Recomindation: New Functions for Augmented Memories

Carolin Plate, Nathalie Basselin, Alexander Kröner, Michael Schneider, Stephan Baldes, Vania Dimitrova, and Anthony Jameson*

DFKI, German Research Center for Artificial Intelligence
Saarbrücken, Germany

Abstract. Advances in technological support for augmented personal memories make possible new ways of enhancing the process of product recommendation. Instead of simply analyzing information about a user's past behavior in order to generate recommendations, a *recominder*[1] system can additionally supply various types of information from the user's augmented memory that allows the user to take a more active role in the search for suitable products. We illustrate the paradigm of *recomindation* with reference to a prototype implementation of the system SPECTER in a CD shopping scenario and the results of a study with 20 subjects, who found most of the recomindation functionality to constitute a useful enhancement of their shopping experience.

1 An Introduction to Recomindation

One development of the past few years that promises to bring substantial innovations in the area of adaptive hypermedia (among others) concerns augmented personal memories: It is becoming feasible for a system, with the user's consent, to store a vast amount of information about the user's actions, experiences, and contexts over a long period of time (see, e.g., [1] for a pioneering effort; [2] and [3] for more recent influential projects; and [4] for a collection of recent papers). As is discussed in a recent survey by Czerwinski et al. ([5]), these technological possibilities raise the questions "Why bother?" and "What might I do with all the stuff I collect?" (p. 46). These authors list several possible answers, ranging from helping users to find lost objects to improving their time management.

The aim of this paper is to advance a new answer to this question: Augmented memories can add a new dimension to the process of product recommendation. On the one hand, most approaches to recommendation (which are surveyed, e.g., in several chapters of the edited volume [6]) rely heavily on stored information about the past behavior of the user: products purchased, ratings made, web sites visited. On the other hand, this access to information about the past occurs largely outside of the user's awareness.

* This research was funded by the German Federal Ministry of Education and Research (BMBF) under contract 524-40001-01 IW C03. Vania Dimitrova participated during a sabbatical leave from the University of Leeds. The character shown in Figure 2 was designed by treschique digital arts.

[1] The third syllable is pronounced as in *remind*.

To be sure, the recent trend toward the explanation of recommendations sometimes (though by no means always; see, e.g., [7]) involves exposing the user increasingly to information about his or her past. The most familiar examples are the recommendations of amazon.com of the form "We recommend product A because you once bought product B". A different type of explanation along these lines was introduced in the *reflective history* mechanism of Zimmermann et al. ([8, 9]): When recommending a new TV show, the 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.

We will argue that explicit reminders of the user's past experiences and behavior can enhance the recommendation process in various ways; and that, with the technical feasibility of extensive augmented memories, the time is ripe to explore these possibilities. We introduce the new term *recomindation* to capture the blend of recommendation and reminding that we propose.

We will introduce our vision concretely and concisely by describing the proof-of-concept system SPECTER that we have developed, along with a user study that reveals how people use and evaluate the new functions that SPECTER offers.[2]

2 Specter and the Method of the User Study

The basic scenario is as follows: You have been using the personal assistant SPECTER for a long time, and the system has collected information about things like CDs that you have bought and movies that you have seen. You have just received a gift certificate that entitles you to buy CDs at two music stores, called *Bonnie's* and *Clyde's*, respectively. These stores specialize in CDs of movie soundtracks. Since you are not very familiar with these genres, you decide to spend about 20 minutes browsing in the "Soundtrack CD" section of amazon.com's web site before going to the stores. SPECTER will keep a record of your browsing behavior and, once you have finished browsing, will allow you to edit this record. When you actually visit the two stores, you will be able to access SPECTER's record in various ways, even as the record is being extended with SPECTER's further observations of your behavior in the two stores.

The subjects who performed the tasks of this scenario in our user study were 20 persons between the ages of 18 and 32, 11 male and 9 female. All subjects had considerable computer experience, and 8 subjects had significant experience with PDAs.

2.1 Pre-Shopping Phases

Entering of Information About Earlier Experiences. For the study, the only feasible way of finding out about relevant biographical events such as the watching of movies was to ask the subjects to enter this information via a web form before coming to participate in the experiment. Presented with a list of 100 popular movies, the subject

[2] An earlier presentation of SPECTER that focuses largely on other functionalities is given by [10]. An earlier version of SPECTER was tested in a similar user study with 30 subjects, whose results led to improvements in the system's design and in the method of the study described in this paper.

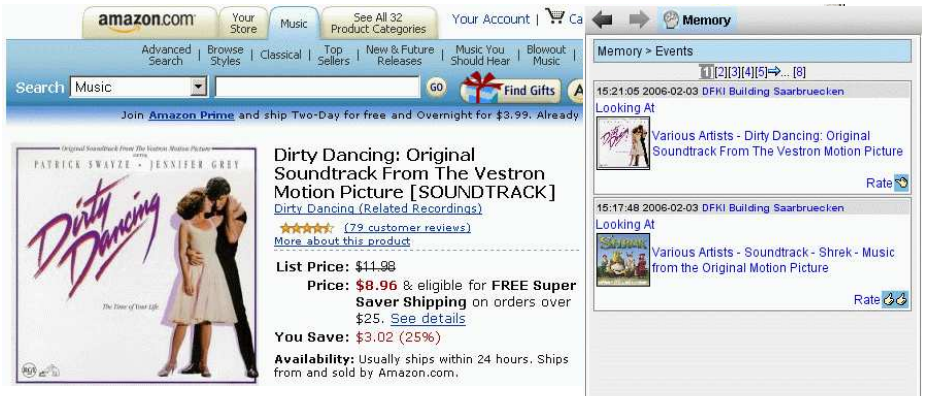

Fig. 1. Snapshot from the phase of browsing in amazon.com, with SPECTER's record of observed events shown on the right. (Ratings suggested by SPECTER have a darker background color than those made by the user herself, though this difference is barely noticeable on a monochrome printout.)

was asked to rate the movies that she[3] had seen and state when and under what circumstances she had seen them (e.g., "with friends/family"); she also answered similar questions about the soundtrack CD for the movie.

The remaining phases of the study took place at our laboratory. Before each phase, the subject was given any instructions, advice, and practice (lasting up to 30 minutes) that were necessary to enable her to operate the interfaces in question without spending much time on trial and error.

Browsing in amazon.com. As is illustrated in Figure 1, each subject spent 20 minutes browsing on a PC in the special section of amazon.com's web site for soundtrack CDs. In addition to reading texts and reviews, the subject could listen to excerpts of individual songs. In a window next to the web browser, an instantiation of SPECTER was running. Each time the subject looked at a CD's page, a record of this action was added to SPECTER's reverse-chronologically ordered list of descriptions of actions (performed by the user or by SPECTER itself). For each CD viewed, SPECTER entered an estimated rating on a 5-point scale: The highest rating of "two thumbs up" was entered if the subject spent more than 1 minute looking at the page, while "two thumbs down" was entered if they spent less time. The subject could immediately change any estimated rating; but she could also postpone such adjustments to one of the next two phases (as most subjects in fact did).

Reflecting on the Browsing Phase. In the *reflection* phase, the subject was encouraged to explore for up to 20 minutes the records that SPECTER had built up concerning the CDs that the user had encountered so far. (Most subjects did so for 5–10 minutes.) As can be seen in the right-hand side of Figure 2, each of these records refers to a par-

[3] Since gender-neutral language is often imprecise and/or cumbersome, we arbitrarily use a feminine pronoun for each generic reference to a subject in the study.

Fig. 2. Screen in the reflection phase with which the user could select a subset of previously seen CDs to review. (In this case, she has chosen those rated "one thumb up"..)

ticular CD—in contrast to the records shown in Figure 1, which describe actions and events). An animated character representing SPECTER encouraged the subject actively to examine various subsets of these CDs, such as the ones corresponding to the movies that the user had seen or the ones for which SPECTER had suggested a rating during the browsing phase. Here again, the subject could choose whether to change the rating estimates of SPECTER that did not reflect her actual evaluation (which tended to be numerous, because of SPECTER's crude estimation method).

2.2 Shopping in the Stores

Setup and Instructions. The final and most complex phase involved shopping for a total of 30 minutes in *Bonnie's* and *Clyde's*. Each store was mocked up with a CD rack containing 250 CD cases that looked like those found in real stores, arranged in alphabetical order by title. The two stores were mocked up on two sides of the same large room; but so that subjects could not move back and forth between them at an unrealistic rate, they were required to walk along a circuitous 50-meter trajectory to get from one store to the other.

A subset of 50 CDs was present in both stores; in each such case, the prices in the two stores differed noticeably, making it worthwhile for subjects to compare prices.

The subject was told that she was to pick out a total of 6 CDs that she would like to own; at the end of the study, two of these CDs would be picked at random by the experimenter and awarded to the subject to keep. The difference between the total price of these 2 CDs and 35 euros was paid to the subject in cash.

In this phase, the SPECTER application was operated on a PDA that was linked via Bluetooth and VNC to the central SPECTER server. Some of the functions to be described below made use of web services from amazon.com that provided information about music CDs: most importantly, a service that listed CDs that were "similar" to a given CD (or list of CDs) in the sense of amazon.com's item-to-item collaborative

Table 1. Overview of functions available to the user in the hand-held SPECTER system used in stores.

In any situation:
- List all events that have occurred so far
- List all previously encountered CDs which ...
  - ... were encountered in a given location
  - ... have a price lower than a given amount
  - ... have a rating (by the user, by Specter, or by either one) above a particular threshold

Given a list of CDs (cf. Figure 3 B):
- Filter the list to keep only those which ...
  - ... have been encountered by the user
  - ... are available in the current store
- List similar CDs which ...
  - ... have been encountered by the user
  - ... are available in the current store

Given a particular CD (cf. Figure 4 A)
- Suggest related CDs:
  - List similar CDs in the current store
- Give further information about the CD:
  - Show the details (e.g., artists) of the CD
  - List the prices of the CD at all places at which it has been encountered
  - Show the rating of the CD (and allow the user to change it)
  - List similar CDs that the user has encountered
  - List all events involving the CD

Given a list of events:
- Filter the list to keep only those which ...
  - ... occurred in a given location
  - ... involved a particular type of action

Fig. 3. Sequence of screens illustrating how CDs previously rated positively can be used as a starting point for exploration in a store.

filtering (see, e.g., [11]): Roughly speaking, CD A is similar to CD B if customers who buy A also tend to buy B.

Each CD case contained an RFID tag that identified the CD uniquely. When a subject removed a CD from the rack to look at the information on its cover, this action was detected by an RFID antenna that transmitted information about the action to SPECTER, allowing SPECTER to display the CD on the handheld's screen along with several options for obtaining additional information about it (see, Figure 4 A).

Functions Offered by SPECTER. Table 1 gives an overview of the functions that were made available to the user during the shopping phase. The design philosophy involved including: (a) promising novel functions that make use of SPECTER's augmented memory; and (b) a few additional functions that can make effective use of the results of the memory-related functions.

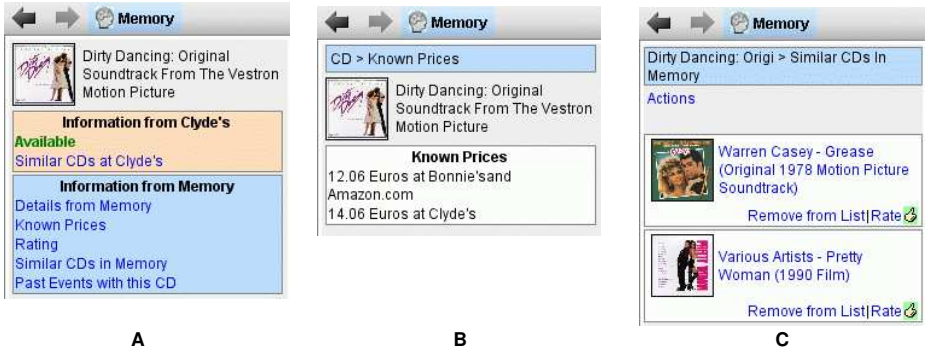

Fig. 4. Screens illustrating information from the augmented memory that the subject can request for a given CD.

Figure 3 A illustrates how a user can acquire a starting point for exploring the current store by requesting a list of all CDs that she has previously rated very positively. As Figure 3 B shows, she could filter this list to include only CDs that are available in the current store (Clyde's); but in this example, she instead requests a list of *similar* CDs that are available in Clyde's. This list will probably include a number of CDs that are unfamiliar to the subject. For each such CD, the user can either (a) tap on its entry in the list or (b) retrieve the physical CD in the store and examine its contents. In either case, a screen like the one shown in Figure 4 A will appear, offering various types of additional information about the current CD, including: the prices of the CD at all locations where the user has encountered it (Figure 4 B); and a list of similar CDs that the subject has already encountered (Figure 4 C).

The potential utility of this third function may not be obvious. But consider an attentive salesperson who is aware of the customer's past purchases and can therefore introduce a new CD by pointing out that it is similar to one or more particular CDs that the customer already knows.

Another typical function is illustrated in Figure 5: For a given CD, the system lists all events that it has recorded which involve that CD. In Figure 5 A, both events occurred in Clyde's during the shopping phase. In Figure 5 B, the earlier event is the original viewing of the movie, which the user reported via the web interface. One hypothesis underlying the introduction of this function was that reminders of past experiences might bring to mind relevant thoughts, evaluations, and even emotions that the user had experienced in the past.

3 Use and Evaluation of Recomindation Functionality

After this selective presentation of SPECTER's functionality, we now turn to some results about how subjects actually used and evaluated this functionality.

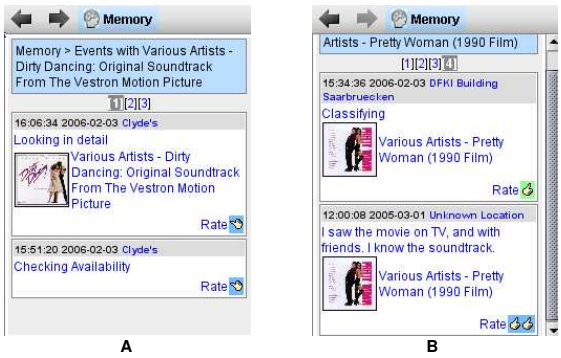

Fig. 5. Screens illustrating information about past events that can be requested for a given CD. (On the right, "Classifying" refers to the user's action of explicitly rating the CD.)

3.1 Responses to the Reflection Functionality

When asked "Judging from your experience, what is the reflection phase useful for?", 10 subjects noted that it was convenient to be reminded, before going into the stores, of the CDs that they had seen while browsing. Apparently, these subjects were not content simply to have SPECTER keep track of their encounters with CDs; they wanted to have the most important CDs in their mind when they entered the stores. The other main benefit of reflection that subjects mentioned concerns the ability to adjust the tentative ratings that were made by SPECTER during the browsing phase. Subjects widely recognized that these ratings were (understandably) often inaccurate; but most of them found these tentative ratings, combined with an opportunity to revise them, to be an effective way of generating a list of promising CDs.

When asked where they would be most likely to engage in this type of reflection in everyday life, most of the subjects expressed a willingness to reflect in a bus or a train (85%), or a waiting room (75%); by contrast, only 35% expressed a willingness to do so in an office or at home in their free time. We can conclude that, despite the generally favorable response to the activity of reflection, the design of a system like SPECTER should ensure that reflection can be performed on a mobile device and in a variety of contexts.

3.2 Responses to Recomindation Functionality Used in Stores

For each of the questions represented in Figure 6, the user was asked "How often did you try to answer the following question?"; they were also asked to state which of three information sources they used (at least sometimes) to answer it: the hand-held SPECTER system, their own memory, or information available in the store. A separate question about each of the corresponding SPECTER functions asked subjects to rate its utility for their search for CDs on a scale from 1 ("not useful at all") to 5 ("very useful"). Figure 6 also shows objective data concerning the subjects' actual frequency of use of these SPECTER functions.

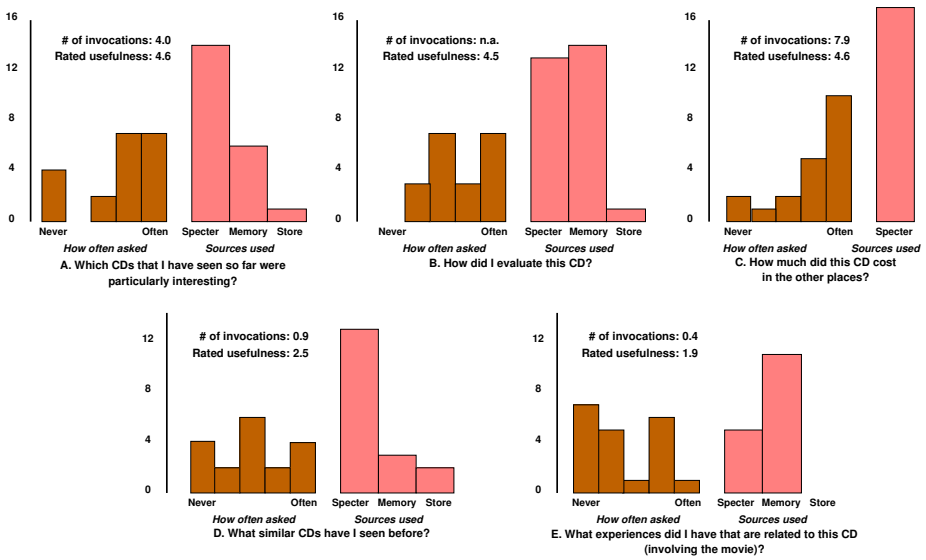

Fig. 6. Frequency distributions of subjects' estimates of how often they asked themselves particular questions during the shopping phase (left-hand side of each graph) and what source(s) of information they used (right-hand side of each graph): the SPECTER system, their own memory, or the CDs in the store). (Also given are the average number of invocations of the corresponding SPECTER function during the shopping phase, as revealed by the system logs; and the subjects' average rating of the usefulness of this SPECTER function for helping them to find appropriate CDs, on a scale from 1 (worst) to 5 (best).)

Use of Previously Encountered CDs as a Starting Point. As Figure 6 A shows, when subjects wanted to remember which CDs they had previously found promising, they were inclined to ask SPECTER; by contrast, remembering their previous evaluation of a given CD appears to be something that they could often accomplish with their own memory (Figure 6 B), though they still found SPECTER's automatic display of their previous ratings useful.

Use of Information About Known Prices. As can be seen in Figure 6 C, subjects were understandably often interested in recalling prices seen earlier, and they invariably used SPECTER, rather than their own memory, for this purpose.[4]

Use of Pointers to Previously Encountered Similar CDs. One of the more novel and less obvious functions offered by SPECTER is the one that was illustrated in Figure 4 C: informing the user about similar CDs that she has encountered before. Figure 6 D confirms that the overall frequency with which subjects considered this ques-

[4] When the subject was in a given store, her SPECTER could not access information about the availability or prices of CDs in the other store unless that information had already been retrieved and stored during a visit by the user to the other store. This restriction corresponds to a real-life situation in which each store offers complete information about its products only to customers who are physically present in the store.

tion was relatively low, though a few subjects did report considering it "often". Note that SPECTER offers a considerable advantage over the subject's own memory here, in that SPECTER can provide an answer even when the CD in question is totally unfamiliar to the user. It is therefore not surprising that subjects reported answering this question much more often with SPECTER than with their own memory.

Use of Representations of Past Events. When we turn to questions about specific events in the past (as shown in Figure 5), there is convergent evidence that this function was less popular than the other functions. With regard to the special case of retrieving information about movies seen at some earlier date, most subjects reported a tendency to consult their own memories rather than SPECTER (Figure 6 E). It is in fact difficult for SPECTER to compete with a user's own memory in this regard, unless it really has been collecting data about the user's movie-going behavior over a period of years, especially since all of the relevant information was supplied by the subject herself via the web interface shortly before the main part of the study.[5] On a different and more generalizable level, several subjects noted that displays such as those shown in Figure 5 are relatively cluttered, relative to the amount of useful information that they provide.

In sum, detailed representations of past events appear to have less compelling utility in the specific scenario investigated than the less concrete information about the products encountered and the ways in which they were evaluated. Any attempt to increase the utility of concrete event representations should devote more attention to ways of representing them more compactly and selectively, perhaps with some degree of aggregation. A more general lesson for research on augmented memories is that it should not be taken for granted without specific evidence that detailed representations of past events will be useful for any particular purpose.

3.3 Global Evaluations of Specter

In any user study of a novel system, global evaluations by users must be interpreted with caution (cf. [12]). Still, it is worth reporting that all of the questions that requested an overall evaluation of SPECTER yielded responses that fell mostly within the top two categories of a 5-point scale. For example, 70% of the subjects indicated that they would be interested in using a system like SPECTER for shopping if it became available.

4 Conclusions

The paradigm of *recomindation* can be seen as a way of exploiting the technological possibilities of augmented memories in such a way as to put the "mind" of the user into the process of recommendation to a considerably greater extent than has been possible so far with recommendation approaches that are based on records from the user's past. Even this selective presentation of our implementation and user feedback shows that a number of aspects of the recomindation paradigm tend to be recognized as appropriate, effective, and even enjoyable—though some limitations and preconditions have been exposed and of course no novel approach is likely to satisfy all users equally.

[5] In a future study, we intend to acquire information about the subjects' past experiences in a way that will not refresh the subject's own memory of these experiences.

Comments of subjects in this and other user studies conducted within the SPECTER project have indicated that the added value of recomindation is likely to be greater in settings (such as those involving shopping for food and cooking) in which the user needs to make more complex decisions (e.g., exactly how to prepare a meal on a particular occasion) and more detailed information from the past can be usefully presented as reminders (e.g., what particular ingredients were used on previous occasions; how the diners evaluated the results). This type of scenario is currently being investigated in SPECTER's successor project SHAREDLIFE, in which the possibilities for sharing memories among friends and acquaintances (cf., e.g., [3, 4]) in the context of recomindation are also being explored.

References

1. Lamming, M., Flynn, M.: "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)
2. Gemmell, J., Bell, G., Lueder, R., Drucker, S., Wong, C.: MyLifeBits: Fulfilling the Memex vision. ACM Multimedia (2002) 235–238
3. Sumi, Y., Sakamoto, R., Nakao, K., Mase, K.: ComicDiary: Representing individual experiences in a comics style. In Borriello, G., Holmquist, L., eds.: UbiComp 2002: Ubiquitous Computing. Springer, Berlin (2002) 16–32
4. Mase, K., Sumi, Y., Fels, S., eds.: Special issue of the journal "Personal and Ubiquitous Computing" on "Memory and Sharing of Experiences". Springer, Berlin (2006) In press.
5. Czerwinski, M., Gage, D., Gemmell, J., Marshall, C., Pérez-Quinones, M., Skeels, M., Catarci, T.: Digital memories in an era of ubiquitous computing and abundant storage. Communications of the ACM **49** (2006) 44–50
6. Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: The Adaptive Web: Methods and Strategies of Web Personalization. Springer, Berlin (2006)
7. Herlocker, J., Konstan, J., Riedl, J.: Explaining collaborative filtering recommendations. In: Proceedings of the 2000 Conference on Computer-Supported Cooperative Work, Philadelphia, PA (2000) 241–250
8. Zimmerman, J., Kurapati, K.: Exposing profiles to build trust in a recommender. In: Extended Abstracts for CHI'02, Minneapolis, MN (2002) 608–609
9. Zimmerman, J., Kurapati, K., Buczak, A., Schaffer, D., Gutta, S., Martino, J.: TV personalization system: Design of a TV show recommender engine and interface. In Ardissono, L., Kobsa, A., Maybury, M., eds.: Personalized Digital Television: Targeting Programs to Individual Viewers. Springer, Berlin (2004)
10. Schneider, M., Bauer, M., Kröner, A.: Building a personal memory for situated user support. In: Proceedings of the First International Workshop on Exploiting Context Histories in Smart Environments (ECHISE 2005) at Pervasive 2005, Munich (2005)
11. Linden, G., Smith, B., York, J.: Amazon.com recommendations: Item-to-item collaborative filtering. IEEE Internet Computing **Jan/Feb** (2003) 76–80
12. Landauer, T.: Behavioral research methods in human-computer interaction. In Helander, M., Landauer, T., Prabhu, P., eds.: Handbook of Human-Computer Interaction. North-Holland, Amsterdam (1997) 203–227