

A Multi-Dimensional, Unified User Model for Cross-System Personalization

Claudia Niederée
Fraunhofer IPSI
Dolivostr. 15
D-64293 Darmstadt
+(49) (0) 6151 869 831
niederee@ipsi.fhg.de

Avaré Stewart
Fraunhofer IPSI
Dolivostr. 15
D-64293 Darmstadt
+(49) (0) 6151 869 806
stewart@ipsi.fhg.de

Bhaskar Mehta
Fraunhofer IPSI
Dolivostr. 15
D-64293 Darmstadt
+(49) (0) 6151 869 882
mehta@ipsi.fhg.de

Matthias Hemmje
Fraunhofer IPSI
Dolivostr. 15
D-64293 Darmstadt
+(49) (0) 6151 869 844
hemmje@ipsi.fhg.de

ABSTRACT

Personalization is an established method for reducing information overload and facilitating targeted access to relevant information objects in an information system. However, the use of the information collected about a user and his current context in a user profile is traditionally restricted to a single system, although many user tasks require information handling activities that span several systems, like searching in different collections. The potential advantage of reusing profiles in other systems is increased by the advent of flexible service architectures for information systems like digital libraries that dynamically incorporate services with specific personalization characteristics.

This paper discusses an approach in support of cross-system personalization. Building upon an ontology-based unified user context model, which describes the relevant dimensions of the user and his working context(s), this approach uses the metaphor of a context passport that accompanies users on their travel through the information space. When interacting with a system the relevant “context-of-use” is extracted from this context passport and is used for improved support of the respective information-related activity. The approach is discussed in more detail for the relationship dimension of the unified user context model.

Categories and Subject Descriptors

H.1.2 *User/Machine Systems - Human factors*, H.1.2 *User/Machine Systems - Human information processing*.

General Terms Human Factors

Keywords:

Cross-system personalization, User Context Model, Context Passport, User Context, Personal Web

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AVI'04, May 25, 2004, Gallipoli, Italy.

1 INTRODUCTION

Personalization approaches rely on the fact that the user revisits the same system several times with similar requirements, since this is the prerequisite for collecting useful, user-specific data. An improved re-use and refinement of the information gathered in the user profile can be achieved when exploiting the same user profile in several systems. This is especially valuable when a user's interaction with an information system is part of a larger task that covers several such interactions possibly with different systems [16]. Examples of such information and knowledge intensive tasks are: writing a proposal, preparing a course or, leisurely planning a holiday. Such cross-system or cross-service personalization becomes even more important with the current trend to build upon architectures that are dynamic federations of services rather than a static integration (e.g. service architectures, Grid-based architectures and peer-to-peer architectures for digital libraries [1,17,18]). However, the reuse of user profiles requires an alignment of the personalization approaches. In more detail, the three basic challenges of personalization [28, 46] have to be reconsidered:

- How to model the user and his situation in an extensible and unified way that can be interpreted and used for personalization by different systems?
- How to collect data for filling and updating user profiles from user interactions with different systems based on the unified user context model?
- How to conceptually and technically enable different systems to make use of the collected user profiles for improved support of the user?

In this paper we present an extensible Unified User Context Model (UUCM) that can be used for modeling characteristics of the user and his situation, i.e. the user context [20], along different dimensions. In the UUCM several working contexts of a user can be modeled. The extensible set of UUCM facets, not only describes characteristics of the users themselves like interests and skills, but also aspects of the users' situation and environment. UUCM incorporates several dimensions such as Task, Relationship, and Cognitive Patterns dimensions. This supports flexibly modeling aspects of the user like the: tasks of a user and information objects related to the user. For easing the interpretation of user profiles based on the UUCM, we rely on Semantic Web Technologies for the representation of the user context model as well as the concrete user context profiles. The UUCM can not only be used to represent a user context model within a system, but also provides an intermediate format to exchange user profiles between legacy personalization systems.

We illustrate our approach by discussing in more detail a use case in the context of community events that combines the relationship and task dimensions of the UUCM.

Our approach on cross-system personalization uses a context passport metaphor based on the UUCM. The user profile information is encapsulated in a context passport that accompanies the user when traveling through the information space. When performing an activity the relevant part of the context passport, what we refer to as the context-of-use, is selected and used to transform the activity in order to better support the needs of the user.

The rest of the paper is structured as follows: In Section 2 we discuss related work in the area of user and context modeling and its dimensions. Our approach to cross-system personalization is described in Sections 3, 4 and 5. Section 3 describes the Unified User Context Model (UUCM) together with an ontology for the UUCM facets and dimensions, whereas Section 4 presents a use case for the UUCM, combining the relationship and the task dimensions. Section 5 introduces an approach for cross-system personalization based on a passport metaphor. The paper concludes with a summary of the most important issues of the paper and some ideas for future work.

2 RELATED WORK

Context can be thought of as the “extra”, often implicit, information (i.e. associations, facts, assumptions), which makes it possible to fully understand an interaction, communication or knowledge representation [6]. Historically, context was either ignored or treated as a black box [33]. One effort to build context models for knowledge representation was carried out by Cyc [13], where the knowledge is structured into a large number of different contexts and rules are used to support the logical consistency of assertions for a given (closed-world) situation and importing assertions across different situations. A more general approach to context modeling can be found in [6], where the metaphor of a Box is introduced to systematically discuss different approaches to context modeling and contextual reasoning.

An important issue of context modeling is the question on which parameters or dimensions are to be taken into account for modeling context. Some researchers [30] consider context as an n-dimensional space, with each individual context as a region in this n-dimensional space. Lenat [30] identifies twelve dimensions that they consider sufficient to describe context. It should be noted however that a context may not be 100% specified by giving its coordinates in 12-dimensional space, but only approximated. In our work, the user has a task at hand and context here is likely to be more clearly defined as compared to the more general situation in knowledge representation. Dimensions like Culture, Sophistication, Granularity, Modality, Argument-Preference and Justification defined in Cyc are less relevant for user context modeling.

Similarly, in our work, we consider a context model along a set of dimensions. We selected the dimensions Task, Cognitive Pattern, Relationship and Environment from the dimensions described in literature. Additionally, we define relevant subsets of the user context model as a user’s working contexts distinguishing the different roles he plays. Context-of-use selectively considers only those facets of the user’s context

model which are relevant during a given information related activity. Our definition of context-of-use is in contrast to other researchers [21] who have defined the term more generally, covering the complete environment which characterizes the user’s interaction with a computer system

Relation-based models of a user are information and community models that take into account the salient interrelationships of individuals in a cooperation or community context [35]. Having its roots in social theories, these systems use graph-based, or complex network structures to model interactions between human beings. Specifically, social network analysis (SNA) [64] extends and complements traditional social science by focusing on the causes and consequences of relations between people and among sets of people [15]. One approach to SNA is ego-centric network analysis. This approach focuses on an individual (or ego) and uses this individual’s network of relations to understand the diverse factors contributing to his/her behavior and attitude [34, 41].

More general than social networks, relations as well relation types are considered of high significance in modeling users and information [58,60]. Relation types describe common properties for a class of relation and include, for example, containment relations such as: meronymic (or part-whole) and class inclusion [3,50,66,67] as well as non-containment relations such as thematic roles (or case relations). Having its roots in linguistics studies [4,32], the thematic roles relation types are significant in modeling a user’s interaction with his environment because they represent a function, behavior, or assigned characterization that a participant plays in an association [54,58]. One type of participant from these classifications includes a determinant: an entity which is an active participant who initiates or determines the direction of process. Other types include immanent and recipient. In addition to ontology-based classifications, relational elements theories have been used to describe inherent properties of the relations themselves [9]. For example, a hierarchical property would be inherent for the relations *manages* and *contained-in*. Additionally, ongoing research in using relations include not only examining types for a single relation linking entities, but also types that apply to a relation built by either composition and transitivity of several relations or co-occurrences between relations [2,11,24,42]. The goal of such work is to discover patterns that would reveal a more complex relation in a network, which was derivable from the given ones.

From the domain of ego-centric social network analysis, we will be concerned with building relation based models of a single user and to this end consider an ego-centric approach. However studies have shown that persons, especially professionals, are tied into a manifold network of domain entity of different types [29]. Their connectedness within such a “web” is defined and affected by the work they do, the things and people they know, and the activities they engage in, etc. [42,65]. Such a socio-systemic context, necessitates a model more general than social theories. We, therefore focus not only on the relationships between people, but also relationships among resources within a given domain (of which people are just one type). In our research, a user’s personal web is an important dimension of modelling a user.

Traditional Models of users are based on a mentalist paradigm [36, 48]. Mentalistic paradigms are based on characteristics of the user which we collectively refer to as cognitive patterns.

These patterns represent user-specific aspects and include for example: interests, knowledge, preferences, misconceptions, or abilities [63].

Systems incorporating models of user's interest [16,26] have been widely used to selectively filter information on behalf of users from a large, possibly dynamic information source [5]. A common example of an interest based model is a collaborative filter which infers a user's interest and preferences from the ratings a user applies to an information item and from similarities between user's interests [27,45].

Despite studies which suggest that cognitive pattern models such as interest are insufficient data for accurate models of the user, it seems likely that these systems will continue to be adopted in the future [22]; therefore we consider, this traditional modelling dimension of the user.

Task models of user are considered important [25,57] based on the assumption that the goals of users (who participate in a task) can influence their information needs. When these needs are known in advance, a system can better adapt to its users [61] [62]. Based on these goal-driven theories for information related-activity, we consider the tasks an important dimension in modeling users and their context.

Environmental models are considered a key issue with respect to the interaction between human and computer because they describe the surrounding facts or assumptions which provide a meaningful interpretation (context) to a users' computer usage when their physical environment varies [53]. Furthermore, researchers have suggested that user's future usage scenarios will require more sophisticated support for changes that occur in a user's location and infrastructure. Such scenarios include: multi-computer usage: (e.g. a PC at work, a laptop on the go, and a PC at home); mobile computing: where a user carries a small information devices that can be temporarily connected to a network or ubiquitous information: where the information space can be accessed from information walls, kiosks, or desktops [16] and federated services: where collective information is dispersed among information sources.

Given the aforementioned trends and scenarios, environmental models are an important dimension in adequately supporting aspects of the users' situation and environment.

Generic user models are, in theory, systems which have, among other aspects, two major goals: 1) generality: which would allow a model of the user to be usable in a variety of application content domains; 2) expressiveness: in that the model is able to express a wide variety of assumptions about the user [26].

A number of factors contribute to the proposal in support of generic or unified user models. On the one hand, given the number of aforementioned dimension that is possible when modeling users, researchers have considered a generic approach to modeling users, because at present, there is no unified theory which systematically integrates all dimensions. On the other hand, current systems which typically model a user along a single dimension suffer from a limited view of users and a significant amount of potentially useful information about the user may be lost; thereby demanding a need for more robust models [7,55].

The trend towards dynamic federated services and multi-usage and ubiquitous environments, system designers can ease development and maintenance, increase reuse of user profiles

and support the alignment of personalization approaches across systems by considering the use of generic user models in which a variety of different dimensions of the user are possible [26].

Finally, similar work in standards to build unified user-related models for dynamic information spaces in RDF [10] and the standardized RDF vocabulary CC/PP (Composite Capabilities/Preferences Profile, [8]). The work in CC/PP describes user preferences with respect to devices and user agents. We follow the standardization guidelines set forth in these areas.

3 UNIFIED USER CONTEXT MODELING (UUCM)

Adequate user modeling is the basis for every kind of personalization. Beyond properties of the users themselves, as they are traditionally captured in user models, information about the users' situation can be used for extended personalization support leading to user context models [20]. Cross-system personalization requires a unified approach to user context modeling that captures the relevant facts about users and their current situation in a way that can be shared by different systems. In order to deal with the variations in personalization approaches such a Unified User Context Model (UUCM) has to be flexible and extensible enough to capture the relevant dimensions of user and context modeling.

The UUCM not only is structured along different dimensions like tasks and cognitive pattern, but also captures the fact that the user interacts with systems in different working contexts by structuring the model accordingly.

3.1 Unified User Context Modeling

Two levels are distinguished in our approach for unified user context modeling. On the abstract level, the basic building blocks for the UUCM are defined: user context, user model facets, core properties for facet description, and user model dimensions. We use the term facet here to represent the different characteristics of the user not as it is traditionally used in knowledge representation to refer to the properties of a slot. This level defines a meta-model for the concrete dimensions and facets used in the description of the user context model. For the cross-system personalization approach, that we are aiming for, it is assumed that this user context "meta-model" is published as a shared ontology and all participating systems rely on this model.

On the concrete level, an extensible set of UUCM dimensions and facets is defined. This is not restricted to just users' interests, but also includes tasks and relations to other entities in the information space and respective user communities. UUCM facets and dimensions are described as part of an additional ontology that is shared by the components involved in the cross-system personalization. The UUCM meta-model, thus, can be combined with different UUCM facet and dimension ontologies to form concrete user context models that provide the schema for the construction of user context profiles.

Structure of the Unified User Context Model

The structure of the UUCM is summarized in Figure 1. A simple but flexible and extensible way of modeling the different facets of the user can be accomplished by the use of name/value

pairs (cf. modeling of context by parameter value pairs in [6]) Following this approach, a name/value pair is used to capture each facet of the user context model (e.g. user preferences) and new facets can be easily added.

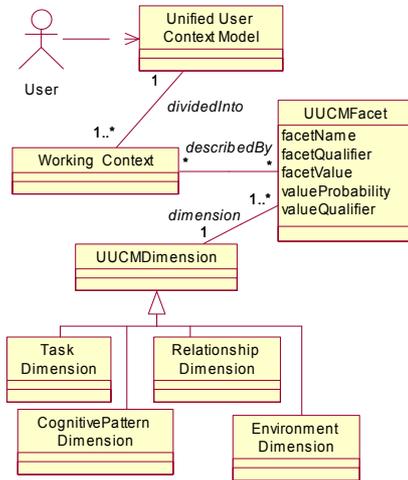


Figure 1: Structure of the UUCM

Since we want to use the UUCM in an open, cross-system personalization approach, qualification of names as well as values is a crucial aspect. Such a qualification binds facet names and values to vocabularies or ontologies. This eases interpretation of user profiles in a global context. In summary, each UUCM facet is described by the following properties:

facet name: Name of the UUCM facet to be described;

facet qualifier: A qualifier for the facet itself; this qualifier can be used to bind the facet to a defining vocabulary;

facet value: The values of the facet; The structure of the value may depend on the respective facet. In the general case, it is a value or a reference to another resource. (Our use of the term resource is comparable to that of RDF [49].)

value qualifier: A qualifier for the value(s) of the facet; this qualifier can be used to bind the value of the facet to a defining vocabulary (in contrast to the facet qualifier that qualifies the facet itself); In the case of a UUCM facet area-of-interest, for example, one might state that the ACM classification schema is used to specify the user's research interests.

value probability: A weight reflecting the probability of the facet value; this property of a facet can be used to express the reliability of facet values that are computed from analyzing user behavior.

facet dimension: Each facet is assigned to one of the dimensions covered by the UUCM; The UUCM facet area-of-interest, for example, is part of the UUCM cognitive pattern dimension.

Two aspects are used for structuring the UUCM. First, the UUCM is structured into several working contexts taking into account the different tasks and roles of a user (see Section 3.2). Secondly, the UUCM is structured along a set of UUCM

dimensions that are discussed in more detail below. The UUCM structure can, thus, be summarized as follows (see figure REF): A user profile is divided into a set of working contexts, where each context is described by a set of UUCM facets. Each facet is assigned to one of the UUCM dimensions. We encode the UUCM as an RDF Schema augmented with OWL expressions.

User Context Dimensions and Facets

The UUCM just defines the principle way a user context profile is described and structured for cross-system personalization. For the description of concrete user context profiles the UUCM relies on ontologies (or vocabularies) for the UUCM dimensions, the UUCM facets and for the facet values.

In our approach, we use four UUCM dimensions (see Figure 1): The Cognitive Pattern dimension, the Task dimension, the Relationship dimension, and the Environment dimension. The selection of the dimensions is based on user models in existing personalization approaches and on context modeling approaches (see related work Section). For each of the UUCM dimensions, our UUCM facet ontology defines a set of UUCM facets, which describe the aspects of the respective UUCM dimension. These facets are presented together with the respective dimensions in the following. However, the UUCM is independent of the selected facet and dimension ontology. It can be combined with other sets of facets and dimensions, but to discuss a flexible and extensible approach for unified user context modeling.

Within our facet ontology, the concrete facets are defined as subclasses of the general class UUCMFacet defined in the UUCM. They inherit all properties of the class UUCMFacet and define facet-specific restrictions like e.g. for the types of resources that are valid facet values. With this approach, there is a large flexibility with respect to which aspects are fixed for all instances of one facet (e.g. the facet name) and which can be selected individually for each facet instance (e.g. the value qualifier, if one wants to allow the use of values from different vocabularies). An alternative modeling approach is to make all facets instances of the general class UUCMFacet. This, however, gives fewer options for a systematic definition of specific types of facets.

The Cognitive Pattern Dimension

The cognitive pattern dimension describes cognitive characteristics of the user. It contains the facets that are traditionally used in personalization approaches. Based on an analysis of existing personalization approaches we selected the following facets to be included into our facet ontology:

- The facet *areas-of-interest* describing the interests of a user typically based on a controlled vocabulary or ontology of subjects (specified by the value qualifier).
- The facet *competence* with two facet subclasses *skill* and *expertise*.
- The facet *preference* that can be used to model preferences of the user.

Each of these facets may have several values. In this case the same facet is contained several times in the user context profile. Alternatively one may also enable the use of multi-value facets

within the facet ontology. However, in this case all values have to share the same value qualifier.

The Task Dimension

When interacting with an information system, the user is involved in a task that determines his/her information needs and the goals of the performed activities. Tasks are described in (domain-specific) task models that structure tasks into subclass hierarchies. The user profile may refer to such task models.

We identified the following useful facets for the task dimension:

Current Task: this facet describes the task the user is currently involved in. This facet has a facet qualifier referring to a task model description, and the value qualifier refers to a concrete domain task model, whereas the facet value is a reference to a concrete task instance based on this task model. Using this approach, any appropriately described task model can be fitted into the UUCM.

Task Role: This facet describes the role of the modeled user in the current task. This facet has a value qualifier referring to an ontology of roles in the current task domain, and the facet value refers to a node in the chosen ontology.

Task History: This facet points to a history of tasks completed so far within the current working context. The task history helps to keep track of completed tasks and subtasks. This facet again is based on a task model (typically the same as the current task) and refers to a sequence of interrelated tasks.

Further task properties are currently considered for inclusion in the set of facets of the Task Dimensions. Since considerable work has been done in task modeling [39,52], the challenge here is not to identify adequate properties to describe tasks, but to decide, which of these properties are required as integral parts of the user context model.

The Relationship Dimension

The requirements and information needs of a user are also determined by the entities the user is related to. The facets of the relationship dimension are based on the relationships the user is involved in. Each facet represents one type of relationship. The facets of the Relationship dimension are, thus, based on one (or more) relationship type ontology (in our application example relationship ontologies from the scientific research community domain). The facet names are names of relationship types, the facet qualifier points to the respective relationship ontology and the facet value refers to the resource the user is related to via this relationship. The value probability, finally, gives a probability for the existence of this relationship. This dimension is discussed in more detail in Section 4 in the context of the UUCM use case.

The Environment Dimension

The environment dimension refers to those parameters which are typically used for context-awareness approaches. Facets like current time, location, device, language etc are parameters which influence and, thus, are important in understanding the interaction between the user and the computer. These aspects are also important in understanding the user's changing requirements in different scenarios. These facets include:

Time: Every working context would be valid in a certain time frame;

Location: This facet refers to the physical location of user;

Device: The device the user is using, e.g. PC, PDA, etc.

Language: The language of choice for the user;

These are only the most central facets of this dimension. Many other facets describing the environment might be important depending upon the specific application. However, the environment dimension is not in the focus of our work. We will rely on existing and upcoming work in this area. We included this dimension to show that it can be part of the UUCM in the same unified way as the other, more traditional user modeling dimensions.

3.2 User Context Modeling, Working Context, and Context-of-Use

In principle the user context can be described by a large set of facets. However, the user interacts with systems in different roles and is involved in different tasks in parallel, each of which is associated with a specific subset of the user context facets (e.g. private vs. job-related activities, see also [42]). To reflect this structuring the user context is divided into multiple working contexts (see Figure 1) grouping together user context facets that are related to and relevant for the same task and/or role of the user.

While accessing an information system and performing an activity to complete a task, a part of the current working context is extracted based on the relevance of the working context's facets for the planned activity (or activities). This subset of the working context is called the context-of-use. Section 5 describes the selection and the use of the context-of-use in more detail.

4 PERSONAL WEB CONTEXT: A UUCM USE CASE

In this Section, we discuss the Personal Web Context (PWC) as a use case of the UUCM. The Personal Web Context combines the relationship and the task dimensions of the UUCM and outlines an approach for collecting information for the relationship dimension of the UUCM.

4.1 The Personal Web Context

The PWC is a model of a user, similar in structure to an ego-centric network (see Section 2), that describes the salient inter-relationships between the user and a set of interacting units that exist within a user's information space. The PWC represents the Relationship Dimension of the UUCM (see Section 3.1). The goal of the PWC approach is to analyze known relationships in order to infer and explicate implicit relationships between resources. This results in new "neighborhoods" that can be exploited for various kinds of services in the context of community events such as:

- Determining whom one should meet at a conference or trade fair in order to expand one's network of people [59];
- Recommending either relevant documents, or other people with whom to collaborate [23,38];

- Supporting a personalized conceptualization of an environment [40,51, 56];
- Supporting analytical activities that are crucial in business intelligence [31];

In the following, we discuss our approach to determine the PWC. We use the scientific research domain as an example application.

4.2 Information Model of the Personal Web

The underlying structure of the PWC is a domain ontology describing relevant entities in a domain. In our case, the domain ontology is built from the integration of existing ontologies describing the scientific community domain [37,44, 47]. Some examples of the domain resources in the ontology are: information objects (i.e. books, periodicals, and recorded presentations, technologies), people, organizations, and events. Typical relationship types in the domain are characterized by: a person being a member-of an organization; researchers who supervise graduate students; projects which produce a technology; a person author-ing a publication.

The domain ontology provides the basis for creating graph-based structures which we refer to as Resource Networks by defining the available resource and relationship types. Resource Networks are graph-based information models in which the nodes represent resources in the domain of various types and the directed edges indicate typed relationships between connected resources. The Resource Networks are used for several purposes. First, the Resource Network describes the relevant information space. Second, it is used as a starting point for expanding the PWC. Ideally, it uses the same or compatible relationship types as the UUCM for describing the Relationship.

Different approaches can be used to collect the data for a Resource Network from information spaces. On the one hand, metadata provided by the user like document authors, keywords etc. can be used to populate the Resource Network and, on the other hand, data for filling and updating the Resource Network can be collected by automatic analysis of information collections like DBLP [14].

4.3 Expanding the Personal Web Context

The goal of the PWC approach is to discover and expose the implicit relationships that link a user, by composition of relationships, to other resources that expand the PWC with new resources that were previously not in their PWC. The first step to extending the PWC is to extract pairs of relationships from the Resource Network which are connected by a common node. These pairs are called composition candidates (see Figure 2). After the extraction, we have a large set of composition candidates which can potentially contribute to expanding the Personal Web Context.

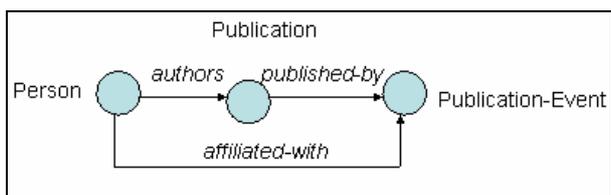


Figure 2: Composition Candidate

The next step in expanding the Personal Web is to consider a match-making process that supports discovering sequences of relations from which a single weighted relation can be compositionally inferred. For this purpose, composition rules are used. In the rules the relationships are represented as quadruples of the form (Subject, Predicate, Object, weight). Subject, Predicate and Object represent an RDF resource [49], Predicate represents the relationship types and weight represents the probability assigned to the Predicate linking the Subject to the Object. Figure 3 shows example rules which are discussed in more detail below.

In the match-making process, the rules are used: to decide which compositions candidates to use; to determine the type of composed relation and to compute the weight of the composed relation. A weight for a composed relation measures the likelihood that the composed relation holds. The measure for a composed relation is determined by: 1) the type of composed relations, and 2) the weights of the composing relations. The composed relationship acts as a surrogate for the sequence it replaces. If one of the resources involved in the composition is part of the PWC, the composition links portion of the user's existing PWC with entities in the information space.

The approach we have taken to manage the set of composition candidates is to incorporate task-specific domain knowledge so we can exclude some of the composition candidates from the search space. Combining the task dimension provides us with further facts that increases the available contextual information and allows us to restrict the search process, for example by fixing the type of the terminus nodes, or only including nodes that are related based on a given subject matter.

The following example illustrates how the relationship and task dimension can be combined. Given a task: "Find potentials collaborators at this conference for a proposal in Grid computing".

The additional facts that we get from including this type of task restrict the search process to finding: 1) the existence of a relation between the information seeker and a potential collaborator, 2) the strength of such a relation and, 3) a measure of the potential's determinant role (see Section 2) in the community. Additionally, since we are interested in proposal preparation, we focus only those relations in the network that exemplify a person as an active or influential community member for the subject Grid Computing. The additional context provided by the tasks, supports the use of example rules shown in Figure 3. For each rule in Figure 3 we make the following assumptions: all a_i are persons, all d_i are documents, all p_i are probabilities, and all e_i are events. Furthermore, we assume that an event, e_i , is a serial event devoted to a given subject matter.

In Rule 1, we assume that the number of authors for a document is known. Domain specific knowledge about publication activities allows us to assert with a certain probability, that one of the authors will attend the event to present the publication. This rule restricts the use of composition candidates to those that lead to persons who attend the event as a presenter.

In Rule 2, restricts the use relations that exemplify a person as an active or influential community member. One domain-specific criterion which can be used to determine influence is the frequency and recency of publications in an event which is devoted to a given subject. The relationships publication-

frequent and publication-recent are inferred by other rules that use the author's relation and the connection between the document and the event (the published-by relation) as a starting point. Other criteria could include heading an organization or project.

<p>Rule 1: attend-as-presenter</p> <p>Assume: $n = d_1.\text{number-of-authors} = \text{number of authors for document } d_1$</p> <p>$(a_1, \text{authors}, d_1, p_1) \wedge (d_2, \text{published-by}, e_1, p_2) \wedge (d_1 = d_2) \rightarrow (a_1, \text{attend-as-presenter}, e_1, p_3)$ where $p_3 = 1/n * (\text{MIN}(p_1, p_2))$</p> <p>Rule 2: determinant-author</p> <p>$(a_1, \text{publication-frequent}, e_1, p_1) \wedge (a_1, \text{publication-recent}, e_1, p_2) \rightarrow (a_1, \text{determinant-author}, e_1, p_3)$ where $p_3 = \text{MIN}(1, p_1 + p_2)$</p> <p>Rule 3: meeting-presenter</p> <p>$n = e_1.\text{number-of-participants} = \text{the number of participant in event } e_1$</p> <p>$(a_1, \text{presenter}, e_1, p_1) \wedge (a_2, \text{presenter}, e_2, p_2) \wedge (e_1 = e_2) \wedge (a_1 \neq a_2) \rightarrow (a_1 \text{ meeting-presenter } a_2, p_3)$ where $p_3 = 10/n * \text{MIN}(p_1, p_2)$</p>

Figure 3: Example Rules

Using Rule 1 and Rule 2, we can discover and recommend the implicit relationships that link the information seeker (for whom the system knows will attend the event) to other prominent persons that have certain probability of attending the same event as a co-participant.

The use of the rules for the above examples can be sketched as follows: we insert a link of type attend into the PWC that links the user to the event, e_1 , under consideration. As a next step, we use Rule 2 to determine the active community members for the event. In the next step, Rule 1 can be used to determine for each of the members, the probability of the active member attending the event in the case he has an accepted publication for the event. This results in a relationship attend-as-presenter with the respective probability. Finally, the relationship will be combined with the initially inserted relationship attend to a co-attend relationship between the user and the community member.

Rule 3, is representative of another kind of rule which allows us to estimate the probability of meeting a person at the event by chance. In this rule, we assume that the number of event participants is known and one meets about 10 people at a given event. These are example rules and we are currently developing further ones.

This example demonstrates that the use of relationship and the task dimensions for the UUCM and can be used to support a match between a systems users needs and the available items in an information space.

5 THE CONTEXT-PASSPORT APPROACH FOR CROSS-SYSTEM PERSONALIZATION

This section introduces an approach for cross-system personalization based on a passport metaphor. The context passport is a compact representation of the user's current context model for cross system personalization. It also contains the activities chosen by the user to be performed in order to fulfill the tasks allotted. It contains ontologically-arranged information about the user's current tasks and related activities, his cognitive patterns (skills, area-of-interest etc), the environment (time, place, device used), and his personal web of the people and relationships involved, following the UUCM model for user context modeling (See Section 3).

In order to use the context passport for cross system personalization, the user takes the context passport and "presents" it to an information system (IS) that the user wants to use. Since the context passport is bound to a shared ontology, there is a chance the IS can partially interpret the context passport using a mediator architecture [19]. As a result of this partial interpretation, two flows of information are possible: one from the context passport to the IS, and secondly from the IS to the context passport. The first flow helps the IS to better understand what the user requires from the IS, since the context passport refers to the task model, activities and also other information about the user context model. The second flow arises due to the interaction between the IS and the user which changes the state of the user's context. The purpose of this flow is to update the context passport with the feedback from the interaction. The approach outlined here so far does not specify the details of second flow.

We describe the approach as follows: we start with the box metaphor for context modeling, and then we outline how to apply this metaphor for the context passport and the formalization for selecting activities and context-of-use. Finally, we describe the approach for using the context passport in a cross system personalization environment.

5.1 The Box Metaphor

In [21], the notion of context dependence is illustrated by introducing the metaphor of the box (see Figure 4).

$$P1=V1 \dots Pn=Vn \dots$$

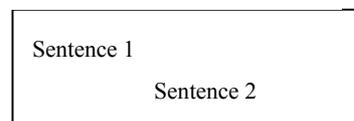


Figure 4: The Box Metaphor

A context dependent representation can be split into three parts: a) inside the box, a collection of linguistic expressions that describe a state of affairs or a domain; b) outside the box, a collection of parameters $P1, P2, \dots$ and finally, c) a value V_i for each parameter P_i . The intuition is that the content of what is inside the box is partially determined by the values of the parameters associated with that box. For example, in a context in which the date is 21st April, 2003 and the location is Paris,

any occurrence of ‘today’ will refer to 21st April, 2003. Thus a statement like ‘It will rain today’ in this context is completed by the presence of date and location as a context parameter.

The meaning of a context dependent representation is partly encoded in the parameters outside the box, and partly in the sentences inside the box. Reasoning mechanisms can be supplied for altering the balance between what is explicitly encoded inside the box and what is left implicit (i.e. encoded in the parameters). Intuitively, we can move information from the collection of parameters outside the box to the representation inside the box, and vice versa. These two operations are called push and pop to suggest a partial analogy with the operations of adding (pushing) and extracting (popping) elements from a stack. In one direction, push adds a contextual parameter to the collection outside the box and produces a flow of information from the inside to the outside of the box, that is part of what was explicitly encoded in the representation is encoded in some parameter. In the opposite direction, pop removes a contextual parameter from the collection outside the box and produces a flow of information from the outside to the inside, i.e. the information that was encoded in a parameter is now explicitly represented inside the box. This is a kind of de-contextualization.

5.2 Using the Box Metaphor with the Context Passport

The context passport contains the facets along the four dimensions described in Section 3 which can each be formalized as a set, and facets for each dimension grouped in the relevant set as a key value pair. For the purposes of formalization, facet properties like facet qualifier, value probabilities etc are not taken into account. However these facet properties will be used in other steps of the approach.

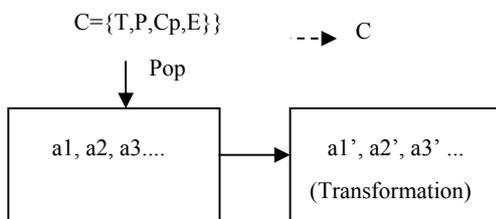


Figure 5: The Box Metaphor for the Context Passport

We apply the box metaphor to the context passport as follows: instead of statements, the chosen activities are inside the box, the working context facets parameters are outside the box, along with their values. The UUCM dimensions and their facets determine the available parameters outside the box. Also, inside the box are the activities that a user might need to perform in order to fulfill allotted tasks. These activities are chosen from the task, or manually chosen by the user from a collection of activities.

In the figure, the dimensions are represented by symbols T (Task dimension), P (Relationship dimension), E (Environment dimension) and Cp (Cognitive Pattern Dimension). Each of T, P, Cp and E are sets with key value pairs of their facets or attributes, along with their assigned values.

$T = \{(t1, tv1), (t2, tv2), \dots\}$ t is a facet, tv is the related facet value of the Task Dimension,

$P = \{(r1, re1), (r2, re2), \dots\}$, r is a facet, re is the related facet value of the Personal Web Dimension,

$C_p = \{(Cp1, Cpv1), (Cp2, Cpv2), \dots\}$, Cp is a facet, Cpv is the related facet value of the Cognitive Patterns Dimension,

$E = \{(e1, ev1), (e2, ev2), \dots\}$, e is an facet, ev is the related value of the Environment Dimension.

The interaction between the user and the IS using the context passport can be summed up as follows

1. The user goes to an information system and presents the context passport.
2. The IS can partially interpret the user's requirements.
 - a. Activities contained in the context passport which can be supported by the IS are selected.
 - b. The relevant context-of-use is extracted.
 - c. The activities are 'transformed' using the context-of-use.

These steps are described in detail in the following paragraphs.

3. The IS performs activities based on information derived from the context passport.
4. The feedback from the IS as a result of interaction is used to update the context passport and keep it up to date.

In a cross-system personalization environment, accessing a given information system may require the use of some part of the User's context. The context-of-use so required will depend on the activities that the user needs to perform. However, the given information system may not support all the activities that the user wishes to perform. Therefore, the applicable activities that can be supported have to be selected before.

Consider the Selection and Transformation functions as defined below:

Activity Selection function:

$$S_1(C, IS) \rightarrow A, A = \{a_1, a_2, \dots\}$$

(IS = Information System, A = Activity list,
C = context passport)

Context Selection function:

$$S_2(C, a_j) \rightarrow C' \text{ (modified context)}$$

Activity transformation function:

$$F_a(S_2(C, a_j), a_j) \rightarrow a_j'$$

After the activities supported by the current information system have been selected, the Context Selection function is used to identify which parameters from the context are relevant for this given activity. As introduced in the box metaphor, the selected parameters are then popped and put inside the box. Once the activity and the relevant parameters are known, the activity can be transformed and the transformed activity performed for the user. For example, if the activity is an information seeking activity, then popping the location from the user context can transform the information seeking into a geography based search (i.e. to find all hotels in the user's current city).

When the user now visits another information system, similar procedure is repeated. The activities not supported by previous

system, but supported by this system can be selected by the Activity selection function. The context-of-use is then extracted from the context passport and the selected activities are transformed and performed. It is also possible that case this IS can support previously unfinished tasks or subtasks. The unfinished tasks are those which have been already performed on an information system, but the result was not satisfactory. In the detailed approach which is under development, the history of the unfinished activity will be used to trace the progress made so far while performing this activity.

What remains at this stage is the process of updating the context. While this is not clearly outlined at this moment, possible solutions involve developing a protocol for passing update information to the mediator which then updates the context passport. Most existing information systems would not be able to follow this protocol directly, so a wrapper-based architecture could be used to implement the update protocol. Newly developed systems supporting cross system personalization across a federation of systems or services, could implement this protocol directly and thus not require a wrapper. As standards like Web services become more popular and information systems start supporting them, the wrapper could combine the user interface and the context update feedback procedure. This might also make tracking user activities related to this system possible.

6 CONCLUSIONS AND FUTURE WORK

In this paper we presented a Unified User Context Model (UUCM) that incorporates the various aspects that are relevant in capturing the characteristics of a user and his current situation in a flexible and extensible way. The UUCM provides a basis for the realization of cross-system and cross-service personalization approaches that enable the exchange and reuse of user profiles across system boundaries enabling improved support for multi-step information handling activities. The introduction of the UUCM, thus, is a contribution to the first question posed in the introduction of this paper. As a concrete use case for the UUCM we presented the Personal Web Context that combines the Task and the Relationship UUCM dimension. For the Personal Web Context we discussed ways to fill a relationship-based user context profile with data (collecting and updating user context profile data) and we also presented possible applications in the context of community events. Also, we will draw upon the inherent properties of the relations to determine probabilities that can be associated with each relation as well as the composed relations. The work presented in this part of the paper, thus, is a way to deal with the challenge addressed by the second question posed in the introduction of the paper.

Of course, the Unified User Context Model is only the first step on the way to cross-system personalization. As first ideas for the next steps in this direction and as a contribution to the third question raised in the introduction, we, therefore, presented an approach based on the context passport metaphor. The context-passport includes user and context information based on the UUCM. The relevant task-dependent context-of-use is selected from the context passport and applied for activity selection and transformation for better meeting the information needs of the user. The presented approaches are work in progress that still requires further investigation in several areas. Some important

issues for further research and development activities in cross-system personalization support are.

- Development of further methods for the systematic collection data for the different dimensions and facets of user context profiles based on the UUCM; This work will take into account existing methods for collecting user profile data;
- Development and evaluation of further rules for the composition of relationships for the Personal Web Context approach based on an analysis of relevant domain knowledge
- For the context passport approach rules for the selection and transformation of information access activities based on the current context-of-use have to be developed.
- Investigation in methods for ensuring the adequate evolution of user context profiles; this includes methods the collection of new profile data in a cross system personalization context, but also the definition of facet-specific aging processes for the profile content. These activities will be based on existing work dealing with information decay [43].
- Evaluation of adequate visualization methods for the Personal Web Context and its interaction with the Resource Network of the examined information space.

Furthermore, we are working on the further refinement and a prototypical implementation of the context passport architecture.

7 ACKNOWLEDGEMENTS

This research is funded in part by the DELOS: a Network of Excellence on Digital Libraries Proposal no. G038-507618 Sixth Framework entitled: Technology-enhanced Learning and Access to Cultural Heritage.

REFERENCES

- [1] Alexandria Digital Library Project. (ADEPT) *Annual Progress Report Of The Alexandria Digital Earth Prototype Project*. 30th June, 2002.
- [2] Anyanwu K, Sheth A.: *The p Operator: Discovering and Ranking Associations on the Semantic Web*. SIGMOD Record 31(4): 42-47 2002.
- [3] Artale, A., Franconi, E., Guarino, N., and Pazzi, L. 1996. *Part-Whole Relations in Object Centered Systems: An Overview*. Data and Knowledge Engineering, 20(3): 347-383. Elsevier.
- [4] Baker, Collin F. and Ruppenhofer J.: *FrameNet vs Frames vs. Levin's Verb Classes*. In Proc. of the 28th Annual Meeting of the Berkeley Linguistics Society, February 15-18, 2002.
- [5] Baudisch P., *Dynamic Information Filtering*. GMD Research Series 2001, No. 16, GMD Forschungszentrum Informationstechnik GmbH, Sankt Augustin.
- [6] Benerecetti, M., Bouquet, P., and Ghidini, C. *Contextual Reasoning Distilled*. Journal of Theoretical and

- Experimental Artificial Intelligence, 12(3):279--305, July 2000.
- [7] Callan J., Smeaton, A., et. al.: *Personalization and Recommender Systems in Digital Libraries*. DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries, June 2003. (http://www.ercim.org/publication/workshop_reports.html).
- [8] CC/PP: Composite Capability/Preference Profiles (CC/PP): Structure and Vocabularies. (<http://www.w3.org/TR/CCPP-struct-vocab/>).
- [9] Chaffin, R., and Herrmann, D.J., Relation Element Theory: A New Account of the Representation and Processing of Semantic Relations. Memory and Learning - The Ebbinghaus Centennial Conference, edited by Gorfein, D.S., and Hoffman, R.R., Lawrence Erlbaum Associates, 1987, pp. 221-245.
- [10] Cingil, I., Dognac, A., Azgin, A.: A Broader Approach to Personalization. Communications of the ACM, 43:8:136-139, August 2000.
- [11] Cohen, P.R., Loiselle, C.L.: Beyond ISA: Structures for Plausible Inference in Semantic Networks Proceedings AAAI88, St. Paul, MN, August, (1988).
- [12] Crabtree, B.I., Soltysiak, S.J.: Identifying and Tracking Changing User Interests. International Journal on Digital Libraries, 2:1:38-53, (1998).
- [13] Cyc.org. (<http://www.cyc.com/cycdoc/course/what-is-a-context.html>).
- [14] DBLP Digital Bibliography & Library Project (<http://dblp.uni-trier.de/>).
- [15] Egonet. (<http://survey.bebr.ufl.edu/EgoNet/>).
- [16] Fink, J., User Modeling Servers - Requirements, Design, and Evaluation, Dissertation vorgelegt zum Erwerb des akademischen Grades Dr. rer. nat. Datum der mündlichen Prüfung: 15. July 2003.
- [17] Foster I., Kesselman C., Nick J., Tuecke S.. Grid Services for Distributed System Integration. Computer, 35(6), 2002.
- [18] FreeLib, <http://p2pdl.cs.odu.edu/>.
- [19] Gerding G., Kuhn, W., A Functional Approach to a Wrapper- Mediator Architecture, (<http://musil.uni-muenster.de/documents/FuncMediation.pdf>).
- [20] Goker A. and Myrhaug H.I. User Context and Personalization. European Conference on Case Based Reasoning, Workshop on Case Based Reasoning and Personalization, 4-7 September 2002.
- [21] Giunchiglia, F. & Bouquet, P. . Introduction to contextual reasoning. An Artificial Intelligence perspective, in B. Kokinov (ed.), Perspectives on Cognitive Science, Vol. 3, NBU Press, Sofia, pp. 138–159. Lecture Notes of a course on “Contextual Reasoning” of the European Summer School on Cognitive Science, Sofia, 1996.
- [22] Hanani U., Shapira B., Shoval P., Information Filtering: Overview of Issues, Research and Systems. User Modeling and User-Adapted Interaction 11(3), 2001, 203-259.
- [23] Huang, Z., Chung, W., Ong, T.H., Chen, H.: A Graph-based Recommender System for Digital Libraries. In Proc. of the second ACM/IEEE-CS joint conference on Digital libraries, pp 65-73, Oregon, USA 2002.
- [24] Huhns, M.N., Stephens, L.M.: Plausible Inferencing Using Extended Composition. IJCAI Vol. 2, 1420 – 1425, 1989.
- [25] Kaplan, C., Fenwick, J., Chen, J.: Adaptive Hypertext Navigation Based On User Goals and Context. User Modeling and User-Adapted Interaction 3. Kluwer Academic Publishers, The Netherlands, 1993, pp. 193-220.
- [26] Kobsa, A., Generic User Modeling Systems, User Modeling and User-Adapted Interaction Journal 11: 49-63, 2001.
- [27] Konstan, J.A., Miller, B.N., Maltz, D. Herlocker, J.L. Gordon, L.R., Reidl, J.: Applying Collaborative Filtering to Usenet News. Communications of the ACM, Vol. 40 No. 4, (1997) 77-87.
- [28] Kramer J., Noronha S., and Vergo J.. A User-Centered Design Approach to Personalization. In ACM Computing Surveys, 08 2000.
- [29] Kraut, H., Galegher, J., Edigo, C.: Intellectual Teamwork: Social and Technological Bases for Cooperative Work, Lawrence Erlbaum, Hillsdale, NJ (1990).
- [30] Lenat D.B.. The dimensions of Context Space. (<http://www.cyc.com/publications>), 1998.
- [31] Lesser, E.L., Storck, J., Communities of Practice and Organizational Performance, IBM Systems Journal “Knowledge Management”, Volume 40, No. 4, 2001.
- [32] Levin, Beth. 1993. English Verb Class and Alternations: A Preliminary Investigation. Chicago: University of Chicago Press.
- [33] McCarthy, J., and Hayes, P. J. Some Philosophical Problems from the Standpoint of Artificial Intelligence. Machine Intelligence 4, Edinburgh University Press. (1969).
- [34] McCarty C. Structure in Personal Networks, Journal of Social Structure 3:1 (2002).
- [35] McDonald D.W.: Recommending Collaboration with Social Networks: A Comparative Evaluation. In Proc. of the ACM 2003 Conference on Human Factors in Computing Systems (CHI'03), Vol. 5, No.1 (2003) 593-600.
- [36] McTear M. F., User Modeling for Adaptive Computer Systems: A Survey. Artificial Intelligence Review, 7(3-4):157–184, August 1993.
- [37] Middleton S., Alani, H., Shadbolt, N.R., Rource, D.C.: Exploiting Synergy between Ontologies and Recommender Systems. In Proc. of the Eleventh International World Wide Web Conference Semantic Web Workshop, Hawaii, USA, (2002).
- [38] Middleton S., DeRoure, D., Shadbolt, N.: Capturing knowledge of User Preferences: Ontologies in Recommender Systems. In Proc. of the ACM K-CAP'01. 2001. Victoria, Canada: ACM Press (2001).

- [39] Motta, E., Reusable Components for Knowledge Modeling. IOS Press, Amsterdam. November 1999.
- [40] Nardi, B.A. et al. Integrating Communication and Information Through ContactMap. Communications of the ACM, 45(4). (2002)205-242.
- [41] Newman M. E. J., Ego-centered networks and the ripple effect. Social Networks 25, 83-95 (2003).
- [42] Neuhold, E. J., Niederée C., Stewart A.: Personalization in Digital Libraries : An Extended View. ICADL 2003: 1-16.
- [43] Niederée, C., Steffens, U., Schmidt J.W., Matthes, F.. Aging Links. In J. Borbinha, T. Baker, editors, Proceedings of the Fourth European Conference on Research and Advanced Technology for Digital Libraries (ECDL 2000), September 18 - 20, 2000, Lisbon, Portugal.
- [44] Ontology of SCIENCE:
http://protege.stanford.edu/ontologies/ontologyOfScience/ontology_of_science.htm.
- [45] Pazzani M.: A framework for Collaborative, Content-based and Demographic Filtering. Artificial Intelligence Review, 13 Vol. 5-6, 393-408, (1999).
- [46] Pednault, E. P. D., Representation is everything. ACM 43,8 (XXX-Month 2000), 80-83.
- [47] Protégé Ontologies Library
<http://protege.stanford.edu/ontologies/ontologies.html>
- [48] Pohl W., LaboUr - machine learning for user modeling, In Proc. of the 7th International Conference on Human-Computer Interaction). Volume B, pages 27-30, Amsterdam, 1997. Elsevier Science.
- [49] RDF/XML Syntax Specification (Revised) W3C Recommendation. (<http://www.w3.org/RDF/>).
- [50] Rumbaugh, J. Relations as semantic constructs in an object-oriented language, Conference proceedings on Object-oriented programming systems, languages and applications, 1987, 466-481.
- [51] Sack, W., Conversation Map: A Content-Based Usenet Newsgroup Browser. The 2000 International Conference on Intelligent User Interfaces (UIU 2000), ACM Press, 233-340.
- [52] Schreiber, G., Akkermans, H., Anjewierden, A., De Hoog, R., Shadbolt, N., Van de Velde, W., Wielinga, B., Knowledge Engineering and Management - The CommonKADS Methodology. MIT Press, pp. 455, Cambridge, Massachusetts 2000.
- [53] Schmidt, A., Beigl, M., Gellersen, S. There is more to context than location. In Proc. of the International Workshop on Interactive Applications of Mobile Computing (IMC98), Rostock, Germany, 1998.
- [54] Shapira, B., Shoval, P., Hanani, U., Stereotypes in Information Filtering Systems. Information Processing and Management, Volume 33, Issue 3, Pages: 273 – 287, (May 1997).
- [55] Sharma, A., A Generic Architecture for User Modeling Systems and Adaptive Web Services. IJCAI 2001 Workshop on E-Business & the Intelligent Web, Seattle, USA, 2001.
- [56] Smith M.A., et al., Conversation Trees and Threaded Chats. ACM 2000 Conference on Computer-Supported Cooperative Work (CSCW'00) (Philadelphia, PA, 2000) 97-105.
- [57] Souchon, N., Limbourg, Q., Vanderdonck, J., Task Modeling in Multiple Contexts of use. in Proceedings of 9th International Workshop on Design, Specification, and Verification of Interactive Systems DSV-IS'2002 (Rostock, 12-14 June, 2002), Lecture Notes in Computer Science, Vol. 2545, Springer-Verlag, Berlin, 2002.
- [58] Sowa, J. F., Knowledge Representation: Logical, Philosophical, and Computational Foundations. Brooks Cole Publishing Co., Pacific Grove, CA, 2000.
- [59] Tomobe, H., Matsuo, Y., Hasida, K., Ishizuka, M., Social Network Extraction of Conference Participants. WWW 2003 poster.
(<http://www2003.org/cdrom/papers/poster/p092/p92-tomobe.html>).
- [60] Trigg R., A Network-Based Approach to Text Handling for the Online Scientific Community , PhD Dissertation. (<http://www.workpractice.com/trigg/thesis-chap4.html>).
- [61] Tyler, S.W. and Treu, S., An interface architecture to provide Adaptive Task-Specific Context for the user. International Journal of Man, Machine and Studies, 1989(30): p. 303-327.
- [62] Vassileva, J., A Practical Architecture for User Modeling in a Hypermedia-Based Information System. In Proc. of the 4th International Conference on User Modeling, pp 115-120, UM'94, Cape Cod, MA, 14-19 August 1994.
- [63] Wahlster, W. and Kobsa, A., User models in dialog systems. User Models in Dialog Systems, pp 4–34. Springer, Berlin, Heidelberg, 1989.
- [64] Wasserman, S., Galaskiewicz, J., Eds.: Advances in Social Network Analysis Sage. Thousand Oaks, California, (1994).
- [65] Wenger, E., Snyder, W.: Communities of Practice: The Organizational Frontier. Harvard Business Review, January-February (2000).
- [66] Wierzbicka, A., Apples are not a "kind of fruit": The semantics of human categorization. American Ethnologist, 313-28, (1984).
- [67] Winston, M., Chaffin, R., and Herrmann, D. A Taxonomy of Part-Whole Relations. Cognitive Science, 11: 417-444, 1987.