

NOVAK, Jasminko  
WURST, Michael  
SCHNEIDER, Martin  
FLEISCHMANN, Monika  
STRAUSS, Wolfgang

## **Discovering, Visualizing and Sharing Knowledge through Personalized Learning Knowledge Maps**

Publiziert auf netzspannung.org:  
<http://netzspannung.org/about/tools/intro>  
24.06.2004

Erstveröffentlichung: Van Elst, Ludger (Hrsg.): Agent-Mediated Knowledge Management. Berlin: Springer Verlag, 2003, S. 231 - 228.



**Fraunhofer** Institut  
Medienkommunikation

The Exploratory Media Lab  
**MARS** Media Arts & Research Studies

# Discovering, Visualizing, and Sharing Knowledge through Personalized Learning Knowledge Maps

Jasminko Novak<sup>1</sup>, Michael Wurst<sup>2</sup>, Monika Fleischmann<sup>1</sup>, and  
Wolfgang Strauss<sup>1</sup>

<sup>1</sup> Fraunhofer Institute for Media Communication,  
MARS Exploratory Media Lab,  
Schloss Birlinghoven,  
D-53754 Sankt Augustin, Germany  
[jasminko.novak@imk.fhg.de](mailto:jasminko.novak@imk.fhg.de)  
<sup>2</sup> University of Dortmund,  
Artificial Intelligence Dept.  
D-44221 Dortmund, Germany  
[wurst@ls8.cs.uni-dortmund.de](mailto:wurst@ls8.cs.uni-dortmund.de)

**Abstract.** This paper presents an agent-based approach to semantic exploration and knowledge discovery in large information spaces by means of capturing, visualizing and making usable implicit knowledge structures of a group of users. The focus is on the developed conceptual model and system for creation and collaborative use of personalized learning knowledge maps. We use the paradigm of agents on the one hand as model for our approach, on the other hand it serves as a basis for an efficient implementation of the system. We present an unobtrusive model for profiling personalised user agents based on two dimensional semantic maps that provide 1) a medium of implicit communication between human users and the agents, 2) form of visual representation of resulting knowledge structures. Concerning the issues of implementation we present an agent architecture, consisting of two sets of asynchronously operating agents, which enables both sophisticated processing, as well as short respond times necessary for enabling interactive use in real-time.

## 1 Introduction

The basic point of departure of our work can be related to the approach which argues that knowledge consists largely of a very personal, difficultly articulable and partly unconscious component, usually referred to as implicit or tacit knowledge [1]. Accordingly, a key to the communication and shared use knowledge, lies in the transformation of implicit knowledge and hidden assumptions to explicit structures perceivable und usable by others.

This recognition leads us to the following question: How can existing, but not yet explicitly formulated knowledge structures, of a given community or a group

of experts be discovered, visualized and made usable for cooperative discovery of knowledge in heterogeneous information pools?

In formulating a practical approach to addressing these issues we introduce the following constraints and definitions. We relate the notion of knowledge discovery to supporting the discovery of semantic contexts and relationships in an information pool which is either 1) too big or too fast growing to be scanned and categorized manually, or 2) consists of too heterogeneous content to impose one fixed categorization structure, or 3) serves different user groups with heterogeneous interests.

This definition immediately reflects the relevance of our approach and research challenge to practical applications. On one hand these conditions apply today to a vast range of Intranet/Internet portals in their own right. On the other hand, they can also be generalized to the problem of connecting existing information sources on the Internet in a way that allows semantic exploration of information and creation of both personalized and shared structures of knowledge.

The paradigm of agents is a very promising approach to overcome some of the problems connected with heterogeneity on the side of the data sources as well as on the side of the users. As agents should operate autonomously and can be loosely coupled, they are well suited for the integration of distributed heterogeneous data sources, building unifying wrappers around them. This becomes especially beneficial, if agents can learn to extract information from an information source automatically (see for example [2]). On the side of the users, the paradigm of Personal Information Agents offers a way to encapsulate the interests, the knowledge as well as the preferences of individual users. This is especially important in a system serving different groups of users. While agents in some systems mainly filter and distribute information (as in [3] for distributing Knowledge Discovery results) they are also very well suited for the task of capturing the (tacit) knowledge of users, as to make it accessible to others. Therefore Personal Agents can take the role of mediators between users and information sources, as well as between users among each other (see also [4] and [5]).

Based on the paradigm of "Agent Mediated Knowledge Management", we present a model for expressing implicit knowledge structures of individuals and groups of users and for using this as a means for semantic navigation and discovery of relationships in heterogeneous information spaces. We will show, how this model enables the implicit, as well as the explicit exchange of knowledge between users through intelligent agents. In particular, we discuss a model for unobtrusive generation and profiling of personalized user agents based on effects of user interaction with information and a related model for visualising and navigating resulting knowledge structures. Furthermore we present an agent architecture consisting of two sets of asynchronously operating agents. This architecture enables us to perform sophisticated data and interaction analysis, without losing the property of short respond times essential for interactive work in real-time.

## 2 Personalized Learning Knowledge Maps

In order to develop a working solution for capturing and visualizing implicit knowledge structures of human users based on their interaction with information, two basic problems need to be solved:

1. a context for user actions has to be created in order to be able to interpret the meaning of user interaction with information items. The lack of a clear interaction context is the main difficulty of general user-tracking and interaction-mining approaches such as [6].
2. a form of visual representation has to be found that communicates to the user both the semantics of the information space in itself (content, structure and relationships) and relates this to the meaning of his actions.

As a practical context for addressing these issues we take the process of information seeking and semantic exploration of a document pool. This can be understood as a process in which the users interaction with information both reflects their existing knowledge and produces new knowledge structures. In the concrete solution we develop a model of agents learning personalized knowledge maps. The notion of a knowledge map in our approach refers to the representation of information spaces in which the individual information items are not isolated but structured according to possible meanings and semantic relationships. This concept serves as a point of departure for both providing an unobtrusive context for interpreting user actions as well as for visualizing the resulting knowledge structures and exchanging them between users.

### 2.1 Capturing User Knowledge

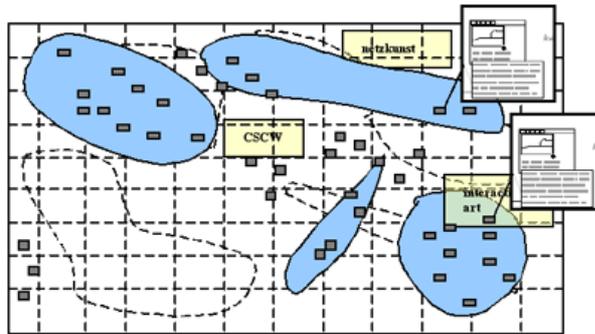
The basic idea is to build agents, that provide the users with a semantically structured overview of a document pool as a basis for their exploration and interaction with information. The results of their interaction can then be taken as the basis for generating user-specific templates. These templates (personal maps) are the basis for generating and profiling personal information agents which can then automatically generate a semantically structured map of a document pool, in a way that reflects a users particular point of view. In our approach the generation of user-specific templates is based on a two-stage model. First the user is presented with an agent-generated knowledge map created by means of methods for autonomous machine clustering such as in [7], [8], [9], [10]. This map serves as an initial context and navigation guide for the users exploration of the document space.

As she explores the information space, the user identifies relevant documents and relationships between them which she can express by selecting individual items into personal collections and by (re-)arranging them according to her personal understanding of their meaning (e.g. by moving objects between groups,

creating new groups, adding labels and relationships). In this way the user creates a personal map as a natural result of her exploration of information. This template can now be learned by a personal information agent by means of methods for supervised learning. Having learned a user-specific template, the agent can semantically structure arbitrary information pools or dynamically classify unknown information items.

## 2.2 Visualizing the Knowledge Structures

The challenge for the visual representation of the knowledge maps is to develop a visual tool for both navigating a large information space as well as for discovering possible contexts and relationships between groups of items. This applies both to relationships uncovered by the machine analysis and those stemming from interpretation and knowledge of human users. To achieve this the two main elements of the knowledge map visualization are: the Content Map and the Concept Map.

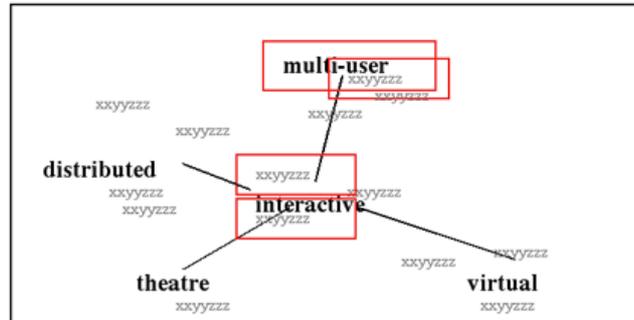


**Fig. 1.** The content map

The Content Map provides an overview of the information space structured according to semantic relationships between information items. In the first realization the Content Map visualizes clusters of related documents and offers insight into implicit relationships between their content. This is the main context for users exploration and interaction with information.

The Concept Map visualizes a concept-network that is extracted from the document pool and redefined by the users. This provides both a navigation structure and insight into the criteria that have determined the semantic structuring in the Content Map. These criteria are a kind of semantic axes that define a given structuring out of a variety of possibilities.

Since the personalized map templates have been produced by a user as an effect of his interaction with information and can be dynamically applied to reflect his



**Fig. 2.** The concept map

point of view, they are a form of representation of the user's knowledge that has previously not been expressed. Visualizing the personalized maps and the related concept structures, and making them available to other users is a way of making the users knowledge perceivable and available to others.

Hence, our claim that this is a way of expressing a user's implicit knowledge resulting out of his interaction with an information space, in a way, which makes it perceivable and usable by others.

### 2.3 Exchanging Knowledge

In our model, there are two major ways to enable the exchange of knowledge between users. Firstly, users can explicitly exchange knowledge maps they have created, secondly, information contained in personal maps can be analyzed implicitly (without the user being involved) and then be used to support the exploration and map editing process of other users. In chapters 4.2 and 4.3, we describe, how both of these possibilities are integrated in our system, the first through a personal assistant to enable search in the set of knowledge maps, the second through interaction analysis used for learning personal maps.

### 2.4 Relationship to Related Work

The basic idea of generating user-specific templates and applying them for personalized structuring and filtering of information has been previously realized in several different ways. In one class of approaches the users have to express their preferences explicitly and as their primary task, such as by voting, preference profiling or initial selection of items from a given information pool (see [11] for an overview). One critical issue here is the bootstrapping problem: the available orientation for users initial identification of relevant items in an information pool (which they are not familiar with) is based solely on already available profiles of

other users (e.g. [12]). A related problem is that of communicating the intention and meaning behind user choices that contributed to the creation of a given profile to other users: the profiles themselves are typically neither explained, nor visualised, nor put in relation to the semantic structure of the underlying information pool. Another typical class of approaches attempts to analyze the users actions in form of click streams and navigation patterns on the web (e.g. [13], [6]). The critical issue here is the lack of a clear context for interpreting the meaning of users actions.

In our approach both of these problems are addressed by introducing a system generated map as 1) a clear initial context for user actions, 2) a structure for semantic navigation in an unknown information pool, 3) form of visualising users personal knowledge structures in relation to the original information space. This approach also allows us to make the expression of personal points of view unobtrusive and not distracting from the users main task: that of discovering relevant information and internalizing it into knowledge. Furthermore, the personalized maps in our approach provide an easy and understandable way for communicating and sharing knowledge between different users both through explicit selection of different maps by the users themselves, as well as through implicit inference mechanisms of the agents that analyze the relationships between individual maps (Chapters 4.2, 4.3)

### 3 Agent-System Architecture

As already mentioned, our system consists of two different kinds of agents (Fig 3). One group of agents is concerned with responding to user requests. These agents have to work very efficient, as interactive work requires very short respond times. To achieve this, we use a second group of agents, which asynchronously preprocess data and store it in intermediate structures. These agents take much of the work load from the first group of agents. Using this strategy we can use sophisticated and costly data and interaction analysis methods and even so have short respond times. In the following, we will roughly describe some of the systems components.

#### 3.1 Data Preprocessing Agents

These agents allow the user to create a pool of documents by connecting heterogeneous data-sources. The user can either choose between readily available data sources or manually connect other structured data-sources (such as databases and semi-structured document repositories). This is supported by a dynamic data adapter for user-oriented semantic integration of XML-based semi-structured information.

Preprocessing includes a text-analyzer for encoding semantic properties of texts into a vector space model, link&reference analysis, co-author relationships and the extraction of other properties.

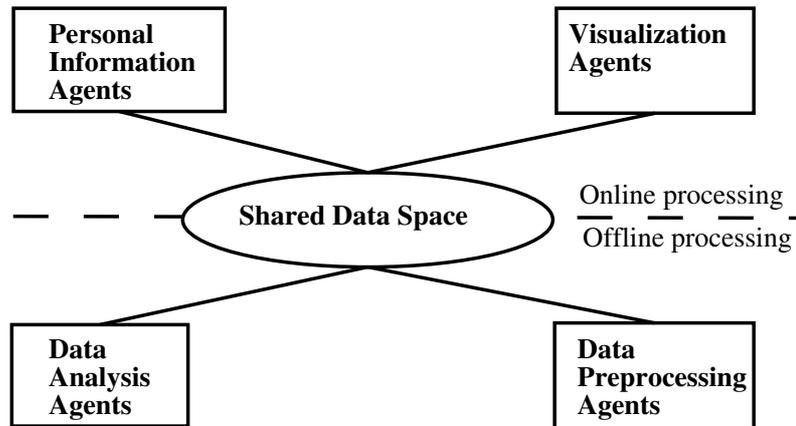


Fig. 3. The Agent system structure

### 3.2 Data Analysis Agents

This layer contains agents for semantic processing of data and for interaction analysis. Interaction analysis processes the personal maps of all users in order to identify relations between objects (see 4.2). While preprocessing is performed only once for an object, interaction analysis is performed at regular intervals, as the set of personal maps changes.

### 3.3 Personal Information Agents

Personal Information Agents have three different tasks. Firstly, they construct knowledge maps, based on unsupervised learning, allowing the user to influence this process by a set of options. We use Self Organizing Maps (SOM) for this purpose (see 4.1). Secondly, personal agents are able to learn a personal map, created by a user and to apply it to an individual object or a whole information pool. For this purpose, we use case-based reasoning, based on content and interaction analysis, as described in 4.2. The third task of a personal information agent is to provide its user with interesting maps of other users, enabling a direct exchange of knowledge between them (see 4.3).

### 3.4 Visualization Agents

The visualisation agents provide necessary post-processing of the data and of the interaction-analysis done by the personal information agents. They take care of collecting all necessary information from different agents, needed to construct all the information layers of the Content Map and the Concept Map described in

the previous chapter. In a typical case, a personalised information agent delivers the logical map of documents grouped into clusters of related content, with basic parameters such as weight of document membership to a given cluster, typical members of each cluster etc. Based on the selected visualisation model, the visualisation agent then retrieves information stored by the data integration assistant and preprocessing agents, in order to fill in additional information (e.g. titles, abstracts, term-document frequencies etc.) and compose all necessary information layers needed for a given visualisation.

### 3.5 Agent Communication and Coordination

We use two classical techniques for agent communication and coordination. The exchange of data between agents is realized as shared data space. The idea is, that on the one hand there are possibly several agents working on preprocessing in parallel. On the other hand, the preprocessing agents can provide data for the request processing agents asynchronously, without direct communication or coordination. Though within each group of agents, there is need for a tighter form of coordination. This is done by a simple event service based on XML and SOAP.

## 4 Personal Agents and Data Preprocessing

In this section, the personal agents used for automatically creating knowledge maps, for learning personal knowledge maps and for searching the set of knowledge maps from other users are described in more detail. Along with these agents themselves, the corresponding agents for preprocessing are described.

### 4.1 Clustering Documents Automatically Using Self Organizing Maps with Interactive Parameterisation

We use Kohonen's self-organizing neural network ([7], [8]) to map the high dimensional word vectors onto a two dimensional map. As the vectors encode semantic properties of texts the map will position semantically correlated texts close to each other.

The information describing the distribution of items and the measure of "semantic similarity" between both individual items and groups of items (clusters) provides the basis for the visualization in form of the Content Map (Fig. 1, Fig. 4)

In addition to the content map, a concept map is generated, which visualizes the relations between different words (Fig 2, Fig 4). We employ an approach similar to that described e.g in [14] to build this map. The idea is to structure the words

by examining which other words appear in the context of a given word. The high dimensional context relations resulting from this are then mapped to a two dimensional space, again using the SOM. In this way we can create an initial set of concepts (words) that serve both as an explanation of the clustering and as a navigation structure. Our system provides the additional feature, that users can customize the aspects according to which the maps are generated by manually selecting a number of words on the concept map. The weights for these words in the vector space are increased making them the most important words. Then the mapping procedure is re-applied using these modified weights.

In this way, by interactively exploring different possible clustering variants, the users can develop an understanding of how the clustering works and what makes out the character of individual document groups. Moreover, they can develop an understanding of the overall semantic structure and relationships between groups of documents (e.g. topics, trends, representatives) and the concepts (words) that determine a particular semantic point of view. This allows semantic navigation across a document pool for identifying relevant pieces of information embedded in contexts and relationships from different points of view. The discovered insights that are internalized by users as acquired knowledge are then reflected in their own personal maps.

#### **4.2 Combining Content-Based and Collaborative Methods to Learn Personal Knowledge Maps**

By creating a personal map, the user defines a set of classes. The idea of learning a personal knowledge map is to find a function, which assign new objects to these classes automatically. After such a decision function has been found, a map can be applied to any single object or information source provided by the system. The question of whether an object can be reasonably assigned to any of the user defined classes or not is to a significant extent subject to individual preference.

As a consequence, the system gives the user the possibility to interactively adjust the threshold of minimal similarity. If there is no object in the personal knowledge map to which the given document is at least as similar as defined by this threshold, the object is assigned to the trash class. Otherwise the decision function is used to assign it to any of the user defined classes. This allows the user to fine tune the personalized classification by exploring the influence of the threshold between two extremes: if the threshold is maximal then all objects are assigned to the trash class, if it is minimal all documents are assigned to some class and trash class is empty.

As method to find such a decision function that assigns documents to clusters we use Nearest Neighbor(e.g. [15]). This methods first identifies the most similar objects on the personal map for an object in question, and then performs a majority vote among them about the class to which to assign the object. This method offers two important advantages in our context. The first one concerns

efficiency and respond time, the second one concerns the problem, that the user usually provides only few training data. The idea is, that the similarity between objects can be pre-computed using sophisticated algorithms based on data and interaction mining. The query processing agent needs only some few access operation to the result matrix making it very efficient. An outdated similarity matrix could make the result sub-optimal, though in most cases this wont affect the performance, as similarities change only slowly. In the remainder of this section, we describe how the preprocessing agents build this matrix based on content and context analysis and how this helps us to deal with the problem of few training examples.

Content analysis uses properties of items (word vectors, authors, etc.) to measure the similarity of these items. The idea of context analysis is the following: If two objects appear together in many user edited clusters, then we can assume, that these objects are in some way similar. This is a very interesting feature of our system, as items are not only rated by users, like in "collaborative filtering" systems, but are put into the context of other items. This is much more powerful, as usually an item is not interesting or relevant per se, but only relevant in a given context. It helps us to deal with the problem, that the user provides only few examples, as the personal maps of all users can be used to support the learning and application of a map, not only the one of the actual user.

Both the content-based similarity and the context similarity are in a first step calculated independently of each other. Content based similarity is a linear-weighted combination of individual aspects.

For context similarity we use the "Dice"- coefficient:

$$sim(x, y) = 2 \frac{|X \cap Y|}{|X| + |Y|}$$

were  $X$  is the set of clusters, which contain object  $x$  and  $Y$  is the set of clusters, which contain object  $y$ .

Using this measure, clusters, which do not contain any of both objects, are not counted, which seems appropriate for the given case. Also co-occurrences get double weight, as we consider them as more important than single occurrences. The membership of clusters and objects to personal maps is not taken into account at all, as it is quite unclear, how objects on the same map, but in different clusters are related.

Beside the direct use of context similarity in the combination with content similarity, there is still another possibility to take advantage of the user interactions. As mentioned above several aspects describing the content of underlying documents are combined using a weighted linear sum. Now, to find optimal values for this function, we can take the context similarity as prototypical similarity and use it to train a linear regression model (or even more sophisticated regression models). In this point our system also differs from systems that seek association

rules [16], which perform a kind of context analysis too, but which do not analyse the content of the underlying objects and put it into relation to their context.

The remaining question is, how content-based and context similarity should be finally combined into a single measure, preserving the advantages of both. The advantage of content-based similarity is, that it is always applicable and does not rely on user generated data. Though content-based similarity can lead to poor results, if the underlying objects are heterogeneous, e.g. make use of different terminology or are even written in different languages. On the other hand, using context similarity, we avoid these problems completely. The disadvantage of context similarity is however, that if only few users add a given object to their maps or if the contexts, in which it appears, diverge, we do not get any reliable evidence on the similarity of this object to other objects.

Consequently, we use a statistical test (chi-square based) to examine, whether the co-occurrences of two objects are significant in a statistical sense. If so, only context similarity is used, as we have a very direct clue of the similarity of these objects. If not, we use only content-based similarity, as it works independent of any object occurrences. First experiments on synthetic data show that the combination of both methods is on average superior to any of the methods in isolation.

### 4.3 Searching the Set of Personal Maps - Matchmaking

In order for a given user to benefit from the possibility of using knowledge maps of other users, there needs to be a way to efficiently identify knowledge maps which are relevant to him from a potentially huge set of such maps. The method we are developing is based on the following idea: on the one hand a user has preferences, long term interests and pre-knowledge. On the other hand, she has a current information need. To capture both, we are developing a search facility, which combines keyword search (current information need) with a similarity analysis between users based on their personal maps (long-term information need). Combining both aspects results in a ranked list of personal knowledge maps available in the system. As this feature is currently under development, we refer to future work for more details.

## 5 Visualization and Interface

The critical issue in visualizing the knowledge maps and using them as a tool for discovering new knowledge is an intuitive interface which allows the user to unobtrusively construct personalized maps as accompanying effect of his exploration of an information space. On one hand, this requires that the results of the clustering and personalized classification mechanisms need to be visualized in a way, which provides clear insight into the meaning and criteria of a given



tation communicating the semantics of individual elements to the human user (e.g. clusters, keywords, labels etc.) it is also a medium of (implicit) communication between the agent and the user. The result of the new, personalised maps generated by the agent is communicated to the user in the same visual way.

A special issue for the visualization and interface has been the handling of navigation in large information spaces. Especially when investigating possible relationships between different groups of documents, the user needs both to be able to keep switching between detailed views of individual groups and the views encompassing larger, global portions of the map. Furthermore, one also needs to be able to move smoothly between different information layers such as titles, keywords (machine and human), abstracts and images. In addressing these issues we built on experiences from previous work on focus+context techniques such as in [17], [18] and [19]. As a concrete solution we have developed a model for semantic zooming with multiple zoom focuses and global and local zoom areas (Fig. 4). It allows the user to select different zoom focuses and pin them down as fixed points of interest without losing the overview. The user can further decide whether the zooming should have only local effect at the given focus area (drill-down mode) or scale through the global environment so as to always keep both focus and overview (progressive-zoom mode).

## 6 Practical Applications

The practical test bed and first application context of the described work is the Internet platform [netzspannung.org](http://netzspannung.org)<sup>3</sup> [20]. [Netzspannung.org](http://netzspannung.org) aims at establishing a knowledge portal that provides insight in the intersections between digital art, culture and information technology. Typical [netzspannung.org](http://netzspannung.org) users are experts and professionals such as artists, researchers, designers, curators and journalists.

The basic requirement of such an interdisciplinary knowledge portal is: a continually evolving information pool needs to be structured and made accessible according to many different categorization schemes based on needs of different user groups and contexts of use. By using the described system this heterogeneous user group will be able interactively compose and collaboratively structure an information pool coming from different data sources, to visualise and explore it through personalised knowledge maps, and to construct a shared navigation structure based on the interconnection of their personal points of view.

The current system prototype has been internally deployed as information access interface to the submissions of the cast01 conference<sup>4</sup> and of the competition of student projects digital sparks. This simulates the use scenario in which users can explore possible relations between information usually isolated in separate archives of different communities in the fields of media art, research and technology. The results can be tried out in the guided tour and partially online available

---

<sup>3</sup> <http://www.netzspannung.org>

<sup>4</sup> <http://netzspannung.org/cast01/>

interactive demos. A very first visualization prototype for browsing system generated maps is still being used as public information interface.

## 7 Summary and Ongoing Work

We have presented an approach of how to use the paradigm of knowledge maps as a central concept to integrate different methods for interactive information search and for realising a model for collaborative discovery and sharing of knowledge. We have shown, how supervised and unsupervised learning can be used to generate knowledge maps, providing users with different views on the content and semantic structure of an information source.

We have presented an unobtrusive model for profiling personalised user agents based on two dimensional semantic maps that provide both a medium of implicit communication between human users and the agents, as well as a form of visual representation of the resulting knowledge structures. Furthermore, we have presented possibilities to use knowledge maps as medium for explicit and implicit exchange of knowledge between different users. As pointed out, our system differs significantly from so called "collaborative filtering" systems, as items are not just rated by the users, but are put into context, in a way which is unobtrusively embedded into users primary activity. In this sense, our system enables "collaborative structuring" rather than just "collaborative filtering".

Agents and Agent Mediated Knowledge Management have been used as paradigms to model and implement the system. This approach has shown to be well suited for the given problem, as it helped to structure the different components not only in an understandable, but also in an extendable way, offering the possibility of future additions and modifications.

Currently we are working on different methods, to extend and optimize the system. Firstly, we aim to add additional similarity aspects for the learning of personal maps. Secondly, editing personal knowledge maps, the user can arrange objects only in flat structures, which is very intuitive and easy to handle, but not always sufficient. Therefore the system will contain a second editor, capable of creating hierarchical structures and other relations between objects. From the point of view of processing, the problem is to develop such methods, which fully exploit the information contained in such structures. Finally, an evaluation workshop is planned for analysing the usefulness of the system and comparing the individual contributions of the different approaches.

The evaluation will proceed in three steps: first the basic model of capturing user knowledge through personal maps created in unobtrusive interaction with the system-generated map, will be evaluated. In the next step the exchange of knowledge between users through explicit sharing of maps, and through implicit agent inferencing as described in chapters 4.2 and 4.3 will be evaluated. Finally, the third test will evaluate the emergence of a shared navigation structure as a concept map network reflecting implicit knowledge of a group of users.

## Acknowledgements

The work described in this paper has been undertaken within the projects AWAKE - Networked Awareness for Knowledge Discovery and [netzspannung.org](http://netzspannung.org) - an Internet Media Lab, both financed by the German Federal Ministry for Education and Research.

## References

- [1] Nonaka, I., Takeuchi, H.: *The Knowledge-Creating Company*. Oxford University Press (1995)
- [2] Krueger, W., Nielsen, J., Oates, T., Finin, T.: Automatically generated daml markup for semistructured documents. In: *Proceedings of the AAAI Spring Symposium on Agent Mediated Knowledge Management*. (2003)
- [3] Haimowitz, I., Santo, N.: Agent-mediated knowledge management for tracking internet behavior. In: *Proceedings of the AAAI Spring Symposium on Agent Mediated Knowledge Management*. (2003)
- [4] Furtado, J.J.V., Machado, V.P.: Improving organizational memory through agents for knowledge discovery in database. In: *Proceedings of the AAAI Spring Symposium on Agent Mediated Knowledge Management*. (2003)
- [5] Tacla, C.A., Barthes, J.P.: A multi-agent architecture for knowledge acquisition. In: *Proceedings of the AAAI Spring Symposium on Agent Mediated Knowledge Management*. (2003)
- [6] Chalmers, M.: Paths and contextually specific recommendation. In: *DELOS Workshop: Personalisation and Recommender Systems in Digital Libraries*. (2001)
- [7] Lin, X., Soergel, D., Marchionini, G.: A self-organizing semantic map for information retrieval. In: *Proc. of 14th ACM/SIGIR Conf.* (1991)
- [8] Kohonen, T., Kaski, S., Lagus, K., Salojrvi, J., Paatero, V., Saarela, A.: Organization of a massive document collection. *IEEE Transactions on Neural Networks* **11** (2000) 574–585
- [9] Sack, W.: Conversation map: An interface for very large-scale conversations. *Journal of Management Information Systems* (2000)
- [10] Becks, A., Sklorz, S., Jarke, M.: A modular approach for exploring the semantic structure of technical document collections. In: *Proc. of AVI 2000*. (2000)
- [11] Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining collaborative filtering recommendations. In: *Computer Supported Cooperative Work*. (2000) 241–250
- [12] Resnick, P., Iacovou, N., Suchak, M., Bergstorm, P., Riedl, J.: GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In: *Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work*, Chapel Hill, North Carolina, ACM (1994) 175–186
- [13] Joachims, T.: Text categorization with support vector machines: learning with many relevant features. In: *Proceedings of ECML-98, 10th European Conference on Machine Learning*, Springer Verlag (1998) 137–142
- [14] Honkela, T.: *Self-Organizing Maps in Natural Language Processing*. PhD thesis, Helsinki, Finland (1997)
- [15] Aha, D., Kibler, D., Albe, M.: Instance based learning algorithms. *Machine Learning* **6** (1991) 37–66

- [16] Agrawal, R., Srikant, R.: Fast algorithms for mining association rules. In Bocca, J.B., Jarke, M., Zaniolo, C., eds.: Proc. 20th Int. Conf. Very Large Data Bases, VLDB, Morgan Kaufmann (1994) 487–499
- [17] Robertson, G., Mackinlay, J.D.: The document lens. In: ACM UIST. (1993)
- [18] Sarkar, M., Snibbe, S.S., Tversky, O.J., Reiss, S.P.: Stretching the rubber sheet: A metaphor for viewing large layouts on small screens. In: ACM Symposium on User Interface Software and Technology. (1993) 81–91
- [19] Bederson, B., et al.: Pad++: A zoomable graphical sketchpad for exploring alternate interface physics. *J. Vis. Lang. Comput.* **7:3** (1996)
- [20] Fleischmann, M., Strauss, W., Novak, J., Paal, S., Mueller, B., Blome, G., Peranovic, P., Seibert, C., Schneider, M.: netzspannung.org - an internet media lab for knowledge discovery in mixed realities. In: Proceedings of the Conference cast01 - living in mixed realities. (1991)