

Augmenting Cognition With a Digital Episodic Memory

Alexander Kröner, Anthony Jameson, Michael Schneider, and Nathalie Basselin

Current technology makes it increasingly feasible for personal assistance systems to create an augmented episodic memory that supplements their users' own episodic memory. After considering the question of what such an augmented episodic memory might be useful for, we present a proof-of-concept system SPECTER that realizes this idea. We consider in turn the challenges of (a) automatically capturing and appropriately representing information about the user's actions and experiences; (b) enabling the user to review and revise the information captured; (c) offering functions based on the episodic memory that can help the user deal more effectively with everyday situations; and (d) determining in what situations to make these functions available proactively.

1 Introduction

The idea that human cognition can usefully be augmented by digital memories is old and familiar: Most people now rely on electronic appointment calendars, contact lists, and to-do lists to cope with the demands of everyday life.

It is a bit less obvious how a particular type of digital memory—a digital episodic memory, which essentially contains records of a person's past experiences—might be useful. Sure, it may be nice to reflect on the events of one's last vacation, especially if the episodic memory in question includes concrete media like photos; but in what ways can recollections from the past help us to perform tasks in the present and the future? This question is becoming increasingly relevant as the technological means for automatically capturing a person's actions and experiences increase (cf., e.g., [7]).

This paper summarizes research in a project whose primary goal was to investigate ways in which an automatically constructed digital episodic memory can be created and exploited to help people deal with the everyday world. Its relevance to the topic of AI and human cognition does not lie in any attempt to simulate human cognition. Instead, the strategy is to (a) understand some of the functions that a particular aspect of human cognition (episodic memory) has in some practically relevant situations; and (b) design and test an augmentation of that aspect of cognition that can serve the same functions, thereby enhancing the person's overall cognitive performance in certain types of situation. As we will see, the added value of this sort of support is greatest (a) with subtasks for which human cognitive processing is least well adapted and (b) in cases where the results of the system's augmented processing can be combined immediately with other types of processing by the system.

The system in question—called SPECTER—is in principle applicable in a wide variety of situations. For concreteness in this brief paper, we will refer to examples mainly from a shopping domain, in which the system was most extensively tested. We will also include occasional examples from other domains to illustrate the generality of the functions of the system.

To get an initial idea of how an augmented episodic memory can be useful in a shopping setting, consider the following scenario: You are browsing through the CDs in a CD store to-

gether with a friend, looking for something that you might like to add to your collection. You pick up a CD that looks appealing in some ways, though you have never heard of it before. Your friend says: "This CD is a lot like that one by 'Heavy Metal, Inc.' that you listened to part of on the Amazon website yesterday and didn't especially like." In this case, the reference to a previous experience is an efficient way of conveying a good deal of information about the newly encountered object. Note that your friend is not only reminding you of an experience of yours; she is also allowing you to benefit from some knowledge that she has that you do not have: knowledge about the properties of the CD that you are looking at.

Note also that part of the useful contribution of your friend is a matter of timing: You would not want to be reminded continually and indiscriminately of past experiences; reminders should come at a point at which you can make some use of them. Accordingly, one of the subtasks that need to be handled by a system like SPECTER is that of identifying appropriate situations for presenting reminders. An appropriate *triggering rule* for a reminder in our example scenario might be "While I'm shopping, remind me of similar products that I've recently seen and either liked or disliked".

As we will see below, there are a number of other ways in which an augmented episodic memory can enhance a shopping experience. Since these functions combine aspects of reminding and recommending, we have coined the term *recomindation* to refer to them.

In the rest of this paper, after discussing some relevant previous research, we will discuss in turn several components of SPECTER which together yield the functionality required for *recomindation* and for analogous functions in other domains.

2 Related Work

2.1 Recommender Systems

The significance of exploiting an augmented episodic memory can be illustrated by a brief comparison of the *recomindation* paradigm with more familiar approaches to product recommen-

dation.¹ Although there are exceptions, the general pattern is for a recommender system to analyze certain types of information about the user and then to recommend one or more actions or objects to the user. Even where the system makes significant use of information about past actions of the user, such as purchases and ratings, the user is typically not reminded of these past experiences. The exceptions that exist illustrate some of the benefits of such reminding: The well-known explanations of Amazon.com (e.g., “We recommend product *A* because you once bought product *B*”) do not only explain the recommendation; they also give the user an opportunity to reflect about the significance of their past experiences (e.g., “I didn’t actually like product *B*, so I’ll try something different this time”). A more specific type of link to past experiences is offered by the *reflective history* mechanism of Zimmermann et al. ([20, 21]): When recommending a new TV show, their system would produce an explanation like “*Xena: Warrior Princess* is produced by Sam Raimi, who produced the TV show *American Gothic*”, the latter TV show being one that the user has seen in the past.

As is already suggested by these somewhat similar mechanisms, by shifting the emphasis from pure recommendation to a combination of recommending and reminding, the recommendation paradigm involves the user more strongly in the process of making the final decision.

Another important benefit of realizing recommendation in a mobile personal assistant is that data about past experiences from a wider range of sources can be taken into account than with existing recommenders: As we will see, SPECTER considers products that the user has encountered while browsing the web, products that she has seen in stores, and (depending on the technical possibilities for capturing experiences) in principle products that the user has encountered in just about any situation.

2.2 Augmented Memories More Generally

The automated building of memory-like structures has been studied for more than 10 years (see, e.g., [13], [6], [9], [11]). Most of this work has focused on capturing and organizing data with the goals of freeing the user from the need to record perceptions manually (e.g., by making notes) and of extending his or her perception (e.g., by capturing events not noticed by the user).

The retrieval of information from such memories has been realized in quite different ways. For instance, to support a user while he or she is browsing in an augmented memory, the system may visually align and cluster the retrieved information (see, e.g., [8]). For explicit retrieval, Lamming et al. emphasize the special value of context for retrieval ([13]), while van den Hoven explores an object-centered approach to recollection ([18]). We contribute an approach that integrates these ideas with results from a study about how users apply such records in an everyday setting—a question seldom addressed before.

Besides such on-demand retrieval, augmented memories can also play a role in proactive user support. For instance, a system may display memories related to some previously retrieved experience in order to assist the user in discovering relationships (see, e.g., [14]). A variation of this approach based on a task

model is discussed in [10]: A knowledge worker is supported via the presentation of task-relevant documents that have been retrieved automatically from a corporate memory. Other example scenarios include the prevention of nursing accidents by means of context-sensitive reminders of instructions ([12]) and social matching on the basis of shared experiences (see, e.g., [17]).

3 Example Domain and System Overview

In order to investigate the application of augmented memories, we defined a scenario with different stages, in each of which the user’s past experiences play a different role. We sketch the scenario and the study only briefly here; more details will be found below and in [16].

The general background story was that the study participant was given a gift certificate worth 30 Euros for shopping at two CD stores (called “Bonnie’s” and “Clyde’s”) that were specialized in a particular genre of CDs: movie soundtracks.

1. In Phase 1, before coming to the laboratory, each participant filled in a web form indicating, for each of 100 movies, what previous experience they had with that movie and/or its soundtrack.

2. In Phase 2, in the laboratory, the participant spent 20 minutes browsing the “Soundtrack CDs” section of Amazon.com’s website, the goal being to prepare for the later shopping phase by becoming familiar with soundtracks and artists. During this phase, the SPECTER system recorded the participant’s actions, as well as guesses about the participant’s evaluation of the CDs viewed.

3. In Phase 3, the participant reviewed SPECTER’s record of her browsing behavior and had an opportunity to correct any inaccurate guesses that SPECTER had made about her evaluations of CDs.

4. In Phase 4, the participant spent 30 minutes looking in Bonnie’s and Clyde’s for interesting CDs that she would like to buy with her gift certificate. During this phase, SPECTER continued to record the user’s actions (e.g., looking at the covers of individual CDs), but the system exploited the episodic memory built up so far with various recommendation functions.

This setting formed the test bed for a framework of components for capturing and processing perceptions made by sensors in a smart environment (see Figure 1). Perceptions obtained from sensors are buffered in a “short-term memory”.² In addition, perceptions are recorded for long-term use in a context log without further processing.

Perceptions become the subject of an abstraction process that summarizes and combines perceptions to yield a more abstract representation. In the case of the short-term memory, the system matches the result of this process to preconditions of services specified in the user model. One of these services generates entries in a *personal journal*, which provides the user with an event-based view on the collected data.

The personal journal is tightly linked to a user interface that offers ways of inspecting and editing the artificial episodic

²Although no simulation of human cognitive processes is involved, our framework includes some terms borrowed from psychology, which suggest loose analogies.

¹A useful high-level overview is given in the first part of [5].

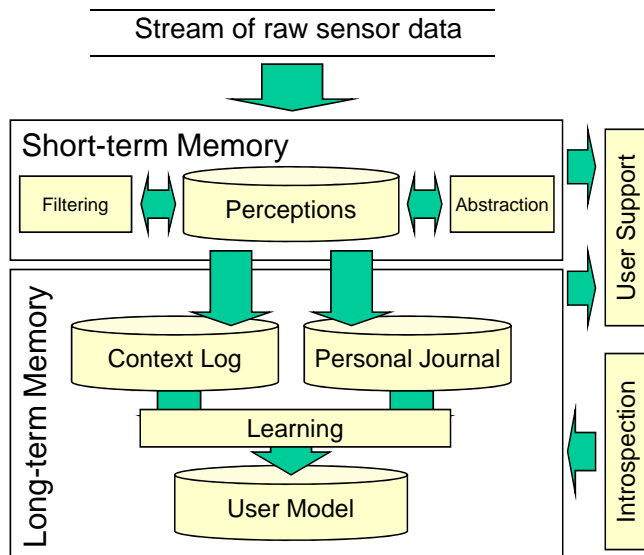


Figure 1: Main components and processes of the augmented personal memory realized in SPECTER.

memory. User feedback gathered during these processes is exploited—in combination with data extracted from the log and the journal—to build the user model. This model then affects the abstraction process, event construction, and situated behavior of SPECTER.

In the following, we will walk through selected aspects of this framework.

4 Populating the Augmented Episodic Memory

4.1 Capturing the User's Experiences

The availability of relevant observations about a user's actions and context is a crucial factor for the success of a system based on an artificial episodic memory. With SPECTER, we use multiple approaches to acquire such information.

An important part of the CD shopping scenario introduced above is the observation of the user while he or she is browsing for music on the Amazon.com web pages (Phase 2 of the study). We use a two-fold approach to capture and semantically represent such interactions: In the first step, we analyze the web pages visited by the user with a web proxy (cf. [19]) and try to extract the Amazon standard identification number (ASIN) of the products viewed by the user. In the second step, we use this number as a key for querying the Amazon product database (cf. [1]) for additional information. From the data returned, we automatically create a new product instance in our object ontology and link this instance to a semantic representation of the browsing act.

Another important aspect of the problem is that of capturing the user's behavior in the real world. To observe the user's actions in our mocked-up stores (Phase 4 of the study) we use radio frequency identification (RFID) techniques. Each CD in the stores is tagged with a passive RFID transponder and given a

unique identification code. Antennas in the shelves record which CDs are taken from the shelves by the user at which times, as well as the length of the time that each CD is held by the user. Additionally, each transponder is annotated with the Amazon identification code of the CD in question. Accordingly, a product instance is automatically created in our object ontology for each CD in the mocked-up stores. This instance enables the system to relate the user's experiences in the real-world stores to his or her experiences on the Amazon.com web pages and vice versa.

To determine the user's position within and between the mocked-up stores, we employ an indoor location tracking system based on active RFID transponders and georeferenced dynamic Bayesian networks (cf. [4]). The system requires a number of transponders to be placed as beacons in the environment. A mobile receiver carried by the user receives the signals from these beacons and computes the user's current position. This information is also used to detect when the user enters or leaves a store. Because the spatial resolution of this system is too low in some cases, we implemented an additional simulation interface that allows the experimenter to generate arbitrary sensor events, especially location-tracking events.

Besides the user's actions, we try to capture the user's general context by exploiting other, external resources. In another scenario loosely related to the CD scenario, for instance, we used the user's current position to query a weather information web service for the current weather conditions and temperature at the user's location. This information was stored together with the observed user actions and used as a cue for retrieval from the digital episodic memory.

4.2 Organizing Perceptions

The result of this approach to capturing context is an extensive record of largely uninterpreted data that can be exploited with the goal of extending the user's perception. But the sheer amount of information, the lack of any relevance measure, and the fine-grained nature of this data will in general make its unaided processing by the user him- or herself infeasible. Therefore, the system has to assist the user not only in capturing but also in interpreting and organizing information about actions and events.

A way to address this challenge is to employ methods that automatically transform perceptions into more abstract structures. In SPECTER, we experimented with methods such as machine learning for the derivation of motion profiles (e.g., to distinguish between ambling and goal-directed walking, cf. [3]) and plan recognition for recognizing complex events made up of several perceptions. These processes do not involve the deletion of perceptions but rather the enrichment of the augmented memory. This approach is driven by several assumptions:

- Long-term effects: Data might turn out to be of value for long-term observations.
- Explanation: The unprocessed data can be exploited for explanations of the system's top-level behavior.
- Limited comprehension: Because of its incomplete user model, the system cannot compete with the user's comprehension of a situation.

Consequently, the original perceptions remain available in the memory. While this principle was less relevant in the CD shopping scenario, it turned out to be of considerable interest

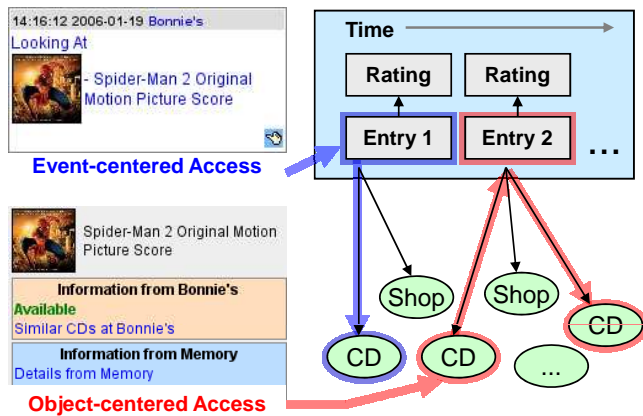


Figure 2: Event- and object-centered access to memories: on the left-hand side, examples from the user interface; on the right-hand side, a sketch of the corresponding data model.

for other settings—for example, in a grocery shopping scenario, where precise records of temperature data can be exploited to explain a warning sent to the user about spoiled food.³ Nevertheless, since this approach does increase the amount of data held in the memory, it complicates the retrieval of experiences.

In search of a data model for the support of retrieval from augmented memories, we had to consider that different user goals might call for different retrieval strategies. For instance, many situations in our scenario focus on the user's interaction with the CD that the user is currently attending to; for these situations, object-centric access is most natural. By contrast, other situations require knowledge about the context of the user's past experiences, for example, to provide the user with examples of CDs seen in similar situations.

We decided to take this duality into account in the data model. Its basis is an ontology describing actions, objects and locations (cf. [15]). The fact that resources such as CDs and locations are bound to the ontology enables the exploration of memories on the basis of the ontology's graph structure. This mechanism is exploited by the system for detecting similarities and by the user interface for providing an object-centered view on memories. This view enables the realization of a user interface that allows the comparison of object features (e.g., price) as well as supporting object-specific functions (e.g., similar CDs in the store). In parallel, the personal journal provides on the data level an event-based index of the same data. In this way, the users are able to browse experiences related to some location or object (e.g., positive and negative experiences bound to a particular CD). The corresponding user interface connects both data views; for instance, the display of an object offers access to related events, whereas the display of an event offers access to the objects involved (see Figure 2).

³This scenario is currently being investigated in the project SharedLife, which builds on the results of SPECTER.

5 Reviewing and Correcting the Augmented Episodic Memory

No matter how sophisticated the capture of information about a user's past experiences may be, the record acquired in this way is likely to contain many inaccuracies, especially where some amount of interpretation on the part of the system is involved. It will not in general be feasible or necessary to correct all of the inaccuracies; but as we will see, in some cases it makes sense to offer the user an opportunity to review the personal journal and to make any adjustments that seem worthwhile. The reasons can be illustrated with reference to the user study in the CD scenario: Before actually shopping in the CD stores, the users spent 20 minutes exploring the "Soundtrack CDs" section of the Amazon.com website to familiarize themselves with some of the CDs and artists that they would be likely to encounter in the stores. During this browsing phase, SPECTER not only kept track of the CDs that the user had looked at (and in some cases listened to briefly) but also made guesses about the user's likely evaluation of each CD encountered. (These guesses were based simply on the amount of time that the user spent on the web page describing the CD.) These guesses were the part of the personal journal that was likely to be inaccurate much of the time. At the end of this browsing phase, each user was given an opportunity to engage in what we call *reflection*: She could look back at her encounters with the CDs on the website, accessing them via various categories (e.g., one including the CDs that SPECTER had guessed that she had evaluated positively). She could then replace the guesses of SPECTER with her actual evaluations. (The record in the personal journal distinguished between guesses made by SPECTER and explicit evaluations recorded by the user, because the latter can be assumed to be more reliable.)

Although correcting estimates of evaluations might be thought to be a tedious process, the participants in the study recognized two benefits that made the process seem worthwhile: 1. Given a reasonably accurate record of the user's evaluations, SPECTER can give more useful reminders and recommendations while the user is shopping in the stores. 2. The process itself of reviewing the experiences—not chronologically but according to categories relevant to the upcoming shopping experience—was found by some to be a useful way of preparing for the shopping activity. Despite these perceived benefits, the participants thought that this type of reflection was the kind of activity that they would prefer to handle in real life during idle time, such as while taking a bus into town in order to shop.

One general principle illustrated here is that each activity that a user is expected to engage in to ensure a reasonable level of performance by a personal assistant like SPECTER needs to seem reasonably attractive in terms of the balance of costs and benefits. Making it possible to engage in reflection during idle time is a way of minimizing the costs; organizing reflection in a way that maximizes the benefit for an immediately following shopping experience is a way of maximizing both the benefits and their immediacy.

6 Exploiting the Augmented Episodic Memory

The functions of SPECTER discussed in the previous sections were designed to ensure that the user has a reasonably rich and accurate augmented episodic memory. We now turn to the key question of what can be done with such a memory. Since space is too limited for a complete discussion of SPECTER's functionality and user interface, we will mention in turn several typical functions and how they were used and accepted by the participants in Phase 4 of the CD study, in which the participants shopped in Bonnie's and Clyde's.

1. Use of Previously Encountered CDs as a Starting Point. Two types of information that SPECTER offered the user when the user entered one of the mocked-up CD stores were the answers to the following questions: (a) "What CDs that I have seen in the past and evaluated positively are available in this store?"; (b) "What CDs in the store are *similar* to ones that I have liked in the past?" The user could also pose explicit queries of these types at any time. This functionality was frequently used and highly rated; it allowed the participants to pick out those CDs among the hundreds in each store that were in some way related to their previous experiences.

We can understand the added value of this functionality by considering the situation of a shopper who enters a store and tries to recall, using her own memory, which relevant products she has seen and liked in the past: Her recall would probably be far from complete, unless the set of products in question comprised only a small number of distinctive items—or she skillfully used information in the store (e.g., signs indicating categories of CDs) as retrieval cues. But even if her recall were perfect, she would then have to proceed to check which of the recalled items were available in the store—in general discovering that some of the items that she had taken the trouble to recall were not available anyway. In other words, SPECTER further increases the value of its support for recall by using the results immediately as input to a second type of processing for which the shopper's cognitive capabilities are even less well suited.

2. Use of Information About Known Prices. Maybe the most obvious practical benefit of reminders of past experiences concerns price comparison: It was natural for participants in the study to want to acquire each CD at the lowest possible price, and the prices sometimes differed between the stores. Accordingly, before deciding to buy a CD, participants regularly checked whether they had already seen it cheaper in the other store. Although price comparison systems on the web perform a roughly similar function, this function of SPECTER provides support that is specifically applicable to the physical stores that the user is visiting (which would not in general be covered in an up-to-date way by any web-based price comparison system, even if the user had access to such a system).

The added value of SPECTER's automatic recording of prices can be understood in terms of the fact that human episodic memory is not very good for the storage and retrieval of detailed numerical information. (By contrast, our shoppers did not especially appreciate SPECTER's ability to remind them of how much they had previously liked a given CD when they encountered that CD again: Evidently their own memory for their past evaluations was seen as largely adequate.) As a further benefit, SPECTER could in principle have provided some

automatic processing of the available price information (e.g., alerting the shopper to cases in which a CD is cheaper in the current store than in the other store). But this functionality did not appear to justify the risk of distracting the shoppers with irrelevant information: The shoppers tended to use price as a secondary consideration, their primary goal being to identify CDs that they would like to have; and they found the active inspection of SPECTER's price reminders to be a straightforward and natural task.

3. Use of Pointers to Previously Encountered Similar CDs. This function is the one that was illustrated in our example in the Introduction: SPECTER presents previously encountered products that are similar to the one that the user is currently examining. Whereas some participants in the study accessed this function fairly frequently and rated it as helpful, about an equal number of participants used it little, if at all. Apparently the benefits of this sort of reminder are not immediately obvious to everyone, perhaps because it is seldom encountered in familiar systems.

Note that in this case as well, the added value of SPECTER's digital episodic memory lies largely in the fact that it is applied in conjunction with a different type of processing that would not in general be possible for the shopper herself: Since the item that the shopper is currently viewing is not known to the shopper, it would be hard for her to recall similar items that she already knows, since she would not be able to assess the similarity of the item to previously encountered items.

4. Use of Representations of Past Events. One function offered by SPECTER that was seldom used or appreciated in our study was the ability to access records of specific previous experiences (e.g., the experience of looking at a given CD in Amazon.com)—to be distinguished from the more abstract reminders of specific products and their evaluations (without reference to any specific event). One possible explanation is that a record of a shopping event such as looking at a product does not contain much useful information beyond the identity of the product and the user's evaluation of it. Events such as having seen a movie a long time ago on a special occasion would presumably be more worth remembering; but the events in question could be recalled in our study by the participants without help from SPECTER (since they had entered this information themselves just a day earlier). Therefore, more interest in specific past experiences might be found in systems that recorded information about more remote and/or interesting events. Still, this result illustrates the general point that systems that offer augmented episodic memories need to present the stored information in a form and at a level of abstraction that makes sense in view of the user's tasks and the capabilities of their own human memories.

7 Determining the Timing of the System's Proactive Actions

The recommendation paradigm requires that the system have access to knowledge about when to present memories to the user. The augmented memory is again of interest for retrieving such knowledge. If the user decides to map some system service (such as recommendation) to some situation, the system needs to know the characteristic features of that situation in order to construct

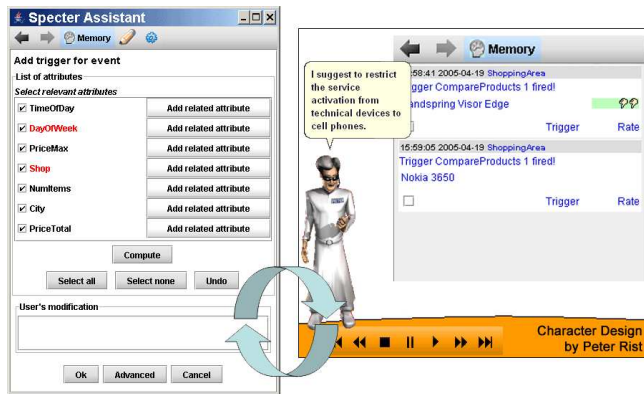


Figure 3: A user interface for critiquing situational features (left-hand side); user feedback may trigger a negotiation about these features during reflection (right-hand side).

a decision tree which is applied to recognize the situation's future occurrences. We realized the specification of these features as a collaborative process comprising several steps, which combines 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:

1. The user selects events from the episodic memory, which represent, in his or her opinion, typical preconditions for triggering the desired service.
2. The system expands this example set on the basis of a similarity measure with additional events from the memory.
3. The system mines the example set for candidate features.
4. The user critiques the feature set.
5. Once the user is satisfied with the feature set, the decision tree is constructed and linked to the system's short-term memory.

The result of our ongoing effort is a user interface which provides several interaction layers of varying complexity for combining services and situations; for a detailed description of this interface, please see [2]. Especially relevant for critiquing a situation's features (and thus for configuring the execution of linked services) is the screen shown on the left-hand side of Figure 3. This screen shows, as presented to the user, a list of features that have been extracted using statistical methods.

The user can 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 be provided only in certain kinds of stores. She might do so by inspecting and refining the store-related features—for example, by replacing the name of a chain of stores with the designation of a particular store in a given location.

Once such a decision tree has been linked to the short-term memory, its effects are recorded in the augmented episodic memory. SPECTER can then take the user's feedback on its supporting actions into account for the reflection process. For instance, if the user rates one of these system actions negatively, the system will ask during the next introspection process whether the decision tree in question should be adjusted. The system makes

suggestions that are based on the ontology's structure—for example, suggesting that the service be triggered only for a more specific product category, as is illustrated on the right-hand side of Figure 3. If the user is not confident with these suggestions, the features have to be adjusted by means of the critiquing mechanism described above.

8 Conclusion

Compared with many articles on the relationship between AI research and human cognition, our brief discussion of SPECTER and the recommendation paradigm has illustrated a different way in which the application of AI (and other) techniques can take into account general properties of human cognition: Instead of trying to simulate human cognition in order to understand it better, we try to support human cognition by providing functionality similar to certain cognitive functions in a way that augments these functions. Although this general strategy has already been applied in many familiar ways, SPECTER's digital episodic memory illustrates how new applications of it can still be found—especially when hardware and software advances make feasible new types of support.

The experience with SPECTER shows that the added value of a digital episodic memory is especially great when (a) the particular recall and recognition tasks involved are ones for which human episodic memory is not very well adapted; and/or (b) the results yielded by the digital memory can be processed further by the system without intervention of the user.

9 Acknowledgements

This research is supported by the German Ministry of Education and Research under grants 524-40001-01 IW C03 (project SPECTER) and 01 IW F03 (project SharedLife). We would like to thank all past and present project members for their invaluable support.

References

- [1] Amazon Inc. Amazon e-commerce service, February 2007. <http://aws.amazon.com/>.
- [2] Mathias Bauer and Stephan Baldes. An Ontology-Based Interface for Machine Learning. In John Riedl, Anthony Jameson, Daniel Billsus, and Tessa Lau, editors, *IUI 2005: International Conference on Intelligent User Interfaces*, pages 314–316, New York, 2005. ACM.
- [3] Mathias Bauer and Matthieu Deru. Motion-based adaptation of information services for mobile users. In Liliana Ardissono, Paul Brna, and Tanja Mitrovic, editors, *Proceedings of the 10th International Conference on User Modeling (UM05)*. Springer, Edinburgh, Scotland, 2005.
- [4] Boris Brandherm and Tim Schwartz. Geo referenced dynamic Bayesian networks for user positioning on mobile systems. In Thomas Strang and Claudia 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.
- [5] Robin Burke. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4):331–370, 2002.

- [6] Brian Clarkson, Kenji Mase, and Alex Pentland. The familiar: a living diary and companion. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 271–272, New York, NY, USA, 2001. ACM Press.
- [7] Mary Czerwinski, Douglas W. Gage, Jim Gemmell, Catherine C. Marshall, Manuel A. Pérez-Quinones, Meredith M. Skeels, and Tiziana Catarci. Digital memories in an era of ubiquitous computing and abundant storage. *Communications of the ACM*, 49(1):44–50, 2006.
- [8] Susan Dumais, Edward Cutrell, JJ Cadiz, Gavin Jancke, Raman Sarin, and Daniel C. Robbins. Stuff I've Seen: A system for personal information retrieval and re-use. In *SIGIR '03: Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 72–79, New York, NY, USA, 2003. ACM Press.
- [9] Jim Gemmell, Lyndsay Williams, Ken Wood, Roger Lueder, and Gordon 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.
- [10] Harald Holz, Heiko Maus, Ansgar Bernardi, and Oleg Rostanin. A lightweight approach for proactive, task-specific information delivery. In *Proceedings of I-KNOW 2005 - Special Track on Business Process Oriented Knowledge Infrastructures BPOKI*, 2005.
- [11] Johan Koolwaaij, Anthony Tarlano, Marko Luther, Petteri Nurmi, Bernd Mrohs, Agathe Battestini, and Raju Vaidya. Context Watcher - Sharing context information in everyday life. In *Proceedings of IASTED conference on Web Technologies, Applications and Services (WTAS)*, Calgary, Canada, July 2006.
- [12] Noriaki Kuwahara, Haruo Noma, Kiyoshi Kogure, Norihiro Hagita, Nobuji Tetsutani, and Hiroshi 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.
- [13] Mik Lamming and Mike 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.
- [14] Brian Lee, Scott R. Klemmer, and Ronen Brafman. Adaptive interfaces for declarative presentation of heterogeneous content. Technical Report CSTR 2006-10 09/29/06, Stanford University, HCI Group, Stanford, CA, USA, October 2006.
- [15] Ian Niles and Adam 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.
- [16] Carolin Plate, Nathalie Basselin, Alexander Kröner, Michael Schneider, Stephan Baldes, Vania Dimitrova, and Anthony Jameson. Recommendation: New functions for augmented memories. In Vincent Wade, Helen Ashman, and Barry Smyth, editors, *Adaptive Hypermedia and Adaptive Web-Based Systems: Proceedings of AH 2006*, pages 141–150. Springer, Berlin, 2006.
- [17] Yasuyuki Sumi and Kenji Mase. Supporting awareness of shared interests and experiences in community. *International Journal of Human-Computer Studies*, 14(6):127–146, January 2002.
- [18] Elise van den Hoven. Exploring graspable cues for everyday recollecting. In *Proceedings of the Workshop on "Memory and Sharing of Experience" at Pervasive 2004*, pages 21–24, Vienna, 2004.
- [19] Harald Weinreich, Volkert Buchmann, and Winfried Lamersdorf. Scene: Ein framework zur evaluativen realisierung von erweiterungen des webs. In K. Irmischer and K.-P. Fähnrich, editors, *Tagungsband Kommunikation in Verteilten Systemen - KiVS 2003*, pages 31–42. Springer-Verlag Heidelberg, 2 2003.
- [20] John Zimmerman and Kaushal Kurapati. Exposing profiles to build trust in a recommender. In *Extended Abstracts for CHI'02*, pages 608–609, Minneapolis, MN, 2002.
- [21] John Zimmerman, Kaushal Kurapati, Anna L. Buczak, Dave Schaffer, Srinivas Gutta, and Jacquelyn Martino. TV personalization system: Design of a TV show recommender engine and interface. In Liliana Ardisson, Alfred Kobsa, and Mark Maybury, editors, *Personalized Digital Television: Targeting Programs to Individual Viewers*. Springer, Berlin, 2004.

Contact

German Research Center for Artificial Intelligence
Stuhlsatzenhausweg 3
66123 Saarbrücken

Email: firstname.lastname@dfki.de

The authors are members of the German Research Center for Artificial Intelligence; the article addresses their joint work in the projects SPECTER and SharedLife.

Bild

Alexander Kröner holds a Ph.D in Computer Science from Saarland University. He directed the project SPECTER; his field of experience comprises the application of constraints, web technology, and ontologies for adaptive user support, with a particular focus on personalized and situated information presentation.

Bild

Anthony Jameson holds a Habilitation degree in Cognitive Science from Saarland University. His research mostly addresses user-adaptive systems, language technology systems, systems based on semantic technologies, and other types of intelligent user interface from a user-centered perspective.

Bild

Michael Schneider holds a Diploma in Computer Science from Saarland University. He is interested in different areas of artificial intelligence with a focus on ubiquitous computing, proactive assistance, and plan recognition.

Bild

Nathalie Basselin holds a master's degree in Human-Computer Interaction from the University of Toulouse, France. She is now a PhD student at Saarland University. Her research focuses on the user-centered design and evaluation of ubiquitous and community-based applications.