

SPECTER: Building, Exploiting, and Sharing Augmented Memories

Alexander Kröner, Dominik Heckmann, and Wolfgang Wahlster
German Research Center for Artificial Intelligence (DFKI)
66123, Saarbrücken, Germany
{firstname.lastname}@dfki.de

Abstract

A comprehensive log of the user's behavior together with corresponding context descriptions allows adaptive systems to learn about users, to identify their habits, and to improve the quality of user support. In addition, users can apply such knowledge to learn from others and about themselves. We describe in this article how these issues can be addressed by the means of augmented memories created from sensor data automatically captured in an intelligent environment. After discussing how to represent and process such memories, we show how they can be exploited by adaptive systems as well as by their users. One promising way of exploiting memories is to share them with others; this is reflected by the last part of our contribution where we explain how ubiquitous user modeling may enable the sharing of information inferred from memories.

1. Introduction

Distributed sensor networks enable systems to perceive information about the user's current situational context. By temporally and spatially dense monitoring of the user's behavior and her interaction with the environment and recording such information a history of contexts can be established, which may then be applied to augment the user's episodic memory. Such *augmented memories* are the foundation of user-centered applications, which take into account the user's habits, extend her perception with precise event logs, allow for the identification of missed opportunities and provide scrutable behavior explained by the course of events.

This promising approach leads to diverse research challenges. For instance, while the foundation of augmented memories is raw sensor data, recording this data alone usually won't be sufficient to build a rich memory. Therefore methods are required to transform sensor data into symbolic representations, which can be processed and understood by both the user and the system. However, depending on the

quality of the underlying data and the applied interpretation methods, a system may perceive a situation differently than the user. Thus a user needs a means by which to reflect on the system's perceptions and to allow for adjustments on personal experience.

These issues are related to the exploitation of augmented memories for personal use. However, humans keep many memories not necessarily for personal reasons. Making small talk about a weekend with colleagues, sending information about a special offer to friends, telling the family the day's events — there are many (and very diverse) situations where people share their memories. The sharing of augmented memories is a way to enrich this process. Its potential benefits are opposed by specific challenges: besides the obvious issues of privacy and trust, the personal nature of memories may result in diverse and contradicting structures, whose exchange requires sophisticated integration mechanisms.

Together with partners from universities and industry, DFKI is conducting research in the field of augmented memories. In this article we focus on work conducted in the project SPECTER. Our main contributions include a model for building augmented memories, and an approach to the collaborative introspection of situated user support. In addition, we discuss how ubiquitous user modeling may be applied to share information from augmented memories between users, and complete our contribution with a look into the issues related to sharing augmented memories, a topic to be explored in the recently started project SHARED LIFE.

Section 2 begins with an example scenario, which illustrates certain ideas of SPECTER. Then, in Section 3, we describe our approach for building augmented memories from sensor data. The introspection of situated user support in Section 4 forms the background for a description about how to exploit augmented memories. We continue in Section 5 with a discussion of existing and planned mechanisms for sharing such memories, and conclude with a brief review of related work.



Figure 1. SPECTER records for each event contextual information (location, time), action, and subjective ratings.

2. A Shopping Scenario

Before we proceed to the more technical aspects of augmented memories we want to provide an example of how they are built and exploited in everyday life. The following shopping scenario provided the background for the implementation and evaluation of the latest SPECTER prototype.

The domain of our test case is shopping for audio CDs, with a particular focus on soundtracks. People come in contact with soundtracks through various situations — e.g., in a cinema with the family, while watching a DVD at home, or while browsing an Internet store. Sometimes they have a precise idea of the music in question, and sometimes, they have never heard it. This background serves as the scenario in which a user exploits augmented memories by means of SPECTER in order to learn more about soundtracks that might be of interest. The scenario basically consists of the following elements:

(1) In order to actively acquire information about the soundtracks, our user browses the Web pages of an Internet store. SPECTER unobtrusively records these actions and assigns to each CD examined by the user a subjective rating based on the user’s attention.

(2) The user reflects about past events in her augmented memories. SPECTER supports this process with various ser-

vices. For example, links tailored for the domain allow the user to acquire information related to the subject CDs from the Web and from her augmented memories. In addition, the system lets the user review and adjust its perceptions (e.g., the mentioned ratings) and its supporting actions.

(3) The user can rely on SPECTER’s support while shopping. Based on situational preferences, the system provides a listing of services related to the CDs being considered — e.g., if she is in her favorite shop and has spare time, the system may inform her of special offers on similar CDs. The user may exploit her augmented memories in several additional ways. For instance, if a CD is unknown to the user, the system may provide a list of similar CDs known by means of augmented memories, and thus provide a clue about its content. The other way around, the user can tell SPECTER to provide some examples of CDs she likes to the shop in order for them to suggest similar CDs as yet unfamiliar to her. All these actions contribute to her augmented memories and may therefore later become the subject of an introspection process.

3. From Sensor Data to Memories

The first step towards the creation of augmented memories is the automated recording of contextual information as perceived from various types of sensors. In our example scenario, each CD is an RFID-tagged smart object, which allows tracking its presence within the store areas (shelf, basket, cashier). Optionally, these objects may be anthropomorphized in order to facilitate the human-environment communication (cf. [22]). The user’s location may be determined using IR, RFID, and/or GPS (see [6]). Biosensors (e.g., electrocardiogram (ECG) electromyogram (EMG), electrodermal activity (EDA), and acceleration (ACC) sensors) provide further information about the user’s state, which is applied for choosing an appropriate communication channel and for automatically evaluating events (cf. [5]). Finally, Web services allow the system to acquire rich context information (e.g., the current weather or important events from RSS feeds), which may later on serve as an access key to the memory. In addition, such services are used to realize certain domain-specific features within the user’s environment. For instance, SPECTER assists its user with services implemented by [1], such as a similarity search for CDs.

Each of these input sources is linked to a so-called *RDF store* (see [18]). Such a store provides an RDF-based interface to a sensor-specific memory, which is decoupled from the user’s augmented memories. Two advantages provide the rationale for this separation:

Efficiency: sensor memories are not bound to SPECTER’s RDF-based implementation of the augmented memory. This is of special interest due to the diversity of the per-




Figure 2. User model dimensions shown in the Ubis Ontology Browser (excerpt).

ceived data, which may range from raw mass data of biosensors to rich information retrieved from Web services. For the same reason, sensors may implement their own abstraction methods — e.g., a simple mapping is performed in order to translate GPS coordinates into semantically meaningful locations, whereas input from biosensors is processed on a mobile system by means of dynamic Bayesian networks.

Flexibility: for various reasons, the connection between the augmented memory and some sensors may sometimes be lost (e.g., technical issues, trust issues). In such a case, the user should be allowed to complete the augmented memories at a later time using records from the sensors’ memories.

3.1. Modeling Perceptions

Sensors provide the system with *perceptions*; their RDF-encoded content contains simple statements such as “user reaching shelf” or “user holding CD”. At any given point of time, the set of all available sensors’ latest perceptions defines the *context* of an event in SPECTER.

Information contained in perceptions references an ontology based on the IEEE SUMO and MILO [17]. The user’s state is modeled using the general user model ontology (GUMO, cf. [10]), a mid-level ontology which provides applications with a shared vocabulary for expressing statements about users. Furthermore, in order to facilitate the exchange of GUMO statements between different appli-

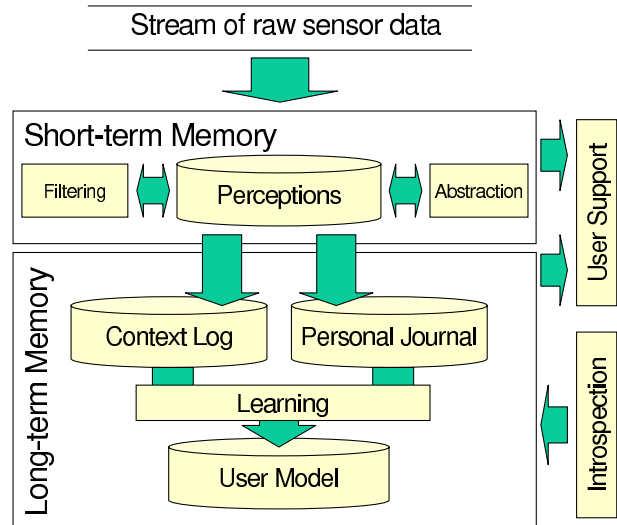


Figure 3. Building memories from perceptions in SPECTER.

cations, the ontology provides a means of combining such statements with meta statements, e.g., about privacy, trust, and expiry issues. GUMO statements reference a variety of dimensions that describe user properties.

Figure 2 shows some of the basic user dimensions within GUMO, which include `contact information`, `personality`, and `emotional state`. These concepts are the topmost level of a broad range of specialized concepts; in our example, several specializations of `emotional state` are shown. The complete ontology can be reviewed and edited online with an ontology browser provided at <http://www.gumo.org/>.

3.2. Processing Perceptions

In order to describe how perceptions are processed and stored we have developed a *memory model* (see Figure 3). While our model uses concepts known from biology and psychology, it is not intended as a simulation of the human brain. Instead, our goal is to offer a model which permits the description of components and processes used in our approach to the realization of augmented memories.

In this model, incoming perceptions are stored in a short-term memory, which serves two main purposes. First, it is of special relevance for the recognition of situations and thus situated user support. The facts stored in the short-term memory model the user’s current context; a BDI planner (JAM, cf. [12]) matches this context against patterns of service bindings specified in the user model (more about this in the next section). The second — and somewhat related — purpose of the short-term memory is the identification

of episodes, which trigger the construction of entries in the user's personal journal.

This leads to our two-fold approach to a long-term memory for intelligent environments. First of all, perceptions are stored in a *context log* without further change. This log works primarily as the system's memory. It is linked to the *personal journal*, which consists of entries representing episodes created from one or more perceptions. For instance, the perceptions "user holding A" and "user holding B" can be combined to the episode "user comparing A and B". The creation of journal entries is an abstraction process, which is performed using pre-authored rules expressing commonsense knowledge and domain knowledge.

The personal journal offers the user the right level of granularity and abstraction when accessing SPECTER's memory. In the course of introspection, the user may interact with journal entries in various ways, e.g., by assigning ratings in order to express a personal impression related to a situation, or by applying them in order to customize SPECTER's user support.

Both context log and personal journal provide input to the processes performed in the short-term memory. In addition, they are exploited for building and refining the user model.

4. Exploiting Memories

In order to illustrate how augmented memories can be exploited we will focus in the following section on situated user support in SPECTER.

4.1. Feedback and Introspection

SPECTER supports its user with memories and services that could be relevant for the user's current context — e.g., if the user is inspecting some CD in a shop, the system might inform her that it is cheaper than other known offers, or perhaps recommend similar CDs in the current shop. Here, relevance is determined by several factors. One of them is explicit or implicit user feedback. Such feedback is captured in ratings (e.g., evaluation, affect), which are attached to entries in the personal journal. Ratings are assigned either manually by means of a GUI or automatically based on feedback from biosensors and domain-specific heuristics.

Ratings assigned by the system may differ from the user's perception, e.g., due to incorrect sensor data or inappropriate heuristics. Especially in the case of an "untrained" system, the mapping of situations to services might require an adaptation to the user's personal needs as well. The user may address such issues by performing an *introspection* of her augmented memories.



Figure 4. SPECTER offers services related to memories and environment (here the CD shop "Clyde's").

We think of introspection as "...a process of inward attention or reflection, so as to examine the contents of the mind..." (cf. [19]). In the case of SPECTER, introspection consists of processes in which the user and/or the system explore the long-term memory in order to learn about the course of events. From the user's point of view this includes the option to explore and to rate journal entries, including those produced in response to system actions. From the system's point of view introspection is an opportunity to refine, collaboratively with the user, the user model (cf. [2]).

Feedback from our test user group indicated that an introspection should be possible in a mobile setting (in order to make use of spare time, e.g., during a train ride) as well as in a desktop setting. This is reflected in SPECTER's user interface to augmented memories, which consists of two major components.

On a mobile device, a *journal browser* lets the user explore and evaluate memories (see Figure 1), and exploit them for purposes of communicating with her environment (see Figure 4). Furthermore, user support can be customized on a high level (services can be disabled or configured to inform the user before execution, etc.). Finally, by setting *reminder points* the user may indicate with a single click the need for a deeper introspection of the current situation.

Such introspection can be performed on a regular desktop PC using an *introspection environment*. It provides a rich user interface consisting of a virtual character, a presentation screen, and the journal browser (see Figure 5,

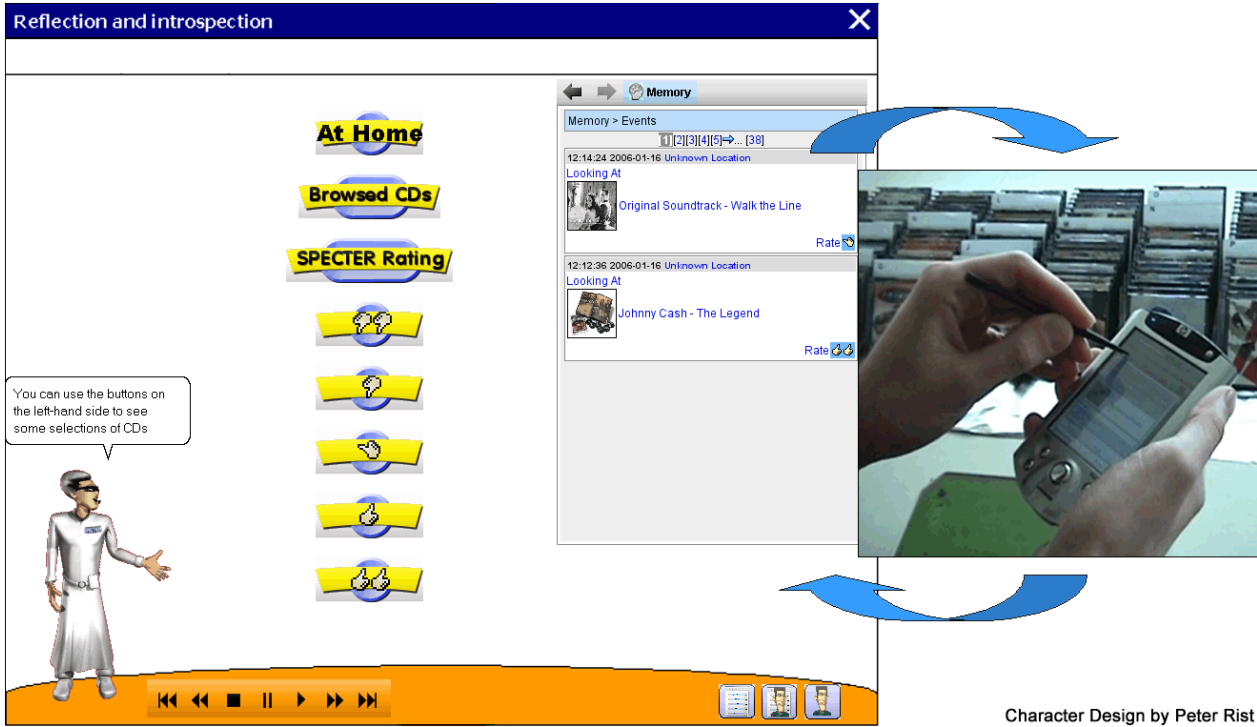


Figure 5. The introspection environment assists the user in reviewing augmented memories and in exploiting these for user support in an intelligent environment.

from left to right). Combined with the planner already mentioned, these components enable a collaborative introspection of augmented memories. If the user is exploring the memory, then the system assists by offering event summaries as well as details and links to the memory and the Web. The latter point to external resources and services, and thus provide another means for retrieving and adding information to the memory.

These activities usually require the user's initiative. The system also proactively checks the memory for situations where clarification might help to improve the user model; examples of such situations are reminder points and user feedback obtained during the execution of some supporting service. When such situations are detected, the user is asked to enter the collaborative process described below.

4.2. Collaborative Critique of Situated User Support

If the user is not confident with the system's suggestions, or wants to set up a new service binding, there is a component explicitly designed for mapping situations to services. This component is described in detail in [3]; at this point we will only summarize its features and focus on its role within

the introspection process.

The purpose of this component is to provide a scrutable and easy-to-use interface that allows the user to interact with complex machine-learning processes without the need to deal with the technical subtleties of feature selection or data encoding. The key to our approach is to combine the system's capability to deal with the statistical relevance of a situation's features with the user's ability to name semantically meaningful concepts that can and should be used to describe the characteristics of a situation.

The result of our ongoing effort is a user interface which provides several interaction layers of varying complexity for combining services and situations. Especially relevant for a critique of a situation's features (and thus for configuring the execution of linked services) is the screen shown in Figure 6. This shows a list of features, which have been extracted using statistical methods from the memory, as it is presented to the user.

The user may critique this set by deselecting features, or by navigating in their semantic neighborhood using a graphical interface to the underlying ontology. In our example scenario, the user might want to inform the system that recommendations should only be retrieved in certain kinds of shops. A way to achieve this goal is to inspect the list of




Figure 6. A user interface for critiquing features selected by the machine learning system based on statistical relevance.

features and then to refine the shop’s features, e.g., by replacing the general shop by a more specific class.

5. Sharing Memories

SPECTER is tailored for the creation of *personal* augmented memories, and therefore implements no specific sharing features. However, the linking of SPECTER and UBISWORLD allows users to adopt information provided by others into their augmented memories and to support others vice versa with perception records and information derived from their memories. Here, the key is that SPECTER’s augmented memories can be interpreted as a collection of statements about users, and therefore realize a complex user model.

5.1. Sharing User Models

UBISWORLD is based on the new concept of *ubiquitous user modeling* which means that networked systems constantly track the user’s behavior at many times and in many ways. Wahlster’s definition in [9] points out that:

Ubiquitous user modeling describes ongoing modeling and exploitation of user behavior with a variety of systems

that share their user models.

These shared user models can either be used for mutual or for individual adaptation goals. Three additional concepts differentiate ubiquitous user modeling from *generic user modeling: ongoing modeling, ongoing sharing, and ongoing exploitation.* As intended in SPECTER, it covers and integrates user modeling for situated interaction in ubiquitous computing environments, as well as ongoing user modeling with a variety of systems and applications.

The concept of *sharing* is further split up within UBISWORLD into *exchanging* and *integrating* statements about users. The former is realized by a user model server (www.u2m.org) that provides a service-based architecture for distributed storage and retrieval of statements about users. The integration of statements is achieved with the accretion model according to [13] together with a multilevel conflict resolution method (see [9]), which also solves the problem of contradictory information. What statements can be retrieved and how they are integrated depends on several layers of metadata attached to the statements by means of reification.

From the outermost to the innermost layer, these are: administration, privacy, explanation, and situation. They establish a sequence of access constraints which have to be met in order to obtain the reified statement. The privacy layer in this sequence is of special interest. It implements the following privacy attributes: key, owner, access, purpose, and retention. The user model server checks these attributes in order to deliver as much information as possible without violating the user’s preferences. Combined with the other layers, complex situational access constraints can be established, such as “friends only & at my home & for personal purposes”.

An example of how users may inspect their user model is given in Figure 7. The two statements *Margeritta is vegetarian* and *Margeritta is female* are shown. The privacy editor is set to *public access*, which means everybody and every system is granted access; *commercial purpose*, which means that this information may also be used for product recommendations; and, *middle retention*, indicating the statement may be stored for one month.

5.2. Future Work

Our future work will address the question of how the sharing of augmented memories can contribute to the communication within small, potentially ad-hoc formed groups. We want to create an application that enables an automated or semi-automated sharing of memories within such groups, and allows users to extend their perception with the other ones of the group and to provide the group with their perceptions. Such an exchange mechanism must not only take into account situated access constraints on privacy and trust




Figure 7. A graphical user interface to inspect shared information and to edit privacy attributes in UbiWorld.

(e.g., distinguish between situations of everyday life and emergency cases), but also the structure of the group (e.g., to define experts or opinion leaders). Users should be put in a position to control such complex constraints. In addition, we want to assist them with mechanisms that extract sharing constraints from the behavior of group members. Furthermore, since we expect users to have unique memory structures, mechanisms will be required which allow them to explore different structures from their own.

These are some of the issues addressed in the project SHARED LIFE, which started in 2006 and will last three years. It will continue the research conducted in SPECTER with a strong focus on sharing augmented memories in an everyday life setting.

6. Related Work

The building of augmented memories in instrumented environments for the purpose of extending the user's perception and recall has been studied for more than 10 years (see, e.g., [16]; [8]). While this research has focused on user interface design for the retrieval of memories, other research has looked into ways of processing the contents of such memories so as to increase their accessibility to their owner (see, e.g., [11]; [7]).

The exploitation of augmented memories has been researched for diverse scenarios. For instance, automatically created nursing records can help to avoid medical accidents in hospitals (cf. [15]). In the project LIVING MEMORY [20], records of people's activities and access to community-related information are automatically processed in support of community-related behavior in relatively complex ways. An example of how memories can support social matching is offered by the system AGENTSALON (see, e.g., [21]). The system uses experience logs of participants in an academic conference in order to stimulate conversation via rather extraordinary means, involving animated characters.

The latter works also address sharing augmented memories. Closely related to our approach to the sharing of user models is *Customer Profile Exchange* [4], which is used in customer relationship management for the purpose of supporting information sharing about customers in enterprise-wide architectures. Another work related to the UBISWORLD data model is the *Composite Capabilities/Preference Profile* (CC/PP), see [14], which addresses the task of sharing user preferences.

7. Conclusions

Histories of interpreted sensor data provide a means of taking the past into account when providing user support in intelligent environments. In this article we discussed the SPECTER system, which addresses the building, the exploitation, and the sharing of such augmented memories.

As we consider perceptions to be the foundation of episodic memories, we began with a review of their RDF-based representation as applied in our approach. With GUMO, we proposed an ontology, which has been especially tailored for user model representation in ubiquitous computing. Following that, we proceeded to SPECTER's memory model — a proposed architecture for software using augmented memories. This model formed the background for our discussion of how to exploit the collected data. At this point, we focused on the collaborative introspection of situated user support, whose central element is a component for the specification of situations. This component combines statistical information, extracted by machine learning from the memories, with explicit user feedback on the relevance of features that are typical characteristically for a given situation.

Another way of exploiting memories is to share them with others. Our contribution concluded with a description of how UBISWORLD can make user models built from or including augmented memories accessible in an intelligent environment. We consider this approach to be an initial, productive step towards the sharing of augmented memories and plan, with the recently started project SHARED LIFE, to intensify our efforts in this area.

8. Acknowledgements

This research was supported by the German Ministry of Education and Research (BMBF) under grant 524-40001-01 IW C03 (project SPECTER) as well as by the German Research Foundation (DFG) in its Collaborative Research Center on Resource-Adaptive Cognitive Processes, SFB 378, Project EM4, BAIR. We would like to thank all project members for their valuable contributions.

References

- [1] Amazon Inc. Amazon e-commerce service, April 2005. <http://www.amazon.com/gp/aws/landing.html>.
- [2] S. Baldes, A. Kröner, and M. Bauer. Configuration and introspection of situated user support. In *LWA 2005, Lernen Wissensentdeckung Adaptivität*, pages 3–7, Saarland University, Saarbrücken, Germany, 2005.
- [3] M. Bauer and S. Baldes. An Ontology-Based Interface for Machine Learning. In J. Riedl, A. Jameson, D. Billsus, and T. Lau, editors, *IUI 2005: International Conference on Intelligent User Interfaces*, pages 314–316, New York, 2005. ACM.
- [4] K. Bohrer and B. Holland. *Customer Profile Exchange (CPEXchange) Specification*, October 2000. <http://www.cpexchange.org/>.
- [5] B. Brandherm, H. Schultheis, M. von Wilamowitz-Moellendorff, T. Schwartz, and M. Schmitz. Using physiological signals in a user-adaptive personal assistant. In *Proceedings of the 11th International Conference on Human-Computer Interaction (HCI-2005)*, Las Vegas, Nevada, USA, 2005.
- [6] B. Brandherm and T. Schwartz. Geo referenced dynamic Bayesian networks for user positioning on mobile systems. In T. Strang and C. Linnhoff-Popien, editors, *Proceedings of the International Workshop on Location- and Context-Awareness (LoCA)*, LNCS 3479, pages 223–234, Munich, Germany, 2005. Springer-Verlag Berlin Heidelberg.
- [7] B. Clarkson. *Life Patterns: structure from wearable sensors*. PhD thesis, School of Architecture and Planning, Massachusetts Institute of Technology, Cambridge, MA, 2002.
- [8] J. Gemmell, L. Williams, K. Wood, R. Lueder, and G. Bell. Passive Capture and Ensuing Issues for a Personal Lifetime Store. In *Proceedings of The First ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE '04)*, pages 48–55, New York, USA, 2004.
- [9] D. Heckmann. *Ubiquitous User Modeling*. PhD thesis, Department of Computer Science, Saarland University, Germany, 2005.
- [10] D. Heckmann, B. Brandherm, M. Schmitz, T. Schwartz, and B. M. von Wilamowitz-Moellendorff. GUMO - the general user model ontology. In *Proceedings of the 10th International Conference on User Modeling*, pages 428–432, Edinburgh, Scotland, Jun 2005. LNAI 3538: Springer, Berlin Heidelberg.
- [11] E. Horvitz, S. Dumais, and P. Koch. Learning predictive models of memory landmarks. In *Proceedings of the CogSci 2004: 26th Annual Meeting of the Cognitive Science Society*, Chicago, USA, August 2004.
- [12] M. J. Huber. Jam: a bdi-theoretic mobile agent architecture. In *AGENTS '99: Proceedings of the third annual conference on Autonomous Agents*, pages 236–243, New York, NY, USA, 1999. ACM Press.
- [13] J. Kay, B. Kummerfeld, and P. Lauder. Personis: A server for user models. In P. de Bra, P. Brusilovsky, and R. Conejo, editors, *Proceedings of the Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, LNCS 2347, pages 203–212. Springer-Verlag Berlin Heidelberg, May 2002.
- [14] G. Klyne, F. Reynolds, C. Woodrow, H. Ohto, J. Hjelm, M. Butler, and L. Tran. *Composite Capability/Preference Profiles (CC/PP): Structure and Vocabularies 1.0*. W3C Recommendation, January 2004. <http://www.w3.org/TR/2004/REC-CCPP-struct-vocab-20040115/>.
- [15] N. Kuwahara, H. Noma, K. Kogure, N. Hagita, N. Tetsutani, and H. Iseki. Wearable auto-event-recording of medical nursing. In *Proceedings of the Ninth IFIP TC13 International Conference on Human-Computer Interaction (INTERACT 2003)*, September 2003.
- [16] M. Lamming and M. Flynn. “Forget-Me-Not”: Intimate computing in support of human memory. In *Proceedings of FRIEND21, the 1994 International Symposium on Next Generation Human Interface*, Meguro Gajoen, Japan, Meguro Gajoen, Japan, 1994.
- [17] I. Niles and A. Pease. Towards a standard upper ontology. In *FOIS '01: Proceedings of the international conference on Formal Ontology in Information Systems*, pages 2–9. ACM Press, 2001.
- [18] M. Schneider, M. Bauer, and A. Kröner. Building a Personal Memory for Situated User Support. In *Proceedings of ECHISE 2005 - 1st International Workshop on Exploiting Context Histories in Smart Environments*, Munich, 2005.
- [19] M. Sharples, D. Hogg, C. Hutchison, S. Torrance, and C. Hutchinson. *Computers and Thought: A Practical Introduction to Artificial Intelligence*. Bradford Book, 1989.
- [20] K. Stathis, O. de Bruijn, and S. Macedo. Living memory: agent-based information management for connected local communities. *Interacting with Computers*, 14(6):663–688, 2002.
- [21] Y. Sumi and K. Mase. Supporting awareness of shared interests and experiences in community. *International Journal of Human-Computer Studies*, 14(6):127–146, January 2002.
- [22] R. Wasinger and W. Wahlster. The anthropomorphized product shelf: Symmetric multimodal interaction with instrumented environments. In E. Aarts and J. Encarnação, editors, *True Visions: The Emergence of Ambient Intelligence*. Springer, Heidelberg, Berlin, New York, February 2006.