Location-Aware Shopping Assistance: Evaluation of a Decision-Theoretic Approach

Thorsten Bohnenberger[1], Anthony Jameson[2], Antonio Krüger[1], and Andreas Butz[3]

[1] Department of Computer Science, Saarland University
[2] German Research Center for Artificial Intelligence (DFKI)
[3] EyeLed GmbH

Abstract. We have implemented and tested a PDA-based system that gives a shopper directions through a shopping mall on the basis of (a) the types of products that the shopper has expressed an interest in, (b) the shopper's current location, and (c) the purchases that the shopper has made so far. The system uses decision-theoretic planning to compute a policy that optimizes the expected utility of a shopper's walk through the shopping mall, taking into account uncertainty about (a) whether the shopper will actually find a suitable product in a given location and (b) the time required for each purchase. To assess the acceptability of this approach to potential users, on two floors of a building we constructed a mock-up of a shopping mall with 15 stores. Each of 20 subjects in our study shopped twice in the mall, once using our system and once using a paper map. The subjects completed their tasks significantly more effectively using the PDA-based shopping guide, and they showed a clear preference for it. They also yielded numerous specific ideas about the conditions under which the guide will be useful and about ways of increasing its usability.

1 Introduction

Mobile systems offer great potential for providing novel services to users. But this potential can be fulfilled only if the systems are designed in a thoroughly user-centered way. User-related considerations can range from fine-grained problems in dealing with novel input and output modalities to high-level considerations concerning the way in which a system fits into users' customary patterns of activity.

Systems for mobile commerce are a case in point. Especially when equipped with location-awareness, they in theory have a good deal to offer their users. But mobile commerce systems have so far largely been rejected by consumers, even by those who were initially eager to try them out.

Part of the problem lies in the technical limitations of current wireless technology, which lead to long waits and frequent interruptions of connections, in addition to clumsy interaction with small devices (cf. [10]). One remedy is to aim for designs that work well with the current limited technology, checking with users to see whether they really do work well enough.

Another type of problem concerns designs that do not adequately take into account the conditions under which mobile systems are used (e.g., while the user is in motion and/or is simultaneously engaging in some other activity). Only realistic tests conducted early enough in the design process can reliably prevent such problems (see, e.g., [9, 6]).

In this paper we report on an effort to apply a user-centered strategy in the design and development of an intelligent shopping guide based on artificial intelligence techniques of decision-theoretic planning.

The rest of this Introduction describes the basic ideas of the shopping guide, and the next section explains how the current prototype was realized. The remaining sections describe and discuss a study that we conducted with 20 users in a semirealistic setting with a view to assessing and improving the usefulness and acceptability of the shopping guide.

1.1 Combining Product Location With Navigation Support

Location-aware mobile commerce systems that aim to bring together customers and products[4] have so far fallen into two main categories (see [5] for more extensive discussion and examples).

1. Product location services: A user $\mathcal{U}$ queries her system $\mathcal{S}$ about the availability of a particular type of product near her current location (see, e.g., the ADAPTIVE CLASSIFIED SERVER of AdaptiveInfo; http://www.adaptiveinfo.com).
2. Location-dependent alerting services: When $\mathcal{U}$ arrives at a given location, she is automatically notified about nearby products that are known to be of special interest to her.

In both of these schemes, the location of $\mathcal{U}$ is seen as being determined by other factors, and the job of $\mathcal{S}$ is to help $\mathcal{U}$ do as well as possible in or near that location.[5] Consequently, $\mathcal{U}$ may discover that no especially desirable products are available near her current location even though there exist more attractive ones at some location that she might just as well have visited.

The alternative approach to be examined here takes into account the fact that people often have considerable freedom in determining what particular locations to visit. To take an example that involves a small geographic region, consider a user who wants to walk through a large shopping mall on the way to some other destination and would like to pick up a few items along the way—providing that she can do so without investing too much extra time. Instead of taking the shortest route through the shopping mall and finding out about products that happen to be available along that route, she may well be willing to take a longer route that leads her by places where she is more likely to find the desired products. Suppose now that $\mathcal{U}$ is to be guided through the shopping mall by a system running on a handheld device: It makes sense for the system to apply techniques similar to those used in the more sophisticated route planning systems for automobiles, which choose a route according to a mixture of criteria such as driving time, gas consumption, and scenic attractiveness.

[4] To facilitate exposition, we will use the term *products* even when what is really meant is products, services, and combinations of both. The symbol $\mathcal{S}$ will denote the system under discussion, while $\mathcal{U}$ will refer to a user of that system.

[5] A system may provide navigation support to guide $\mathcal{U}$ from her current location to the exact nearby place where she can obtain the recommended product; but the current location is still treated as given.

On the other hand, the planning problem in the shopping mall has an important difference from the typical route planning problem: the need to deal with the uncertainties that are inherent in almost any attempt to match customers with products. For example, if $\mathcal{U}$ is looking for a specific book, it is in general hard for a system to know with certainty in advance whether a given bookstore has that book in stock. If $\mathcal{U}$'s interest in books is more indefinite (e.g., "Something amusing to read on the plane"), there will be uncertainty for another reason: Even if $\mathcal{S}$ has exact information on the available books, $\mathcal{S}$ cannot predict with certainty if $\mathcal{U}$ will find one of them that fulfills her requirements.

As is explained in [2], $\mathcal{S}$ can deal adequately with this uncertainty only if it computes not a fixed plan but a *policy* in the sense of decision-theoretic planning. A policy can specify, for example, that if $\mathcal{U}$ does not find a suitable book in Bookstore *A* she should make a slight detour to try the nearby Bookstore *B*.

One of the issues to be addressed in the current paper is whether this type of decision-theoretic planning can actually lead to more effective shopping in the context of a working system. And even if the shopping is more effective, how do users evaluate it subjectively? Although we will address this question within a mock-up of a shopping mall, the basic method is applicable on very different scales, such as shopping within a city.

A second main issue that is specific to smaller scenarios concerns the usability of a particular solution to the problem of providing precise location-awareness indoors. We test a relatively simple, robust technology that makes use of the infrared ports found on many handheld devices.

2 Realization of the Shopping Guide

For concreteness, we will explain the workings of the shopping guide with reference to the hardware used and the (fictitious) shopping mall that we set up for the user study. Consider a user $\mathcal{U}$ who enters a mall with the intention of buying five products, if possible: a CD, some bread, a newspaper, a PDA, and a plant. As the partial map in Fig. 1 indicates, products of each of these types are available in one or more stores. But when deciding which stores to visit, $\mathcal{U}$ needs to take into account the facts that (a) it is not certain that she will find a truly suitable product even in a potentially relevant store; and (b) stores differ in the time that it typically takes to wait in line to pay for a product.

The goal is to allow $\mathcal{U}$'s PDA to determine at each point in time which store $\mathcal{U}$ should visit next. Since the necessary planning (described below) is computationally expensive, it is performed on a stationary computer at which $\mathcal{U}$ enters a specification of the desired products through some suitable form of interaction (e.g., menu selection on a touch screen). The computer then computes a *policy* for moving through the shopping mall. The policy is compact enough to be transferred to $\mathcal{U}$'s PDA, and it contains all the information that the PDA-based application requires.

As Fig. 2 indicates, the PDA guide should at any moment direct $\mathcal{U}$ toward the next store by displaying a navigation instruction on the PDA screen; $\mathcal{S}$ will determine the

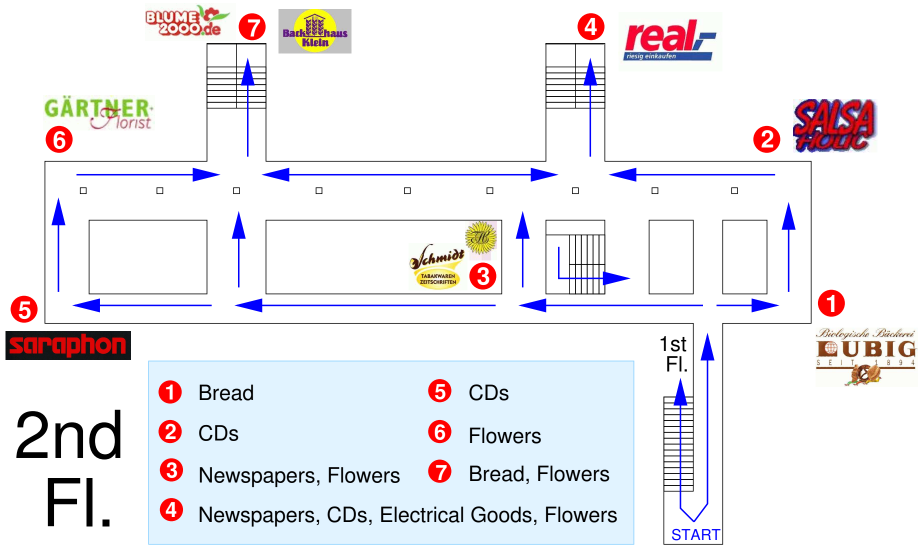

Fig. 1. *Part of one of the maps representing the shopping malls for the user study. Icons represent stores, arrows indicate the directions $\mathcal{U}$ is permitted to walk to.*

current location and orientation[6] by receiving an infrared signal from one of a number of beacons affixed to the walls, and $\mathcal{S}$ will know what products $\mathcal{U}$ has already found, because $\mathcal{U}$ will check each one off after purchasing it.

2.1 Hardware

The particular hardware configuration employed in the user study (see Fig. 2) gives an idea of the typical hardware requirements: The specific PDA was a Compaq iPAQ 3620 Handheld, which can receive signals from the small infrared beacons via its built-in IrDA interface. We mounted 30 of these beacons at decision points (i.e. turns, crossings, stairs) in the building in which the study was conducted, each at a height of about 2.5 meters. The beacons emit a 16 bit ID twice a second, and they have a transmission range of about 7 meters. Running on batteries, they do not need access to the power supply system or to a data network.

The PDA software for the shopping guide is implemented in TCL/TK. When a beacon signal is received, the browser displays the corresponding navigation recommendation according to the precomputed navigation policy, taking into account both the beacon's location and the products that $\mathcal{U}$ has checked off on the shopping list shown on the right-hand part of the screen. If a beacon ID corresponding to a store location is received, the items offered in that store are highlighted in the shopping list.

[6] A beacon transmits signals only within a small angle. Therefore, receiving signals from a beacon means pointing with the PDA's infrared interface in the direction opposite to that in which the beacon transmits signals.

Fig. 2. *The hardware of the PDA-based shopping guide.*

(Left: Example display on the PDA, which receives location signals from a beacon (middle). Right: Close-up of part of the display in the situation where the user has arrived at a relevant store. $\mathcal{U}$ checks off the items found in the current store and clicks "Okay" to proceed with the next navigation recommendation.)

2.2 Decision-Theoretic Planning

The techniques used for the decision-theoretic planning are closely related to those described in [2] and [7] (see [1] and [3] for discussions of the more general family of decision-theoretic planning techniques). Since these techniques are not the focus of the present paper, they will be described only in enough detail to permit an overall understanding of how the shopping guide works.

The result of the planning process is a navigation policy that allows $\mathcal{S}$ in each situation to recommend the action with the highest expected utility for the user. We use a Markov decision process (MDP) to model the dynamics of the shopping trip. Our basic assumptions are the following: (a) $\mathcal{U}$ moves from a starting point to a destination (entry and exit of the shopping mall) via distinct positions in the building. (b) On her way she wants to find as many items of her shopping list as possible in a relatively short time. (c) In each of the stores she has a chance to purchase some of the items needed. (d) In each of the stores in which she finds an item, she will spend a certain amount of time waiting in line to pay for it.

We encode this information in the state features and the transitions between the states as follows: A state is defined by (a) the position of an infrared transmitter, (b) a list of the desired items and the items already purchased, and (c) a flag indicating the status of $\mathcal{U}$'s shopping procedure at that position—a feature that is significant only if there is a store at the position in question. This flag has one of four values: (1) NEUTRAL, indicating that $\mathcal{U}$ has not yet checked if any of the desired articles is available at the current store; (2) NOT FOUND, indicating that none of the articles needed is available in the store; (3) FOUND, indicating that $\mathcal{U}$ has determined that at least one of the articles needed is available; and (4) BOUGHT, indicating that $\mathcal{U}$ has bought at least one article.

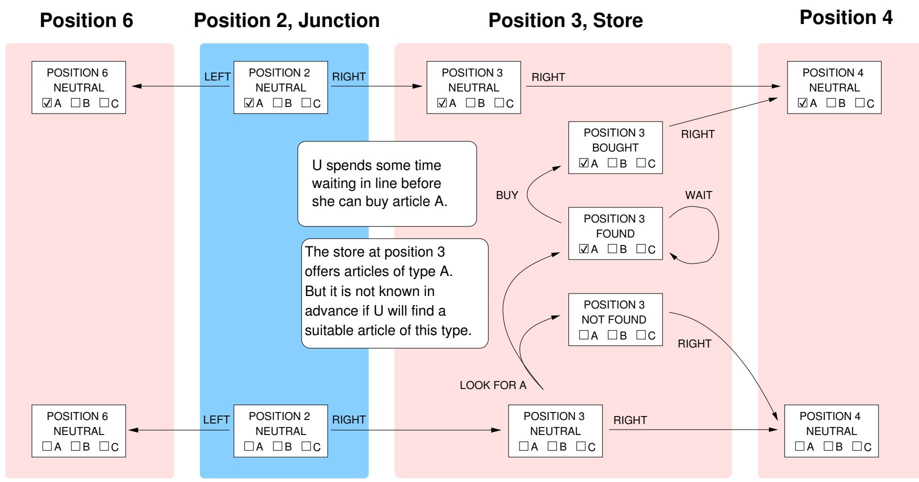

Fig. 3. *Illustration of the modeling underlying the decision-theoretic planning employed in the shopping guide.*
(Explanation in text.)

Figure 3 shows an example of the system dynamics. Suppose $\mathcal{U}$ approaches a junction located at position 2. Whether or not she has already purchased article A (corresponding to the two states shown for position 2), $\mathcal{S}$ has the choice between recommending LEFT or RIGHT, which would lead $\mathcal{U}$ to position 6 or position 3, respectively. Each of these transitions has a cost that depends on the time it will take $\mathcal{U}$ to reach the next position.

Suppose that $\mathcal{U}$, who has not yet purchased article A, approaches the store at position 3, which offers article A. In this case, $\mathcal{S}$ can recommend RIGHT to have $\mathcal{U}$ pass by the store and go ahead toward position 4. But $\mathcal{S}$ can also recommend that $\mathcal{U}$ enter the store and LOOK FOR ARTICLE A. In this case, $\mathcal{S}$ needs to consider what $\mathcal{U}$'s chances are of finding article A in this store. Moreover, looking for article A is associated with a cost: the usual time needed to check if article A is available or not. If $\mathcal{U}$ finds article A in the store, it remains uncertain how long she will have to wait in line to pay for it. This waiting is modeled in a simple fashion by a circular WAIT-transition which is associated with a particular probability of the form $(n-1)/n$ for an integer n. This model would be entirely accurate if (a) the cashier required a constant amount of time to deal with each customer, (b) there were always n customers waiting, and (c) the cashier always chose the next customer at random from the n who were waiting. These assumptions will seldom be satisfied exactly, but they do make it possible to approximate a broad range of more realistic waiting processes. Once $\mathcal{U}$ has finally BOUGHT article A, $\mathcal{S}$ will give the next navigation recommendation: In this case the only option is RIGHT.

The Markov decision process described here is *fully observable*: The system always knows with certainty $\mathcal{U}$'s position, the items needed and the items already purchased. The rewards for the goal states, which are associated with the position at the mall exit, are a function of the number of articles that $\mathcal{U}$ has bought during her shopping trip.

Given this representation of the planning problem, we can use a standard algorithm, such as value iteration (see, e.g., [3]), to compute the navigation policy. This computation takes only a few seconds, since the number of state features is small in the problem at hand.

The modeling underlying the decision-theoretic planning can in principle take into account a variety of factors other than the ones included in our current implementation, such as the difficulty of following each instruction correctly (cf. [7]). The main constraints on the modeling are due to complexity considerations, as will be discussed at the end of the paper.

3 User Study: Method

The overall purpose of the user study was to check the objective effectiveness and subjective acceptability of the shopping guide just described, with regard to both its planning methods and its hardware solution.

So as to have a point of comparison for the results with the PDA guide, we also had subjects perform the same basic task with a conventional paper map (see the example in Fig. 1). The map was designed to be at least as helpful as the best maps that can be found in shopping malls: It was portable (in contrast to the frequently found wall-mounted maps, which burden users' memory), and it showed all of the relevant information that could reasonably be expected to be found on a paper map. This information did not include probability distributions for waiting times or for the likelihood of finding suitable items in particular categories.

3.1 Mockup of the Shopping Malls

One configuration of 30 beacons was installed on two floors of the main Computer Science building at Saarland University. Two alternative shopping malls (of equal navigational complexity), each comprising 15 stores, were defined, so that each subject would be able to perform a distinct shopping task with each of the two types of guidance. Figure 1 shows the upper floor of one of these malls. It was not feasible to decorate the building with physical representations of stores, but the subject could always see from the map or the PDA display when she had reached a given store.

In addition to the physical constraints on movement enforced by the architecture of the building, we added some further constraints that might be found in a more complex setting: Some segments of hallways were designated as "one-way passages", as is indicated by the arrows in Fig. 1. Subjects using the map were instructed not to walk in the wrong direction on these segments, and the instructions given by the PDA guide likewise obeyed these constraints. Analogous constraints are found in real shopping malls in the form of (a) escalators that run in only one direction at a given location; (b) entrances through which shoppers are not allowed to exit, and vice-versa; and (c) one-way arrows drawn in aisles within stores. Since we designated a relatively large number of one-way segments in our malls (partly for technical reasons that will be discussed below), our malls were presumably untypically challenging in this respect.

Experiments on consumer behavior in which subjects are only pretending to buy products are frequently criticized on the grounds that the subjects' motivation is untypical. But in the present study, the focus was not on the buying of products but rather on the process of navigating to find the products.

3.2 Shopping Tasks

Each subject used the paper map and the PDA guide in two successive trials.[7] The subjects' task in each case was to buy as many products as possible from a shopping list with 5 categories: bread, a plant, a CD, a newspaper, and a PDA.

Each of the following combinations and temporal orderings was employed with 5 of the 20 subjects:

- PDA in mall 1; map in mall 2
- PDA in mall 2; map in mall 1
- map in mall 1; PDA in mall 2
- map in mall 2; PDA in mall 1

Therefore, neither learning effects nor differences between the two malls could bias the overall pattern of results.

3.3 Subjects

We recruited 20 subjects (15 of them female). Of these subjects, 14 were students, 3 were employed persons, 1 was a homemaker and 1 was an unemployed person. Only 2 of the subjects had used a PDA before, neither of them for a navigation task. Most of the subjects knew what a PDA is and what it is normally used for. The subjects' reward was made up of (a) a fixed amount of money (15 German marks) plus (b) an additional performance-dependent amount: They received 2.00 marks for each product they managed to "buy", but 0.01 mark was deducted for each second that they spent shopping. (This cost of time was explained by analogy to the parking costs incurred during shopping.)

3.4 Procedure

Each subject performed the tasks individually under the guidance of an experimenter. First, the experimenter explained the nature of the tasks. Then he familiarized the subject with the paper map (e.g., how the paths between stores were visualized) and with the use of the PDA guide (e.g., how to hold the device to ensure good reception of the beacons' signals). During the explanation of the reward policy, the subject was instructed not to race through the building in order to maximize her financial reward but rather to walk at a normal pace, obeying the constraints shown on the map (when using the map) or following the PDA guide's directions.

[7] A within-subject design was chosen over a between-subject design because of (a) our expectation that individual differences would be large; and (b) our desire to hear the comparative comments of subjects who had used both types of shopping guidance.

In each of the two trials, the subject began at the mall entrance and was followed around by the experimenter. When the subject reached a store, the experimenter took on the role of the sales clerk at that store: He informed the subject about the availability of the desired product(s); and if at least one was available, he determined how long the subject would have to wait in line to buy it, enforcing the waiting time with a stopwatch.

After both trials had been completed, the subject filled in a questionnaire and was interviewed briefly.

4 User Study: Results

4.1 Objective Results

As is illustrated in Fig. 4 (left), the PDA guide did enable subjects to complete their shopping task 11% faster on the average than with the map, the difference being statistically significant according to a paired t-test ($t(19) = 2.65, p < .05$). This difference in itself may not seem to constitute a practically important advantage, but it seems more noteworthy when we consider two factors:

- In the map condition, subjects were allowed to study the map as long as they wanted before starting to shop, and this study time was not counted. In the PDA condition there was nothing to study, and the subjects set off immediately.
- Whereas all subjects were familiar with paper maps, none had previously used a PDA navigation aid like the one in this study, and only 2 had used PDAs at all. So the time for the PDA condition includes any time that subjects might have required to get used to using the PDA guide and to dealing with the infrared signals from the beacons.

The map and PDA conditions were just about identical in terms of the number of products that subjects found. It is not surprising that there is no reliable difference here: The number of products found is a relatively coarse-grained variable, the possible values being just the integers between 0 and 5; and the situation was set up in such a way that in general most of the desired products could be found, so there was little room for differences.

It is unclear how well these quantitative results would generalize to other concrete scenarios and to qualitatively different users. Still, the results for this study demonstrate that the PDA guide can fulfill its promise of enabling shoppers to complete their task more effectively despite the two handicaps just listed.

4.2 Subjective Results

While the objective effectiveness of the PDA guide may be a necessary condition for its ultimate acceptance, it is not a sufficient condition. An important question is how attractive subjects would find the prospect of using such a guide in real shopping (bearing in mind the fact that they had no opportunity to experience it in truly realistic conditions). As an assessment of overall preference, the questionnaire included a question that asked which method they would prefer to use if they had to perform further shopping tasks

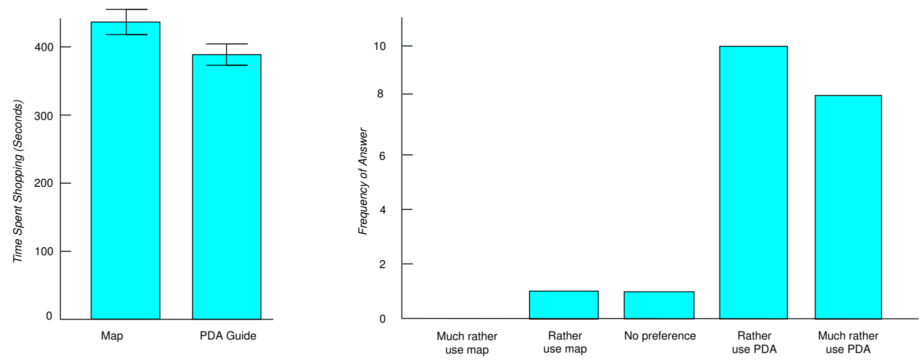

Fig. 4. *Left: Mean values for the quantitative performance index for each of the two types of shopping guidance. Each bracket indicates the standard error of the corresponding mean. Right: Frequency distribution of answers to the question "Which method would you rather use if you had to perform more shopping tasks like the ones you performed in this study?"*

like the ones they had just completed. Figure 4 (right) shows that the subjective preference is much clearer than the objective advantage. It is possible that some subjects were biased in favor of new technology or wanted to please the experimenter (cf. [8]).

It might be suspected that this difference (and perhaps the objective speed difference as well) were due to specific inadequacies in the design of the paper map. But subjects' responses to the relevant questions made it clear that they saw fewer correctable problems with the map than with the PDA. Nevertheless, the usage of multiple maps put up in important locations of the building (and designed according to principles such as those in [4]) instead of a map carried along by the subjects might facilitate the subjects' orientation in the map condition.

On the questionnaire, subjects were encouraged to add verbal comments to their ratings. In addition, the questionnaire contained open questions asking about (a) problems and sources of irritation with the map and the PDA guide; and (b) ways of improving the PDA guide. Since in many cases closely related comments were made to different questions, we summarize them here not in terms of the individual questions but in terms of the overall pattern of ideas that emerged.

4.3 Reasons Given for Preferring the PDA guide

Saving Time and Effort The most obvious advantage of the PDA guide is that the shopper can typically complete her tasks in less time than with a map. As one subject mentioned, there is also a corresponding savings of effort in moving around the shopping mall—an advantage which some users may find appealing even if they are not concerned about time, for example if their mobility is for some reason limited or if the mall is much larger than the one used in this study. These comments confirm that subjects recognized the superior ability of the PDA guide to deal with complexities such as uncertainties about product availability and waiting time—despite their unfamiliarity with systems of this sort and with the underlying technology. The comments also confirm that the potential advantages conferred by the PDA guide's knowledge and reasoning were not seen by subjects as being outweighed by usability problems or unfamiliarity.

Less Cognitive Effort and Frustration A less objectively measurable advantage is the reduced need for mental exertion with the PDA guide. As one subject put it, "You don't have to plan where you're going to go, or think for yourself." When using the map, most subjects recognized the advisability of planning ahead, though it would not have been feasible for them even to attempt the type of decision-theoretic planning that is performed by the PDA guide. A typical strategy was to

- plan a few steps ahead, assuming that the products they wanted to get in the nearby stores would be available; and
- replan whenever an product proved unavailable.

In addition to leading in general to suboptimal solutions, even this simple strategy requires more mental exertion than simply following the advice given by the PDA guide. Moreover, users found it frustrating to have to abandon part of a plan that they had already spent some effort devising.

Less Need for Thinking About Orientation A small number of subjects pointed out that with the map, they continually had to ensure that the map was oriented correctly and to keep in mind where they were. By contrast, the PDA guide simply gave directions relative to their current position and orientation; it can take into account the direction the user is facing to a certain extent because each beacon transmits signals only within a limited angle.[8] Our subjects' problems when navigating with the map are similar to those discussed in a more general analysis by Thorndyke and Hayes-Roth ([11]) about problems that people may have with orientation under different conditions.

4.4 Perceived Drawbacks of the Current PDA guide

Even subjects who were on the whole enthusiastic about the PDA guide mentioned some limitations, some of which can be eliminated through improvements in its design.

Feeling of Not Having the Big Picture Although, as was just mentioned, some subjects appreciated not having to think about their orientation and location, for some subjects this consequence was associated with a drawback: They felt almost as if they were being led blindfolded through the shopping mall—a feeling that was evidently not entirely satisfactory to them even if they were convinced that the system was leading them in an optimal way. Proposals made by subjects for improving their overall feeling of orientation included some ideas that are familiar from GPS-based navigation systems:

[8] Where uncertainty about orientation was possible, the PDA display included an identifiable landmark, such as a flight of stairs, located in the immediate vicinity of the user's current position.

- Display overview maps of (parts of) the mall.
  For example, while $\mathcal{U}$ is walking from one position to the next, an overview map showing her position within (some part of) the shopping mall could be displayed.
- Make it possible for $\mathcal{U}$ to be given longer sequences of directions.
  Instead of always being led from one beacon to the next, subjects sometimes prefer longer sequences of instructions (e.g., the graphical equivalent of "Go straight ahead, then take the next right and then the next left"). In addition to giving users a greater sense of knowing where they are going, such sequences can save time, because they reduce the number of occasions on which $\mathcal{U}$ needs to acquire information from a beacon and consult the PDA. On the other hand, longer sequences place heavier demands on $\mathcal{U}$'s working memory, and they lead to a greater likelihood that some instruction will be executed incorrectly. This type of tradeoff was investigated experimentally in a different setting in [7]: This study showed that finding an optimal instruction sequence length is a task that can be handled naturally within the framework of decision-theoretic planning used here.
- Give advance announcements of upcoming stores.
  Some subjects wanted to know not only what direction they should proceed in but also what they could expect to find there. For example, if the PDA displayed along with the arrow a message like "Saraphon (CDs), 40 m", $\mathcal{U}$ could utilize the walking time to prepare mentally for the task of looking for a suitable CD.

In addition to the advantages already mentioned, users who have more than the minimally necessary amount of information about their surroundings are in general better able to cope with problems with the technology, which will be discussed next.

Need to Attend to the PDA and to the Beacons The current realization of the PDA guide requires users to attend to a certain extent to the handheld device and/or to the beacons even while walking: $\mathcal{U}$ needs to know when she has reached the range of the next beacon, so that she can attend to the next instruction. $\mathcal{U}$ may therefore

1. look out for the physical beacons themselves,
2. aim the PDA explicitly at a beacon, and/or
3. keep looking at the PDA to see when a new instruction appears.

Both the first and second activities can be made unnecessary if the beacons are simply made more powerful, so that their signals reach the PDA whenever the user comes into the relevant area.[9] The third activity can interfere with $\mathcal{U}$'s walking, for example increasing the danger of colliding with other persons or stumbling on stairs. An obvious solution, mentioned by some subjects, is the use of acoustic stimuli, which could be delivered via a loudspeaker or a headset. The acoustic output can range from (a) acoustic signals that alert $\mathcal{U}$ to a change in the display to (b) the use of speech for all navigation instructions.

[9] This technical solution is currently being implemented by the company that supplied the beacons. But as this solution does not overcome the problem if $\mathcal{U}$ points interface at the floor, other technologies, such as Bluetooth, might be employed to indicate the presence of a transmitter regardless of the position of the PDA.

5 Further Issues

The aspects of the shopping guide that could be tested in the present study passed their tests fairly well, and the improvements that proved desirable can be realized without much difficulty. But the goal of deploying the system in real shopping situations raises some issues that were not (fully) addressed in the present study.

5.1 Increasing the Volume of Products Purchased

One important potential advantage of the PDA guide was *not* demonstrated in the results of the present study: the advantage that shoppers might find more products that they want to buy. But this advantage might well be achieved in situations with somewhat different properties than the situation of the present user study:

First, if the time available is more severely limited (e.g., because stores are about to close), shoppers with the slower map method will have to stop shopping before they have exploited all possibilities. In this case, shopping times will tend to be equalized, with the difference appearing in the number of products found.

Second, shoppers often have a choice between shopping or not shopping (e.g., on the way to their gate in an airport) or between shopping in mall *A* or shopping in mall *B*. If a given shopping mall, by providing a shopping guide, offers the prospect of getting the job done more efficiently and with less cognitive effort, this mall has an advantage over competing malls in attracting customers to shop there in the first place.

5.2 Applicability to Spontaneous Shopping Behavior

Especially recreational shopping is often highly spontaneous: A shopper may suddenly abandon any plans she may originally have had in order to pursue some attractive new option. Yet in situations where it is very hard to predict what the shopper might do, it is hard to specify an accurate model as a basis for decision-theoretic planning. The approach presented here is therefore most applicable when the shopper has particular goals that she wishes to achieve within a limited amount of time.

Our approach does not, however, presuppose that the shopper is mechanically tracking down specific products on a shopping list with maximal efficiency. If $\mathcal{U}$ is allowed to specify vague product descriptions like "something to read" or "a present for my 3-year-old daughter", $\mathcal{S}$ will lead $\mathcal{U}$ to relatively promising stores, and $\mathcal{U}$ can be as creative as she likes in deciding what (if anything) to buy there.

It is possible to extend the model that underlies the decision-theoretic planning to account for possible events such as $\mathcal{U}$'s spontaneously stopping to spend time at stores that she passes (see [2] for further examples of possible extensions). The proportion of shoppers whose behavior is adequately fit by the model's assumptions will depend on the sophistication and the accuracy of the modeling effort, an issue to be discussed next.

5.3 Complexity of Decision-Theoretic Planning

Decision-theoretic planning based on Markov decision processes has an inherently high computational complexity. Since in the present study the focus was not on computational issues, we made—and enforced—some assumptions about the reality being modeled which ensured that planning would be computationally tractable. For example, the introduction of "one-way passages" made it impossible for subjects to visit any store twice, thereby making it unnecessary for the MDP models to include information about the stores that $\mathcal{U}$ has already visited.

For the modeling of real shopping situations, it will be necessary to deal with models which have not been selected primarily because of their tractability. A great deal of research has been done in recent years on ways of making decision-theoretic planning more tractable in certain classes of situations (see, e.g., [3]). We are examining how these approaches can help with the type of planning problem involved here.

5.4 Learning of Parameters

For the modeling, a great many specific quantitative parameters need to be specified, concerning matters like walking and waiting times and the probability of finding a given type of product in a given store. In our study, we were able to construct a reality that corresponded to our model. For real use, one approach would be to estimate the parameters initially on the basis of a more or less thorough analysis of the shopping mall. If this initial specification is accurate enough to ensure that the system is of some use to shoppers, further data can be obtained from the PDAs of shoppers who have used the system, provided that they are willing to take a few seconds to allow data about their searches to be (anonymously) transmitted back to the central system. In addition to supporting estimates of the parameters considered in this paper, this method would open up the prospect of introducing some form of social navigation or collaborative filtering.

6 Conclusions

We have provided an initial demonstration of the feasibility and acceptability to users of several aspects of a new approach to intelligent location-aware shopping guidance. The method is characterized by:

- an integration of navigation assistance with product location;
- the exploitation of decision-theoretic planning techniques that perform computations too complex for humans to approximate mentally; and
- the use of robust location-awareness technology which is already available at reasonable cost.

Our study with 20 users in a semirealistic situation yielded a number of insights into the reasons why users find the system acceptable—and the conditions on which this acceptability depends (e.g., unfamiliarity with the environment or time pressure). It also turned up some problems and limitations, along with ideas for overcoming them.

Studies of the use of the shopping guide with real shoppers in a real shopping mall will undoubtedly bring to light new challenges that have not been addressed in the present paper. Still, the intermediate results yielded by our study seem likely to prove useful in connection with both (a) more realistic application scenarios of the shopping guide and (b) other types of mobile system that raise similar issues.

Acknowledgements

The research of the first and third authors was supported by the German Science Foundation (DFG) in its Collaborative Research Center on Resource-Adaptive Cognitive Processes, SFB 378, Projects B2 (READY) and A4 (REAL). Beacons for the study were supplied by Eyeled GmbH (http://www.eyeled.de).

References

1. J. Blythe. Decision-theoretic planning. *AI Magazine*, 20(2):37–54, 1999.
2. T. Bohnenberger and A. Jameson. When policies are better than plans: Decision-theoretic planning of recommendation sequences. In J. Lester, editor, *IUI 2001: International Conference on Intelligent User Interfaces*, pages 21–24. ACM, New York, 2001.
3. C. Boutilier, T. Dean, and S. Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.
4. R. P. Darken and B. Peterson. Spatial orientation, wayfinding, and representation. In K. Stanney, editor, *Handbook of Virtual Environment Technology*. Erlbaum, Mahwah, NJ, 2001.
5. S. Duri, A. Cole, J. Munson, and J. Christensen. An approach to providing a seamless end-user experience for location-aware applications. In *Proceedings of the First International Workshop on Mobile Commerce*, pages 20–25, Rome, 2001.
6. A. Jameson. Usability issues and methods for mobile multimodal systems. In *Proceedings of the ISCA Tutorial and Research Workshop on Multi-Modal Dialogue in Mobile Environments*, Kloster Irsee, Germany, 2002.
7. A. Jameson, B. Großmann-Hutter, L. March, R. Rummer, T. Bohnenberger, and F. Wittig. When actions have consequences: Empirically based decision making for intelligent user interfaces. *Knowledge-Based Systems*, 14:75–92, 2001.
8. T. K. Landauer. Behavioral research methods in human-computer interaction. In M. Helander, T. K. Landauer, and P. V. Prabhu, editors, *Handbook of Human-Computer Interaction*, pages 203–227. North-Holland, Amsterdam, 1997.
9. N. Sawhney and C. Schmandt. Nomadic Radio: Speech and audio interaction for contextual messaging in nomadic environments. *ACM Transactions on Computer-Human Interaction*, 7:353–383, 2000.
10. M. Singh, A. K. Jain, and M. P. Singh. E-commerce over communicators: Challenges and solutions for user interfaces. In *Proceedings of the First ACM Conference on Electronic Commerce*, pages 177–186, Denver, 1999.
11. P. W. Thorndyke and B. Hayes-Roth. Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14:560–589, 1982.