Knowledge Analysis and Mapping

Professional Services - White Paper

intellix
Intellix A/S
H. C. Ørsteds Vej 4
1879 Frederiksberg C
Denmark

Contents

Introduction

Over the past few years, much has been written about the possibilities for companies to share knowledge among their employees. One of the great challenges in this regard is how we can tackle the problem of distributing knowledge that one group of employees possesses to other groups. How can we ensure a correct and accurate identification of an employee's knowledge, and a correct interpretation and use of this expert knowledge by the other employees in the company? One possibility is to formulate knowledge as the ability to process information, and subsequently make this processing ability available for others by automating it in an electronic system.

This white paper describes the possibilities that Intellix's products offer to map and automate a business expert's ability to process information. The white paper takes its departure in the concepts that are contained in Intellix Designer 4.0. To gain the maximum benefit of the white paper, it is recommended that you have some experience with Designer 4.0, equal to the level that is achieved after completing the course *Knowledge Mapping Using Intellix Designer*.

The white paper is divided into two sections. In the first section, a conceptual introduction to the problem of sharing knowledge is provided. Key concepts such as explicit and tacit knowledge will be defined with the goal of showing the possibilities that exist for building decision support systems that do not work exclusively using rule based representations of knowledge. A considerable part of a business expert's knowledge is shown to be of such a nature that it only allows itself to be represented in the form of an example.

The second section consists of two chapters where the practical aspects of building decision support systems using Intellix's products are explained. Chapter 2 advises you how to undertake an interview of a business expert to cover their explicit and tacit knowledge. The chapter gives an introduction of the different roles that a knowledge mapper assumes during the building of a decision support system.

With the purpose of lessening the work that is necessary to model a business expert's ability to process information, the chapter also describes the structural and design-related aspects of the knowledge-mapping task. Finally, chapter 3 contains a description of the experience and practice in the construction of decision support systems possessed by the consultants employed at Intellix.

1. Knowledge and knowledge mapping

1.1 Knowledge

Over the past few years, much has been written about the change from the production society to the knowledge society. The prevailing view is a result of the increased speed whereby information is exchanged between people, which follows the concurrent spread of IT. Particularly the commercial use of the Internet by companies is considered as a prime force behind this change.

The greater focus on the role that is played by knowledge in a company's ability to improve the productivity of their employees has in general terms not been supported by a greater focus on the necessity of strictly defining *knowledge* as a concept. The majority of the people who study society's shift from a production society to a knowledge society share the view that knowledge is an absolute quantity that can be determined based on observations. This view, however, can be seen currently as being at odds with a realistic explanation, and, hence, understanding, of the role that knowledge plays for companies' possibilities for achieving increases in productivity.

If a company's ability to improve the productivity of its employees is influenced by the speed whereby the same company exchanges information, then knowledge cannot be an absolute and hence perfect quantity. Such perfection would mean that new information would not be able to influence existing knowledge, conflicting with the thesis that a company's improvements in productivity are based on a greater stream of information. It is therefore necessary to define the concept of knowledge, so that it can take over the central role it has already been accorded by many people in the description of the ability of companies to achieve improvements in productivity.

In a realistic description of the concept of knowledge, observations play another role altogether. Instead of consisting of a foundation for formalising and supporting knowledge, observations function as a test background for the correction of knowledge. Knowledge therefore should be considered to be correctable hypotheses (Popper (1972)), that is to say, "justified true belief" (Nonaka (1994)) that continually becomes adjusted relative to observations of reality. Knowledge reflects therefore a process rather than something permanent. (Bartley (1987))

A person's knowledge therefore reflects the processes that he uses to interpret and classify information. Only a few of these processes have such a nature that a person could write them down and communicate them in clear language. Nonaka (1994) observes that this so-called *explicit knowledge* is complemented by a *tacit knowledge* that exists in the unconscious processes that a person uses to process and classify sense impressions. This tacit knowledge has a subjective quality that relates to the person's situation-dependent actions, and is expressed for instance by intuition or by physical reactions.

Explicit knowledge, for instance, can consist of the procedures that a bank advisor uses to calculate a customer's net income. This knowledge the bank advisor would actually be able to write down and hence communicate to others. Tacit knowledge is harder to grasp. A bank advisor would find it far more difficult to put into words the intuitive feeling he has about a specific loan. Without being able to identify a reason, the bank advisor's experience tells him that something is wrong with the loan.

Using tacit knowledge, a person compensates for not being able to gather all the relevant information in situations where he has to make a specific decision. By taking his explicit knowledge as a point of departure, the person is able to make a decision. The tacit knowledge fills in the person's "blank spots of information" based on comparison of the particular decision-making situation with the decision-making situations where the person earlier has made decisions.

1.1.1 Conversion of knowledge

With the existence of two forms of knowledge, there are four ways to convert (and generate) knowledge; see figure 1 below. In practice, these four ways overlap.

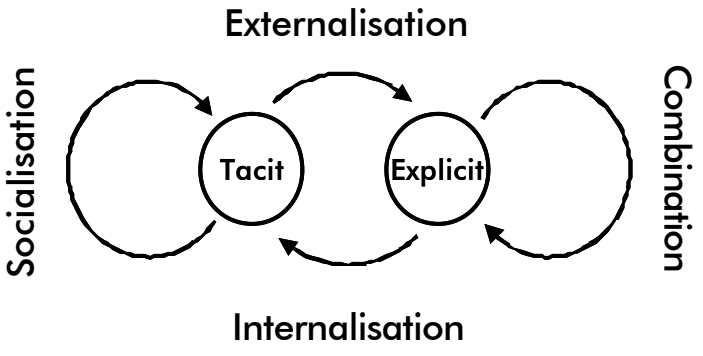

Figure 1. The two forms of knowledge and the processes whereby conversions between them occur (Nonaka (1994)).

As an example of a conversion of knowledge, we could take a bank advisor who develops written guidelines for an advisory process based on reports from different sources. As his point of departure, the bank advisor will take the explicit knowledge that is documented in the different reports and combine this knowledge, thereby creating the new guidelines for advice. The bank advisor converts existing explicit knowledge to new explicit knowledge.

On the other hand, the bank advisor utilises his tacit knowledge when he determines which of the reports are relevant to include in the creation of the new guidelines for advice. The bank advisor can quickly decide by skimming a report whether or not it is relevant to the business area where he works. However, it would be far more difficult for the bank advisor to write down general rules that describe when it is relevant to include a report, because the knowledge he uses to do that is primarily of a tacit nature. A part of this tacit knowledge is converted to explicit knowledge (externalised) in connection with the bank advisor's creation of the new guidelines for advice.

The bank trainee who follows the advisor will learn to make the distinction between relevant and irrelevant reports by looking at examples of how the experienced bank advisor evaluates the reports. This transfer of skill occurs as the result of the socialisation of the trainee in the established bank's working environment. The trainee creates tacit knowledge by observing how experienced advisors utilise tacit knowledge in their work.

At the same time as the bank trainee is taught to observe the actions of his more experienced colleagues, he also receives new knowledge through courses and education. All this knowledge to start with is explicit, but considerable parts over time will become pure routine, where the person is not able to explain his actions. The knowledge that was taught has been internalised.

According to Nonaka (1994), it is in connection with the conversion of knowledge that a basis is created for generating new knowledge. This is often seen in connection with knowledge mapping processes, since this work requires social interaction between business experts and knowledge mappers. The task of knowledge mapping often results in ideas arising that improve the mapping process.

1.2 Knowledge mapping

It is our ability to utilise tacit knowledge that makes us able to act in a complex and ever-changing world. By simulating this skill using computers, it is possible to develop tools that can support a decision maker in bringing forth the tacit part of the knowledge he uses as a determinant for his decisions.

Intellix Designer is a computer program that is built up around a neural network. The program contains tools that make it possible to bring out parts of a person's tacit knowledge. The process this involves is called knowledge mapping. When such knowledge mapping has been undertaken it is possible to formalise and refine a decision-making process. The entire job associated with formalising and subsequently refining this is called knowledge engineering. Figure 2 illustrates the entire knowledge engineering process.

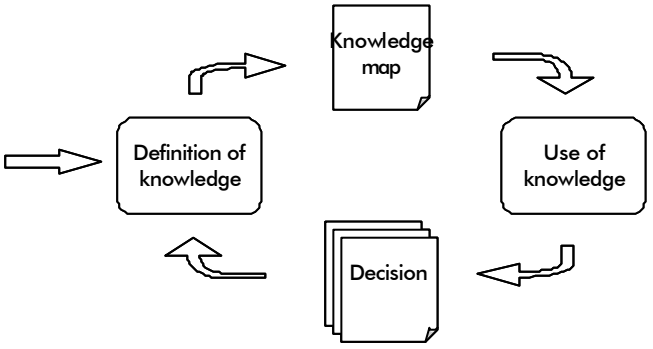

Figure 2. *The main processes in knowledge engineering*

In the start phase (definition of knowledge), where the contents of the area of knowledge are made more specific with the goal of formalising a decision-making rule (knowledge map), the knowledge mapper must form an overview of the decsion-making rule's logic including the information that is necessary for the running of the logic. Finally, he must map this necessary knowledge. When the rule of action has been formalised it is used (use of knowledge) to make decisions with. Gradually, as the number of such decisions increase (decision), in combination with their consequences, they will influence the actual area of knowledge and therefore ought to form a basis for a continual assessment of the current relevance of the formalised area of knowledge.

Intellix's terminology operates with knowledge domains as the general environment for representing knowledge. In a knowledge domain, knowledge is either represented in the form of pure rule sets and example sets, or as a combination of these. In section 1.2.1, the difference between rule based and example based knowledge domains at the same time will be shown by identifying types of problems within which the two forms of knowledge representation have their strength.

1.2.1 Knowledge representation in knowledge domains

In a knowledge domain, it is possible to represent knowledge in both explicit and tacit forms. Explicit knowledge is represented using rules such as "the customer cannot be given credit if he is registered in RKI". This representation can be used when complete information is available.

Tacit knowledge in the same direct manner cannot be converted to a representation that is comprehensible to machines. An indirect method for doing this is to convert the tacit knowledge to an explicit form, which can then be entered as a set of rules. It is important to note, however, that such a conversion bears the risk of distorting the tacit knowledge. Unless you are entirely sure that the converted knowledge accurately represents the tacit knowledge, you ought to stay away from this form of representing knowledge.

Alternatively, it is possible to use a method that does not carry this risk. The same process that is used to transfer tacit knowledge between people, in particular, by using examples, defines this method.

From a set of cases, the Intellix system's SOUL (Self Optimising Universal Learner) core creates a model that imitates the function that tacit knowledge plays for a person's ability to act on the basis of imperfect information. The SOUL core, for instance, "teaches" how a bank advisor works, by looking at examples (cases) of his (past) advice, and based on these examples, is able to generalise this knowledge for new cases. When the SOUL core is presented with information about a new bank customer, for instance, its advice will be based on a generalisation of the cases that the core has been taught. The advice to the customer primarily will be based on the cases that most closely resemble the customer's profile.

For an case-based model of knowledge to "learn" the right context of the business process that it represents, it is important that the cases the model is based on are representative of the business process. The cases therefore should cover situations within the area of the business process as well as possible. If the model, for instance, has not been presented with a given recommendation *X*, it will have no ability to determine the profile that characterises the customers that should be recommended the

product *X*. In the same way, the bank trainee will not be able to recommend the product if he has not be informed of its existence.

1.2.2 Which form is best?

Since it is possible with Intellix's products to represent both tacit and explicit knowledge, it is necessary to determine which form of representation is best suited to each situation. Explicit knowledge can be represented in the form of rules, as described in section 1.2.1, but it is also possible to set up a number of cases that represent the rules. This last instance is the equivalent of converting parts of the explicit knowledge to tacit knowledge, in particular, the parts that are not bound by cases. How far this is to be preferred must lie on an evaluation of the costs connected with identifying all the rules of explicit knowledge against the costs associated with letting the system generate advice based on cases. Case-based representation of knowledge cannot generally be recommended if the knowledge to be modelled already exists in explicit form. As humans, we typically seek to achieve an explicit representation of knowledge, since this is often considered easier to grasp than the tacit form.

If the knowledge to be modelled in the form of a knowledge domain only exists in tacit form, it is possible to attempt to convert it to an explicit form by using different techniques in a dialogue with the business expert, see section 2.2. This leads to the identification of a number of cases/rules, but if the expert afterwards considers that these are not representative for the area of knowledge, a case-based representation of knowledge is preferable. The result could also be a hybrid domain containing both rules and cases.

In some instances, it is possible to determine the quality of a domain by presenting the domain with a number of earlier cases to control afterwards whether the domain's recommendations are the same as those actually realised. Such a quality assessment can be made for both rule-based and case-based domains. The assessment would be an objective measurement of quality on the condition that the actually realised recommendations also were objectively correct.

Another reason to chose an case-based model over a rule-based model could be because there is no access to the "expert". In an assessment of credit, it is in reality impossible to decide whether a specific customer is able to repay a loan or not, and especially not whether the customer *plans* to repay the loan. Certainty about this is first achieved once the loan has been issued and the customer either has repaid it or defaulted on it.

It is therefore unavoidable to accept a certain margin of error no matter what form of representation of knowledge is used on which to base the credit assessment. Time related aspects also can be determining factors for the choice of form of the representation of knowledge. In tax reconciliation, there is actually in the law a set of rules from which it is possible to reconcile tax returns. In reality, this reconciliation is a time-consuming manual process that moreover uses subjective valuations, despite the explicit form of the relevant tax laws. In this case, a model of the process of reconciliation based on cases of similar tax returns quickly could provide a suggestion of the result. If the model is of a sufficient quality, it could be used to select the tax returns that had the greatest need for attention. Even though the model is not perfect, it can still be used to channel the most efficient use of limited resources. To make a purely rule-based model in this instance would be too cumbersome.

The disadvantage of many models is that they seldom can be interpreted by a reasonable number of rules. There exists in particular a trade off between the manageability of an case-based model and its precision. The more precisely a model represents a given area of knowledge, the more complex and consequently unmanageable it becomes. (A completely rule-based model will of course have the greatest complexity and unmanageability.) In such instances, there is naturally a close relationship with the complexity of the modelled area of knowledge. For a specific area of knowledge, it is necessary therefore to assess whether a simple set of rules should be used that explains the process, or if precision is more important.

A model based on examples is just as consistent as a model based on a set of rules. A specific customer profile will always result in the same recommendation. In the rule-based system, it will be possible to identify the cause (the rule) for the recommendation, whereas the case-based model seems to be less transparent.

Consequently, neither of the two methods for representing knowledge can generally be said to be the best – it depends completely on the area of knowledge under consideration.

1.2.3 Validation

Whether a case-based or rule-based form of knowledge representation is used, it is important to ensure the quality of the domain that is built. With case-based knowledge domains, it is especially important that the model, in cases where the examples do not cover the modelled area of knowledge, will be able to find patterns in data that can explain the relation in the training examples but that aren't representative for the modelled area of knowledge. For example, if all the customers that have contacted a bank on Wednesdays have been recommended the same product, a case-based model based on past recommendations will easily (and incorrectly) decide that customers that contact the bank on Wednesdays should be recommended this product. It is possible to guard oneself against such situations, though. In these cases, it is necessary to test the model using a number a situations that are different from those that the model is based on, and thereafter determine whether the model covers the set of cases satisfactorily.

2. Knowledge mapping in practice

This chapter describes the process of knowledge mapping and covers the analysis of business-oriented processes as well as the structured design of systems and construction of systems. Some of these skills are universal system developer disciplines and consequently will not be described in detail. On the other hand, the chapter will examine the techniques that are related to Intellix's products, having their point of departure in the previous chapter's definition of explicit and tacit knowledge.

2.1 Tasks and roles

The knowledge mapper's focus changes between keeping his eye on the over all relationships, while understanding details, and can best be illustrated by a cone where the cone's diameter reflects the degree of detail such that the narrow waist reflects the highest level of detail (see Figure 3 below).

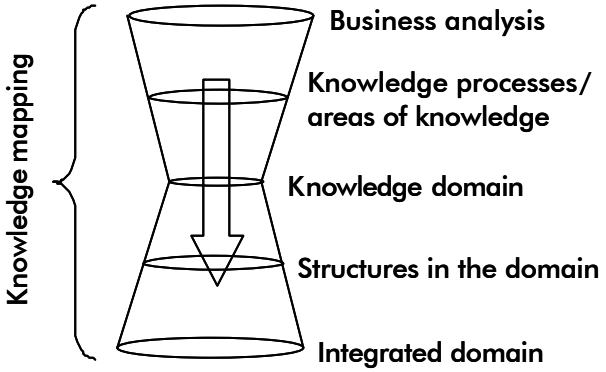

Figure 3. Knowledge mapping cone

To start with, the focus is centred on the business analysis. It is important that the knowledge mapper understands the business relationships in the decision-making process that are to be automated, and the relationships that the system will cover after construction. Compared with figure 2, this corresponds to the "definition of knowledge."

The following sections deal with the development of the knowledge map, as described in figure 2. It is in this phase that the business expert helps to divide the business area into smaller areas of knowledge and knowledge processes. In other words, the knowledge mapper's task is to identify logically delimited areas and processes within the overall area of knowledge. Only once this stage has been reached does the knowledge mapper begin to work as a consultant for the business expert.

It is now time to start building the knowledge domain using Intellix Designer. This is done by identifying the elements in the decision-making nodes that have been found in the previous analysis of the areas of knowledge.

Gradually, as the decision-making nodes and the knowledge domain is built, the knowledge mapper identifies overall structures and relationships within the knowledge domain. Finally, the knowledge domain is integrated with IT systems, so that the finished product stands ready as an assembled and integrated system.

As has been described, the knowledge mapper has many points to consider in his work. The knowledge mapper has the role of integrator in the knowledge mapping process, and typically must act as the liaison between the customer's business experts and the IT department. This co-ordinating role must be anchored in a solid understanding of the business processes and the IT systems, and include both the customer's and Intellix's own systems. This co-ordinating role is represented graphically below.

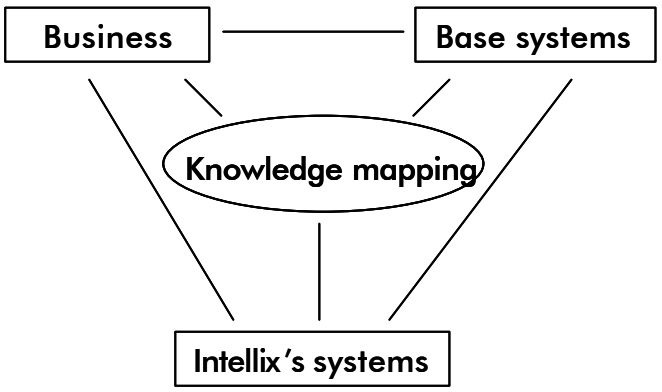

Figure 4. The knowledge mapper as integrator between the business, the customer's IT systems, and Intellix's systems.

2.2 The dialogue with the business expert

In the following section, the dialogue that the knowledge mapper carries on with the business expert is made specific. This is done using the "cone" presented in Figure 3. Before we begin looking at the specific dialogue, it is necessary to define precisely some of the concepts that already have or will be introduced.

Depending on the complexity and size of the knowledge domain, the individual area of knowledge is represented by one or more decision-making nodes. Decision-making nodes are comprised of a collection of pieces of information, called *factors*, that together with a collection of cases and rules describe how the decision in the decision-making node is made. The concept of a knowledge domain includes the total amount of information that is part of the given process of gathering decisions. The first step in the mapping of a knowledge domain is to identify sub domains, or areas of knowledge. (Very simple domains can be represented using a single area of knowledge, but this is generally unsuitable.)

2.2.1 Business analysis

The business expert and the knowledge mapper have to work together to define the knowledge model and make it specific.

This is done by taking the current working process as the point of departure, and finding areas that could benefit from automation in relation to the company's strategy, visions and other requirements.

Further to this, it is also necessary to decide which rights should be given to the knowledge model, and which situations the model should be able to deal with. Who is the end user of the final knowledge model, and what is their purpose for using the model (should an advisor be supported in decision-making, while a customer be given a specific product offer)?

Example

If a bank wants to formalise its credit policy with credit assessments, two possible scenarios can be established:

1. The bank advisor interviews the customer and decides which product he thinks best suits the customer. Then the knowledge model is used to control whether the chosen recommendation contains an acceptable risk and price. If the knowledge model finds that the advisor's recommendation is in order, the recommendation can be presented to the customer as the bank's offer. The knowledge model is thus used internally in the bank as a tool to give advice.

2. The customer is interviewed online to find out about their requirements and what they are willing to buy. If the case is simple, the knowledge model presents the customer with an offer and closes the case. If the case is complex, the gathered information is sent to the advisor in the bank, who contacts the customer and closes the case. The solution is therefore used externally as a new distribution channel.

The same decision-making process can be approached in several different ways, and the ultimate solution will be decisively different and have decisive significance for the work of the organisation.

2.2.2 The dialogue about knowledge processes and areas of knowledge

The first thing to be established is which over all decision-making nodes are necessary to represent the area of knowledge under consideration. Next, these areas must be mapped with their point of departure in our awareness of explicit and tacit knowledge.

When an area of knowledge is to be mapped, the first thing to be defined is the amount of possible result values. Then the amount of factors that are necessary to be able to differentiate between the different result values is identified. The way that factors are connected with the results is by cases and rules. It is at this stage that it is important to be able to differentiate between areas of knowledge that are based on explicit or tacit knowledge.

If the business expert is able perfectly to describe the regular relationships between factors and results, the area of knowledge and the decision-making node is said to be rule-based. No estimated decisions are made in a rule-based knowledge model.

Most frequently occurring are situations where the business expert is not able to establish a complete set of rules for the decision-making process. The daily work is based on estimates that stem from the experience that the expert has built up. The expert's knowledge of the area is, in other words, tacit. Typically, however, it will be possible for the expert to identify archetypical situations where the result is easy to identify. These combinations are rules on the same footing as the rules in a rule-based model, but because the set of rules is incomplete, the model is called case-based and the rules are seen as cases. A knowledge model is created with its point of departure in these cases, and the neural network is used to estimate the results for all other factor combinations. In a model such as this, the cases typically will not be evenly distributed over the entire knowledge model, which is necessary to achieve reliable estimates from the model. By means of dialogue with a knowledge mapper, and using Intellix Designer, the expert is helped to formulate the result values for the combinations in the areas of knowledge that the system has difficulty dealing with. The expert's tacit knowledge is externalised (see figure 1, and the related description). When the system is able to estimate reliably in all instances, the knowledge mapping is complete.

2.2.3 Structures in the knowledge domain

Knowledge domains can become very large and extensive. In the same way knowledge domains also become correspondingly difficult to hold in overview, unless the knowledge mapper constantly keeps his focus on the structure of the domain.

The biggest threat to this overview is the great amount of information that is typically contained in a domain. A well-known tool for obtaining an overview in such situations is "information hiding," which is the principle that information that is attached to a specific area of knowledge cannot be accessible to other areas of knowledge. The following section shows how "information hiding" ought to be used to achieve an overview of knowledge domains, and how Intellix Designer supports this.

Folders in knowledge domains

Folders should be used to divide the knowledge domain into areas of knowledge and knowledge objects. An area of knowledge is characterised by gathering all knowledge and information that is related to a specific business area or decision. A folder can include sub folders that represent subdivided areas of knowledge, partial processes, or partial decisions.

Folders should include the documents that logically are connected within a given business relationship. Moreover, the domain ought to be documented at the folder level (at a minimum), since these represent the lowest units of knowledge in the domain (a good idea is to use "description" in the folder for this purpose). In this way, an overview is ensured of where in the domain documents with a given type of knowledge can be found, which again supports the over all understanding of the dynamics of the knowledge domain.

2.2.4 The completed knowledge domain

When the cone is thus run through, starting from the wide understanding of the business area and extending to a focus where every single detail is defined, where the whole is sought by building up manageable structures, the knowledge domain is complete.

3. Experiences and Best Practice

This chapter contains a description of selected experiences that knowledge mappers at Intellix A/S have made in the knowledge mapping process.

3.1 Necessary competence

The business expert should possess the following competence:

- Profound understanding and knowledge about the business area that is mapped, in other words, a person who can communicate their explicit as well as their tacit knowledge about the business area to a third party without the help of formal policies.
- Decision-making competence that makes him able to decide which decisions the model should make in relation to the company's policies and strategy.
- Understanding how to use Intellix's tools, for which reason it is especially advisable that the expert participate in Intellix's courses. This knowledge and understanding is necessary both for the construction of the knowledge domain and for its later maintenance.

The knowledge mapper should possess the following competence:

- Fundamental knowledge of the industry and of the business area that is dealt with.
- Have a personality that corresponds with the very integrative role that the knowledge mapper has in a project (please see figure 4).

Furthermore, it might be necessary to gain help from the following profiles:

IT expert

- Good awareness of the company's systems, and hence of what information can be obtained where, and of what information is necessary to provide the systems.
- For maintenance reasons, this person also ought to have a good knowledge of Intellix's products from course work.

End user

- The end user can be involved to ensure satisfaction with the new system (where the knowledge model is found).
- By involving the end user early in the process, the subsequent time that is needed to teach the use of the system is reduced.

3.2 Commencement of the analysis and knowledge mapping phases

Depending on the customer's organisation, it might be advantageous to hold workshops before and after the construction of the knowledge domain. The first workshop will be a kick-off for the project where all the participants in the project on the customer side gather and are informed about the forthcoming work and of the different activities, expectations, and of the framework for the project, and its limitations and risks. After this, 80% of the knowledge domain is built with the customer's business experts, using the boundaries that were sketched at the project commencement workshop. Then another workshop is held with the project group (as well as possibly a management group) from the customer, where the expectations for the knowledge domain are finally adjusted. This is done because the implementation of a system that influences a company's ways of doing business typically has several non-technical consequences that first become visible during the process. These can be handled should the occasion arise, after which the work on the domain can be completed.

3.3 Interviews during the knowledge mapping

It is often an advantage to have two knowledge mappers. One of them can focus on working with Intellix Designer and the model that is specifically built, while the other functions as an integrator with the job of integrating into the model useful discussions and useful information that otherwise would be lost because the person at the keyboard only focuses on the detail at hand.

The advantage of using several experts is that experts typically have different views in non-trivial situations. It is precisely these clarifications between experts that help make tacit knowledge explicit. On the other hand, participation by too many experts can lead to too many discussions, whereby focus is easily lost.

4. Literature

Bartley, W.W. (1987). Philosophy of Biology versus Philosophy of Physics. In *Evolutionary Epistemology, Rationality, and the Sociology of Knowledge*, edited by Radnitzky, G. and W.W. Bartley (1987). La Salle, IL: Open Court.

Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science, 5, 14-37*

Popper, K.R. (1972). *Objective Knowledge, an Evolutionary Approach*. Oxford: Oxford University Press.