or the agent profile is matched with other users' or agents' profiles and the matching users are listed as seen in [Figure 2a]. The user may now display the email address or the profile of the found users (by clicking the *e-Mail* or *More info…* buttons, respectively) and can use this information to contact them.
The search results from the agents are displayed in a simple list, similar to those generated by most search engines [Fig. 2b]. However, when the user has verified that one of the returned documents is indeed relevant, the user can provide the agent with *positive feedback* by marking the document and clicking on the retrain button. The fingerprint of the agent will then be merged with the fingerprint of the selected document(s) and the result will become the new agent fingerprint, replacing the previous one.

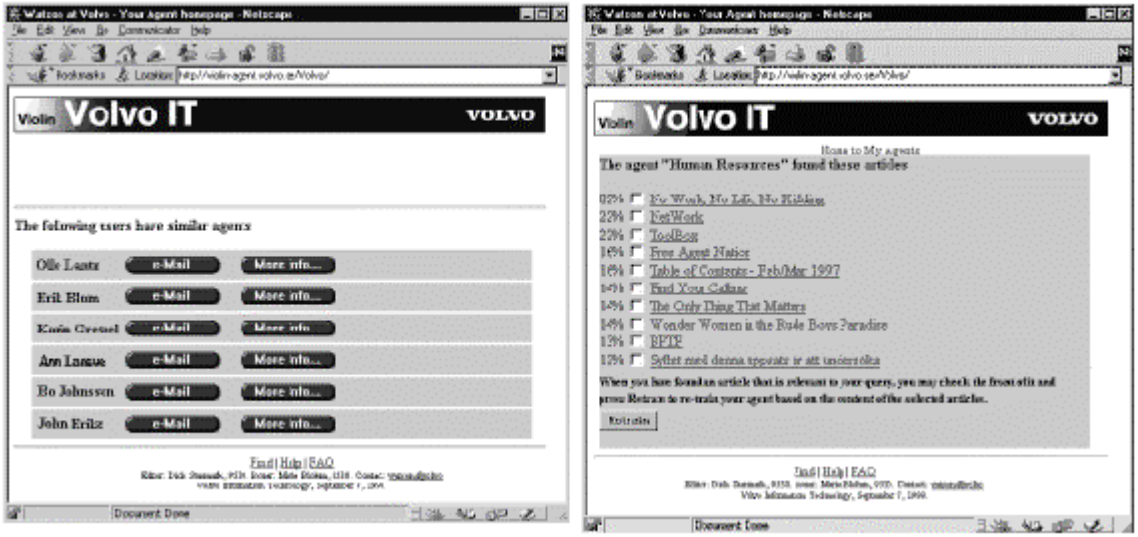

Figure 2: a) A click on the "Similar Agents" button displays a list of users who have created agents with similar profiles (i.e. users who have similar interests). b) The agent shows the 10 best matching results. By checking a document and clicking the retrain button, the agent can be re-trained.

4. Results

The users typically used the application frequently, sometimes heavily, during a couple of days and then stayed away from it for a while before returning for the next session. Usage was especially high right after the release of the application, and then declined slightly before settling on a stabile level throughout the rest of the test.
Overall, user reactions were very positive. 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". More specifically, several users expressed their appreciation of not having to come up with descriptive keywords, since "*they never really contain the meaning you have in mind anyway*", to use the words of one respondent.
Despite the general claims that this sort of "unattended help" with information retrieval is welcomed and appreciated, the users actually mainly reported negative experiences. Many users report what they refer to as "strange" or "unexpected" results. "*[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 tend to blame these bad results on their own inabilities. One user having received very little useful information says "*The rather shallow results may depend on me not using the right words. Otherwise I like the idea. Keep improving!*"
Our recommendation system prototype differed from how conventional Web search tools, such as search engines, are operated, and the users initially had problems adjusting. A common problem was that the users overtrained their agents. Some commented that they felt the way the agent worked had to be better understood by

the users in order for them to be able to interpret and anticipate the outcome of the search. The users who attended the introduction all appreciated the information and claimed that it helped them to use the tool in a better way. The answers from the users *not* attending the initial kick-off also clearly indicate that training is required. One non-attending user said: "*It took some time before I understood [how to set up the agents] … I think the instructions should be even clearer, so that you understand what the agent is and how it works*".

Not many of the test users exploited the Community feature, i.e. the feature that locates users with similar role description profiles. The reasons given for not using the function were that the test users already knew enough people doing similar jobs or that most users with similar profiles worked in the same department as the respondents. Alternatively, 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 know stuff I don't know*." This last remark was an opinion shared by several users, who suggested that an opportunity to search for users with *complementary* profiles would have been more useful. Of those who actually did try the Community feature, all but one considered it to be working. One user, however, claimed to have been connected to people with whom he had nothing in common, and this he referred to as "*a bug*".

The Similar Agent feature was however much more frequently used though it was built on pretty much the same sort of knowledge as the Community feature. "*It's really interesting to see who else is searching for these sort of thing*", one of the users commented. Many respondents reported that they were surprised to find certain people sharing their interests, or that the Similar Agent feature returned users not expected to be interested in a particular topic. These comments were not expressed negatively. On the contrary, the users regarded this result as a useful new insight and no one questioned the correctness of the result.

5. 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. 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.

The Community feature, which is based on explicit knowledge, has not been used much at all. There are several possible reasons: Firstly, this research is carried out on a corporate intranet and not on the World Wide Web. The whole organisation may see itself as a community of which the users are already members, and the need to be associated with sub-communities may not be particularly strong. This I find not very likely. Secondly, conventional office tools, e.g. word processors or email programs, do not support a community concept and the users are thus not familiar with this way of working. They do not see themselves as community members, and do not appreciate the full potential of the concept. Tangible business benefits may first have to be experienced for this new way of working to be appreciated. This alternative seems more plausible than the first. However, none of these two are strong candidates.

The third and most feasible explanation is 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. These profiles are presumably already known to the members and experienced as fictitious, since people are often viewed as performing their jobs according to their formal job descriptions though everyday practice provides evidence of the opposite, as shown by [Brown 1998]. This is consistent with the findings of [Argyris & Schön 1974] who refer to the worldview and values that people believe their behaviour is based on as "espoused theory" as opposed to "theory-in-use". The users rightly or wrongly assume that they know what the Community feature will return and they dismiss it as uninteresting.

The "Find similar agents" feature is different from the above in that it does not rely on static profiles but on the dynamics of retrained agents, initially created for a totally different purpose. The prompt "Enter your profile" connotes an official question equivalent to "what is your profession?" and the user enters the supposed official answer that corresponds to his or her role description. The agents, on the other hand, are created for personal benefit only and no official considerations are taken into account. Instead, real professional interests govern the choice of topics, which makes these profiles more "alive" than the previous 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. Setting up and training personal agents create search profiles that are able to

represent valuable tacit knowledge in a non-intrusive way. Since the primary objective for the agent is to find the user relevant information the incentive to create and maintain agents exists naturally, and the otherwise hard to overcome problem of having to set up profiles for somebody else's benefit is neatly avoided.
To be able to find this knowledge is, however, only a first step; it only helps to identify people within the organisation - it does not prevent these people from leaving the organisation nor guarantee that they will have time to share their knowledge on request. [Davenport & Prusak 1997, p.81] observe that "mapping who knows what in an organization creates an essential knowledge inventory, but does not guarantee the ongoing availability of knowledge". To achieve such permanent capturing and storing of knowledge, other measures that fall outside the scope of this paper must be deployed.

6. Summary

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. The are two main benefits of such an approach: i) 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 or not a given document is interesting, and ii) since a good profile results in more accurate information, a natural incentive to maintain the profile by giving feedback is created.
Previous research on agent-based retrieval system has studied how to connect users with information or users with other users. While this study shows 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 to the field of KM research is the observation that profiles based on the tacit knowledge of practice are conveyed as more trustworthy than the espoused theory-based job description profiles. The former profiles can be used to facilitate the sharing of tacit knowledge without having to make knowledge explicit.

7. References

Argyris, C., & Schön, D. (1974). *Theory in practice: Increasing professional effectiveness*. San Francisco, CA: Jossey Bass.

Brown, J. S. (1998). Internet technology in support of the concept of "communities-of-practice": the case of Xerox. *Accounting, Management and Information Technologies*, no. 8, pp. 227-236.

Davenport, T. H. & Prusak, L. (1997). *Working Knowledge: How Organizations Manage What They Know*. Boston, MA: Harvard Business School Press.

Fagrell, H. & Ljungstrand, P. (1998). Make an Agent and you shall find: An Intranet Recommender System. In Buch et al. (Eds.), Proceedings of IRIS21, Department of Computer Science. Denmark: Aalborg University.

Grudin, J. (1987). Social Evaluation of the User Interface: Who does the Work and Who gets the Benefit? In Bullinger, H.-J. & Shackel, B. (Eds.), Proceedings of INTERACT '87. North-Holland: Elsevier Science Publishers.

Hammersley, M. & Atkinson, P. (1995). *Ethnography. Principles in practice* (second edition), London and New York: Routledge.

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

Polanyi, M. (1998). The Tacit Dimension, In Prusak, L. (Ed.) *Knowledge in Organization*. Boston, MA: Butterworth-Heinemann.

Resnick, P. & Varian, H. R. (Eds.) (1997). Recommender System. *Communications of the ACM*, vol. 40, no. 3, (entire issue).