GUIDE: An Intelligent Attraction Recommender System

Kyle Bedell

Bentley College

Abstract

Traditional tourist guidebooks have a number of limitations that make them ill suited for information presentation on an individual basis. In attempting to be a general-purpose reference about all things to all people, they make themselves difficult to navigate. For the busy tourist in an unfamiliar setting who wants to see and do as any interesting things as possible in a limited timeframe, navigating through such a reference can be an incredibly time-intensive task.

The GUIDE electronic attraction recommendation system attempts to solve these problems through the use of an intelligent user interface. With a wireless network, location beacons, a system of locale-based servers, and a tablet-based portable client as its hardware tools, GUIDE's adaptive hypermedia system compares a user's preferences against a list of nearby locations to recommend attractions that are relevant to its user. By filtering out attractions and points of interest that its user would not visit, GUIDE helps users quickly locate relevant information about their surroundings. The intelligent agent monitors the attractions its user visits, keeping its user's profile current even if that user's tastes change.

Other features like mapping software, always up-to-date attraction information, and built-in tours help augment the touring experience, allowing travelers to get the most out of their time in a given city. As it offers a rich, targeted touring experience, GUIDE looks to be a promising alternative to paper guidebooks.

GUIDE: An Intelligent Attraction Recommender System

Introduction

Faced with the problem of "where to go and what to do in the limited amount of time available" (Cheverst et al., 2002, p.1), many tourists rely on guidebooks to help them navigate unfamiliar cities. However, because of their nature as general-purpose references, guidebooks contain large amounts of irrelevant information from the perspective of a single individual. The presence of this unnecessary content forces readers to navigate a larger information space to find items of interest to them. In addition, these guidebooks do not give readers a way to sort through their contents based on topic of interest; they are organized primarily by geographical location alone.

Both of these issues slow down travelers, as they must spend additional time locating information that is interesting to them before they can visit an attraction. Additionally, the nature of a traditional guidebook means that it cannot possess detailed or up-to-date information on each attraction. Operating hours change, and travelers may have difficulty deciding if they want to visit an attraction based on a two-sentence description.

The GUIDE research prototype, developed in May 2002 by Cheverst et. al., attempts to solve these problems of organization and information retrieval with context awareness, adaptive hypermedia, and a model of the user containing interests and other important information. By taking in information from nearby attractions and comparing their characteristics against a user's preferences, the intelligent system provides "up-to-date and context aware hypermedia information while they explore" (Cheverst et al., 2002, p.1). The benefits to the user over traditional guidebooks include improved recommendations on which attractions to visit, faster location of relevant information, and a knowledge base that is always up to date (as opposed to the 'set in stone' information in traditional written media).

In the following pages, an explanation of GUIDE's interface, hardware, infrastructure, and backend is presented to give the reader a better understanding of the prototype. This is followed by an explanation of the underlying intelligent system (how GUIDE goes about solving the problems posed by traditional guidebooks), the strengths and weaknesses of traditional guidebooks vs. the GUIDE prototype, and a justification for the development and use of an intelligent UI for tourist navigation. To conclude, various technical, privacy, and usability issues with the prototype are noted, along with possible solutions (where applicable/feasible).

An Explanation of the GUIDE Research Prototype

Interface and hardware.

Users interact with the GUIDE through a Fujitsu TeamPad 7600 Tablet PC's touch-sensitive screen. Although the hardware inside is extremely dated by modern standards (Pentium I at 166MHz, 16MB RAM, Windows 95), this low-power configuration offers improved battery life over modern tablet units. Each tablet is equipped with an 802.11b wireless card that functions as a communications link between the GUIDE unit and a local broadcast server.

The actual interface closely resembles a web browser, and the GUIDE team found "that users with previous web experience found the interface reasonably intuitive" (Cheverst et. al., 2002, p.2). A sample image of the UI is shown below, courtesy of Cheverst et. al. (2002).

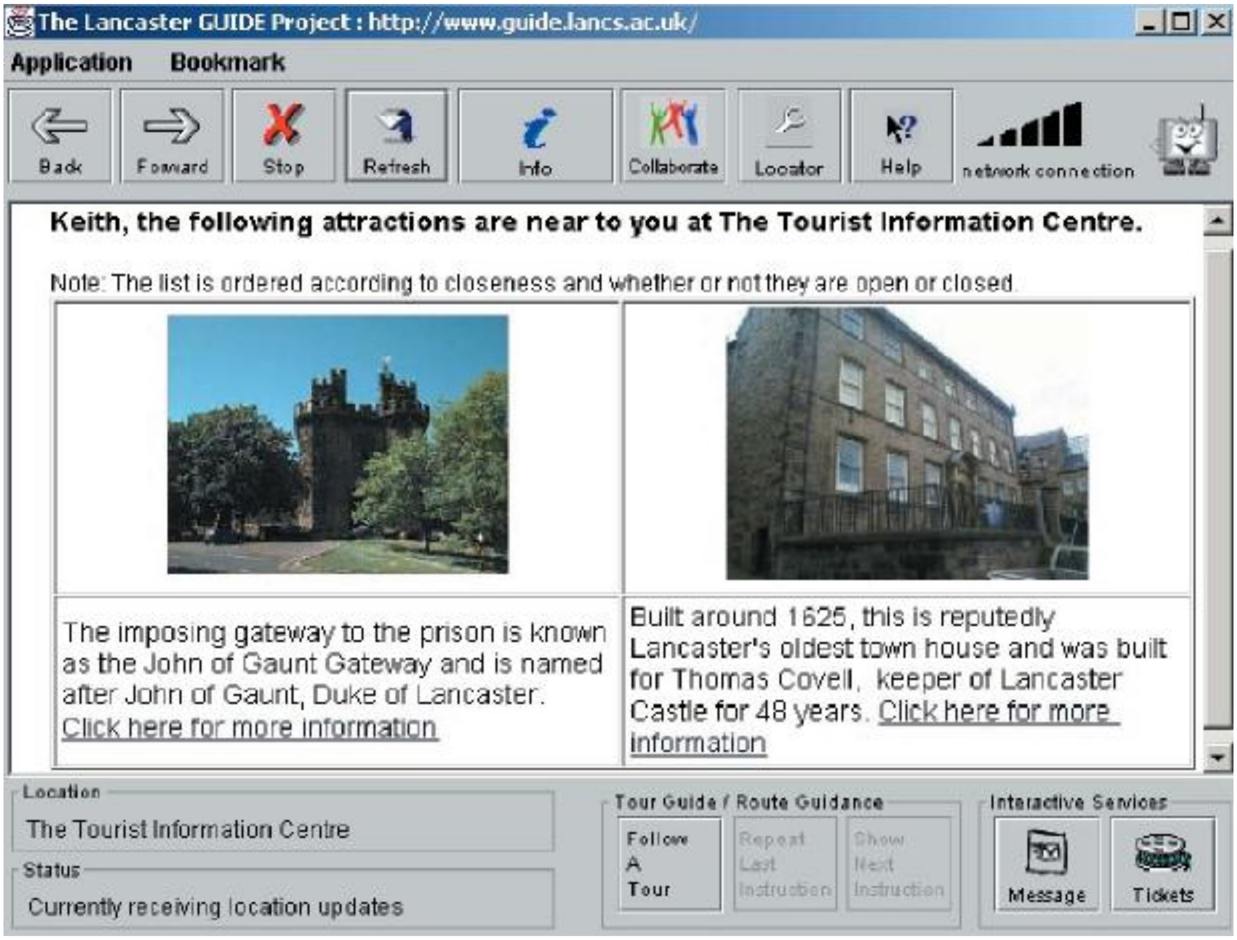

At any point in time, users can press the 'Info' button in the interface to request a list of hyperlinks detailing possible areas of interest; once selected, these hyperlinks allow users to drill down to progressively more detailed levels of information. GUIDE can also offer tours of specific routes in the city, and provides mapping features to allow users to orient themselves in unfamiliar territory.

Infrastructure and backend.

Each attraction connected to the GUIDE system relies on a local server (one for each area of the city) to relay information to nearby GUIDE units. Each local server receives updates from a central GUIDE master server. This information consists of "Web pages, dynamic updates, and coarse-grained positioning information" (Cheverst et al., 2002, p.2). However, the web pages that GUIDE units receive have been customized with additional metadata tags to support the transmission and display of dynamically retrieved (as opposed to static) information. These include the state of an attraction (operating hours, etc.), nearby points of interest (neighbors), and a description of the location.

In its disconnected state (when GUIDE is out of range of a wireless signal), the unit relies on locally cached information and user-inputted location data to determine recommendations. Interactive services (like messaging the staff of a particular attraction or booking a ticket) are also unavailable in this mode.

Underlying Intelligent Technology

Adaptive Hypermedia

Explanation.

Brusilovsky (2001) defines an adaptive hypermedia system as one that "builds a model of the goals, preferences and knowledge of each individual user, and uses this model throughout the interaction with the user, in order to adapt to the needs of that user" (p.1). While a static hypermedia system provides the same content to all users who access it, an adaptive system takes a user's preferences, habits, methods, and/or other pertinent information into account when determining what information to display. A "static virtual museum will offer the same ``guided tour'' and the same narration to visitors with very different goals and background knowledge" (Brusilovsky, 2001, p.1), while an adaptive museum will custom-tailor each visitor's tour to their own unique knowledge bases and goals. The benefits of this are readily apparent; targeted information delivery improves each user's experience interacting with the system by only delivering what the user wants to read or hear.

Adaptive Hypermedia as Applied to the GUIDE System.

The GUIDE system displays its content dynamically depending on a user's location *and* personal preferences, as opposed to simply delivering information to a portable unit based on location. The following diagram from Cheverst et. al (2002) explains the generation of a Nearby Attractions page:

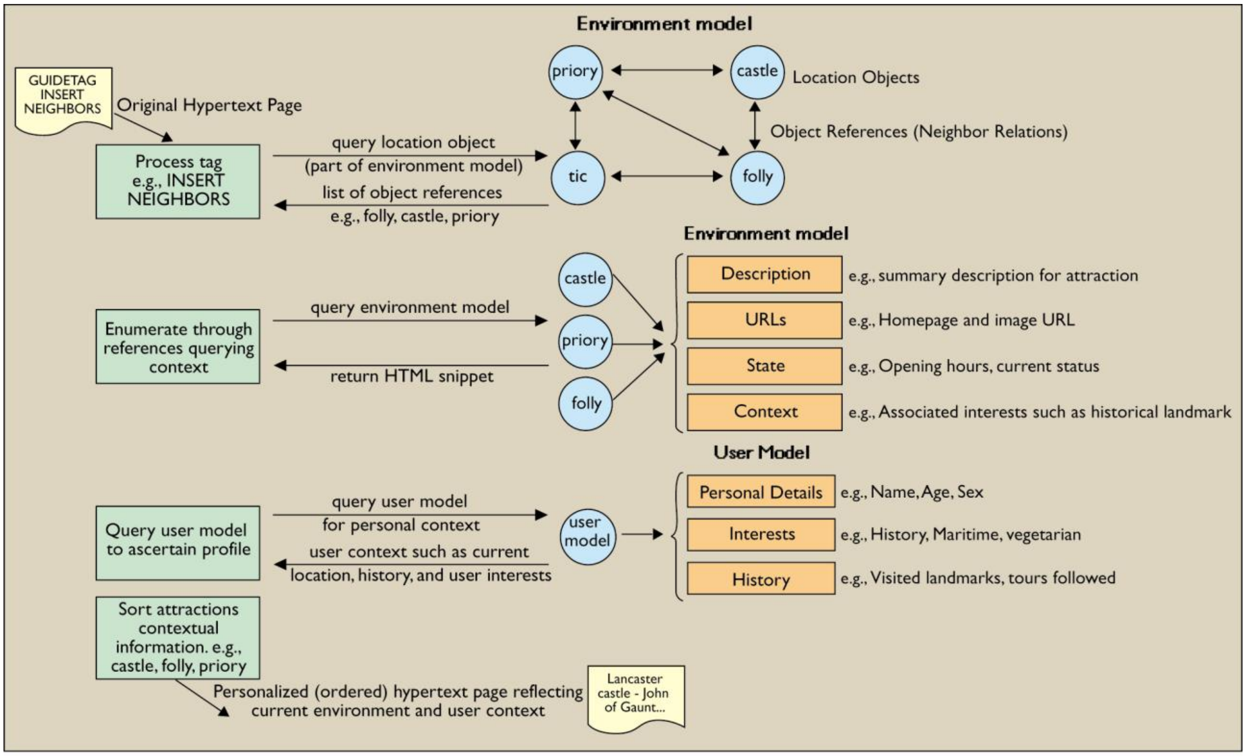

The system relies on an environment model to determine what information to send to a GUIDE unit. Based on available location data, the IUI queries the locale's environment model to determine nearby objects of interest (based on weighted relationships formed between points of interest and one's current location). For each location found, various properties are filled in, including a description of the attraction, a homepage URL where applicable, the state (operating hours and status), and some context (possible associated interests) based on the most current information available from the local servers.

Armed with this 'master list' of nearby attractions, the system then queries the local GUIDE unit for its user's profile. The profile, which contains the individual's personal details, interests, and history (tours followed, landmarks already visited), functions as a filter that eliminates uninteresting or irrelevant attractions from the list. A personalized page with up-to-date and location-aware information is then passed to the GUIDE unit and displayed on the screen.

As a user visits additional attractions, the system attempts to learn additional information about their interests by observing. For example, let us assume that one were to indicate their preferences for historical attractions dating from 1400 A.D. onward. If this individual ended up visiting many attractions dating from 1300 A.D. onward, GUIDE would take note of this and modify their user profile accordingly, expanding the range of listed attractions based on construction date. This learning process is a key part of the *adaptive* hypermedia implementation.

Differences Between Intelligent and Non-Intelligent Versions

The IUI version of the tour guide is focused on depth of information presentation; it takes a general environment model and distills it down to what a particular user wants to read. However, it retains access to a wide breadth of topics (hence, the inclusion of a search feature that does not perform any filtering). A paperback guidebook's purpose is to act as a general reference (breadth of information); it covers many topics, but only in a superficial way. Size and weight limitations prevent a paperback guide from holding all possible relevant information about each attraction.

Intelligent User Interface Version

Strengths of the intelligent version.

The intelligent tour guide offers a number of strengths over the paper version. GUIDE's information is always up-to-date because it retrieves information from a nearby location server that can be updated at any time. The fact that the system is location-aware relieves the user of having to look up his location on a map if lost (as a built-in map is provided), and allows quick access to relevant local landmarks.

GUIDE's greatest strength lies in the use of its customizable user profiles. This means that any individual can turn the GUIDE system into a personalized guidebook, filtering out irrelevant information and improving retrieval capabilities. This leaves more time for a tourist to enjoy attractions that he or she expresses interest in, thus making the touring experience more meaningful.

Weaknesses of the intelligent version.

However, the electronic nature of the intelligent version does pose some problems. Reliance on a battery-powered device means that running out of power is always a possibility, even with energy efficient hardware. The system's recommendations are also only as good as its user model; a poorly fitted user model means that GUIDE's recommendations will suffer. How does the system compensate for groups? Is there a user profile for each individual on the unit that is applied to the environment model in sequence? In addition, as GUIDE's location-awareness abilities only function when in range of a wireless location beacon. Users must enter their coordinates manually when out of range in order to receive relevant information from the GUIDE servers.

The largest weakness of the GUIDE system lies in its knowledge base architecture. As Maes (1994) notes, "a large amount of application-specific and domain-specific knowledge has to be entered into the agent's knowledge base. Little of this knowledge or the agent's control architecture can be used when building agents for other applications" (p.32). In essence, the large amount of customized, agent-specific programming and a narrow-yet-deep knowledge base makes designing an agent like GUIDE an expensive endeavor. Its parts are of little use in other applications.

There is also the issue of user/agent trust to consider. An agent that is not sophisticated will most likely provide poor recommendations, causing the user to doubt its usefulness. However, as Maes (1994) points out, "it is probably not a good idea to give a user an interface agent that is very sophisticated, qualified, and autonomous from the start…since the agent has been programmed by someone else, the user may not have a good model of the agent's limitations, the way it works, and so forth" (p.32). Shneiderman (1983) notes that users are more comfortable with a system when they can "initiate an action, feel in control, and…predict system responses" (p.65). Having a competent (but imperfect) agent allows users to retain a sense of control over the system by to allowing them to learn its limitations through experience.

GUIDE's approach to recommendations does not grant the user large amounts of control; it only relies on initial feedback provided at user profile creation and occasional updates when a user visits multiple attractions that are strongly connected by subject material. In order to alter a user profile after-the-fact, one would have to visit locations that the GUIDE system may not even display based on data from the original user profile. While the implicit approach to collecting information about the user may work well with other hypermedia systems, GUIDE's nature makes it difficult for users to modify their interests. Bra (1999) notes that the benefits of an adaptive hypermedia system come from a "much more fine-grained user model," (p.1) and GUIDE must find ways to better construct that user model if it is to make full use of the benefits of adaptive hypermedia.

Paperback Version

Strengths of the paper version.

A book never runs out of batteries in the middle of a tour, and is much easier to read in direct sunlight. A paperback guidebook is also a comprehensive general information source;

multiple people with different interests could have difficulty with GUIDE’s reliance on an individual user profile.

Weaknesses of the paper version.

Paperback guidebooks can be extremely heavy, and their function as a general-purpose resource means they will lack detailed information on specific topics. In addition, their static form of information presentation means that even when a book first comes off the publisher’s printing press, the information contained within may already be out-of-date. In addition, as previously stated, they rarely support multiple sorting and organizational schemes, forcing the reader to look through each locale’s chapter/section to identify relevant points of interest. A history lover, for example, rarely has a way to pull out *just* historical points of interest out of a paper guidebook.

Justification for the Intelligent User Interface

Merits of the IUI

Despite its flaws, the addition of an intelligent user interface to a tourist guidebook does solve a number of issues. While some of GUIDE’s features (access to a large pool of information) could be replicated with a static UI, the IUI has merits that make it a better solution. The context-awareness features that GUIDE provides means that users do not have to look up their location on a map or flip through chapters of a book to find details on nearby attractions. Nearby attractions are always displayed on GUIDE’s interface, readily accessible when needed. In addition, the filtering process keeps information targeted towards a particular user, even when preferences drift or change over time. If a static UI provides filtering capabilities, they generally do not change with the preferences of the user. Once defined, a static UI’s filters are incapable of learning from a user’s habits.

Print media also lacks the capability to adapt and filter information. If a user dislikes museums, a paper guidebook will still present that information to the user. GUIDE would consider a user's personal preferences and only display things of interest.

Possible Improvements to Efficiency

With the current GUIDE system, users must manually request information about the area and select a specific attraction onscreen before reading information about it. Since the intelligent agent residing in the system possesses knowledge of a user's preferences and location information, there should be an option to automatically update the display with relevant attraction information. Using a radius-based system, GUIDE could automatically display information when a user walks within X distance of an attraction. This eliminates an entire step of the information retrieval process for those willing to place a little more trust in the agent. The fact that the user can toggle this option allows them to retain as much control as they wish, avoiding problems with user/agent trust discussed earlier in this paper.

Privacy, Technical, and Usability Issues

Privacy Concerns

The very nature of how adaptive hypermedia systems collect and store their user information seems to run contrary to the privacy concerns voiced by many users in the internet age. With the prevalence of identity theft and misuse of personal information in our time, it is understandable that users are wary about supplying personal data to any sort of information technology. Kobsa (2002) notes that current adaptive hypermedia systems "appear to be in conflict with privacy laws and guidelines that usually call for parsimony, purpose-specifity, and user awareness or even user consent in the collection and processing of user data" (p.1). Linking user profiles to identifiable or *potentially identifiable* persons without their prior consent could be a violation of many international privacy laws. This raises concerns for adaptive hypermedia, as the foundation of the technology requires the observation and analysis of user patterns and habits.

The solution here (entirely possible in the case of GUIDE) is to allow anonymous or pseudo-anonymous user profiles to exist. Kobsa and Schreck (2003) have proposed a user model that not only allows for anonymity, but also provides the same level of customization that non-anonymous systems do. The proposed system preserves three different types of anonymity, *environmental* (determined by external factors), *content-based* (identification by means of exchanged data is impossible), and *procedural* (senders and receivers cannot identify each other when messages are exchanged between them) (p.157). This is because “to protect users’ privacy through anonymity, all of the above three types of anonymity must be present in a user-adaptive system” (Kobsa and Schreck, 2003, p.157).

To accomplish this, Kobsa and Schreck (2003) relied on a combination of “a permissions server for role based access control, a mix network for hiding the identities of users and of user modeling servers, secure transport, a certificate directory, and a reference monitor that safeguards the access of user modeling clients to user models located in the user modeling server” (p.176). This model allows users to remain unidentifiable to external sources while supporting relevant customization on a user-to-user basis.

Usability Issues

Hardware.

The weight, size, and general specifications of the Fujitsu TeamPad 7600 tablets used for the mobile GUIDE units were acceptable given the hardware requirements of the project. However, with the arrival of high-resolution smartphones and handhelds that weight just a few ounces, there is no reason why the GUIDE system could not be better adapted for mobile use. Outdoor displays of equal or better quality are now available on these portable computing devices, and new capacitive touch-screens like those on the iPhone eliminate both positioning inaccuracies and the need for calibration found on older resistive versions.

The bulkiness of the Fujitsu machine is also an issue. Although the 850g (~2lb) machine is much lighter than many other portable computers, it is not pocketable. Because of this, the user needs to carry the unit around by hand or invest in a backpack/ bump case for transport. A PDA or smartphone-based unit would be easier to transport, and be less obtrusive that a standard tablet. Already, success has been reported in using smaller devices for navigation. Disney reports that its guests have given positive reviews of a Nintendo DS-based navigation/tour guide system for its Disney World resort (Kohler, 2008, p.1).

Interface.

The use of a web-browsing metaphor for GUIDE's interface is a clever one; many people are familiar with the concept of browsing through web pages, and GUIDE's method of information presentation is analogous to that of a web browser's methods on the surface. However, there are a number of problems with using hypertext for information navigation purposes. Wright (1991) states that "in hypertext it is possible to create an irregular patchwork of information that has no structure to be grasped, and therefore no labels that can unambiguously identify where HERE is. [With] such freedom, the reader's cognitive resources used for planning where to go, and keeping track of where they have been, may become overloaded" (p.1). GUIDE's location-based nature minimizes this problem by making the user's 'location' in the information space correspond directly to their physical location in the real world.

However, when deprived of its location-tracking abilities (recall that GUIDE relies on location beacons with limited transmission ranges instead of GPS to track its user's coordinates) users can become lost easily. Forcing users to manually update their location when the system cannot automatically is unintuitive, and defeats the purpose of carrying a location-aware device. A solution to the beacon-based system's issues is discussed in the following section.

Technical Issues

Reliance on location beacons instead of the more accurate GPS.

A mobile GUIDE unit coupled with Bluetooth GPS or built-in GPS eliminates the need for location beacons while reducing the amount of polling the system has to do to continually update its location. Today's GPS devices are more accurate than ever before, and work almost everywhere on the globe. GUIDE's current location beacon system poses a major problem, as a location aware device is not nearly as useful if it cannot identify its location (being out of range of a location beacon).

Use of non-ubiquitous wireless links.

Despite the fact that the creators of GUIDE settled on a "cell-based wireless networking infrastructure" (Cheverst et. al., 2002, p.2), they did not rely on more ubiquitous cellular telephone data infrastructure, instead preferring to use traditional 802.11 wireless access points coupled with location beacons. Today's cellular networks cover more area than any 802.11a/b/g networks do, and assisted GPS can provide rough location information. Using a cell data network would both improve the range of the system and eliminate the need for access points around every point of interest. Bandwidth is not an issue, as the information GUIDE presents to the user is in the form of small images and text, both of which are quickly transmitted over an EDGE (59.2kB/sec) or EV-DO (307.2kB/sec) cellular link.

Older hardware.

The Fujitsu TeamPad 7600 hardware used in the GUIDE prototype was four years old when the project was completed in 2002. In 2008, today's portable devices (smartphones and handhelds) have exceeded the TeamPad 7600's hardware specifications. The GUIDE system could work well if it were adapted to today's mobile computing devices, reducing the size of the system and improving its battery life and outdoor readability. Based on the technical reasons Cheverst et. al. (2000) initially chose their tablet devices ("transreflective screen…light [weight]…performance…and driver support") (p.22), a device like the Apple iPhone (with its high-density 480x320 pixel display, integrated GPS, EDGE data connection, 8-hour battery life, and 620MHz ARM CPU) provides everything an improved GUIDE system would need to function in a smaller package.

Concluding Remarks

Even lacking these improvements, GUIDE offers a compelling alternative to traditional guidebooks. The use of an intelligent agent to display custom-tailored lists of attractions to readers is an excellent solution to the information navigation problem. In addition, GUIDE's ability to download status-related data (like hours of operation) means that users will always be armed with the most current information possible for a given attraction. The built-in mapping software reduces anxiety related to being in unfamiliar places, and being able to communicate with the staff of an area or order tickets on-the-fly is a welcome addition.

Possible technical improvements to the GUIDE system include replacing location beacons with integrated GPS for improved positioning information, reducing the overall size of the system to improve portability, and replacing 802.11b wireless radios with their cell-based equivalents for more ubiquitous data access capabilities.

On the intelligent agent front, privacy issues need to be taken into consideration. The creators of GUIDE may want to consider adopting the anonymity-based user model described by Kobsa and Schreck, lest they conflict with various international privacy laws. The agent also needs to find a way to expand its range of recommendations without hampering information navigation; if the only way to alter one's user profile it to visit different attractions, and the system never displays anything other than what you've expressed interest in, a user never gets the chance to change his or her preferences. Introducing some sort of 'extra' recommendations ("Have You Thought About Visiting Place X?") into the filtered list based on reviews or some other relevant metric would allow the system to take better advantage of its adaptive hypermedia framework.

From a usability perspective, the system should support automatic information displays instead of relying solely on a user prompt. This allows the elimination of an entire step in the information retrieval process, streamlining the experience.

With these changes and improvements, GUIDE 2.0 could bring an even better touring experience to travelers around the world!

References

Bra, P. D., Brusilovsky, P., & Houben, G. (1999). Adaptive hypermedia: from systems to framework. *ACM Comput. Surv.*, 31(4es), 12. doi: 10.1145/345966.345996.

Brusilovsky, P. (2001). Adaptive Hypermedia. *User Modeling and User-Adapted Interaction*, 11(1), 87-110.

Cheverst, K., Davies, N., Mitchell, K., & Friday, A. (2000). Experiences of developing and deploying a context-aware tourist guide: the GUIDE project. In *Proceedings of the 6th annual international conference on Mobile computing and networking* (pp. 20-31). Boston, Massachusetts, United States: ACM. doi: 10.1145/345910.345916.

Cheverst, K., Mitchell, K., & Davies, N. (2002). The role of adaptive hypermedia in a context-aware tourist guide, *Commun. ACM*, 45(5), 47-51. doi: 10.1145/506218.506244.

Kobsa, A. (2002). Personalized hypermedia and international privacy. *Commun. ACM*, 45(5), 64-67. doi: 10.1145/506218.506249.

Kobsa, A., & Schreck, J. (2003). Privacy through pseudonymity in user-adaptive systems. *ACM Trans. Interet Technol.*, 3(2), 149-183. doi: 10.1145/767193.767196.

Kohler, C. (2008, January 29). Disney Guests Love Interactive Nintendo DS Guide. *Wired Blogs - GameLife*. Retrieved April 27, 2008, from http://blog.wired.com/games/2008/01/report-disney-g.html.

Maes, P. (1994). Agents that reduce work and information overload. *Commun. ACM*, 37(7), 30-40. doi: 10.1145/176789.176792.

Shneiderman, B. (1983, August). Direct Manipulation: A Step Beyond Programming Languages. . Retrieved April 20, 2008, from http://csdl2.computer.org/persagen/DLAbsToc.jsp?resourcePath=/dl/mags/co/&toc=comp/mags/co/1983/08/mco08toc.xml&DOI=10.1109/MC.1983.1654471.

Wright, P. (1991). Cognitive overheads and prostheses: some issues in evaluating hypertexts. In *Proceedings of the third annual ACM conference on Hypertext* (pp. 1-12). San Antonio, Texas, United States: ACM. doi: 10.1145/122974.122975.