The Activity Compass

Donald J Patterson, Oren Etzioni, Henry Kautz
Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350
1 206 543 1695
{djp3, etzioni, kautz}@cs.washington.edu

ABSTRACT
In this paper, we introduce the *Activity Compass*, a cognitive aide for early-stage Alzheimer's patients. This device has a simple user interface based on the metaphor of a traditional navigation compass. By following an arrow and an icon, users who are disoriented or forgetful are assisted in reaching their destination. A server-based AI engine learns a model of routine user behavior, predicts their most likely destinations, and then directs the compass interface. By leveraging historic behavior, the interface needs no configuration; the compass automatically improves its suggestions by observing user response over time.

Keywords
Alzheimer's Disease, artificial intelligence, Assisted Cognition, behavior modeling, cognitive aid, ubiquitous computing

INTRODUCTION

Today, approximately four million Americans suffer from Alzheimer's disease. By 2050, the number is expected to rise to 15 million people [1]. In many industrialized countries, such as Japan and Canada, the percentage of people with Alzheimer's is even higher than in the US [2]. In early stages of the disease, patients experience forgetfulness and disorientation followed by an awareness of this cognitive decline and corresponding anxiety. Patients report getting lost when traveling to unfamiliar locations, misplacing items of value, decreased facility in remembering names upon introduction to new people, and an inability to perform complex tasks. The Activity Compass is a novel computer technology that is designed to address the spatio-temporal confusion that accompanies Alzheimer's disease [3].

The Activity Compass utilizes location sensing technology, handheld computers and wireless communications to provide an infrastructure from which to acquire models of patient behavior and routines. These models are utilized to automatically identify patient plans and intentions, even when these plans are incomplete or improperly formed.

Figure 1: A Prototype Activity Compass

An early prototype of the Activity Compass utilizing a PalmVII hand-held computer with a custom-built battery-pack and GPS receiver. Convergence of hand-held computers, cellular telephones and GPS technology promises to make adding external devices unnecessary.

One of the key challenges for the Activity Compass is choosing an appropriate interface method for the user population. Alzheimer's patients are frequently elderly and exhibit declining visual, auditory and manual acuity in addition to the cognitive deterioration specific to the disease. The Activity Compass attempts to overcome these limitations by using the paradigm of a navigation compass to interface with users (see Figure 1). The Activity Compass directs its users along a path toward their most likely destination. There is zero configuration as user models are learned over time, and feedback is (mainly) implied by whether or not the user is following the plan.

MOTIVATING SCENARIOS

Two examples of when the Activity Compass would be useful are as follows. In each scenario the user suffers from early-stage Alzheimer's disease, although the described functionality could be useful for healthy individuals as well.

(1) Upon exiting the library, a library patron becomes disoriented and forgets where she is headed. The Activity Compass, sensing that she is not traveling toward home (her typical routine after leaving the library), beeps in her handbag. She retrieves the Activity Compass and observes an arrow pointing toward a house icon. This reminds her that she is going home and she sets off in the direction of the arrow. After a few feet she is confident of how to get home and returns the Activity Compass to her bag. The device remains silent as long as she continues heading toward home.

(2) Consider a volunteer at a local elementary school. One day after volunteering he leaves the building and begins heading toward his car. However, due to mild confusion, he is headed toward yesterday's parking spot. The Activity Compass senses that he is not headed toward the location where he last left his car, and vibrates in his pocket. He retrieves the Activity Compass and observes the arrow pointing toward the car icon. Because he cannot recall the location of the car, he follows the arrow all the way to the vehicle before returning the Activity Compass to his pocket.

COMPUTATIONAL PROBLEM

INPUT

The information that the Activity Compass receives from the user consists of a series of Global Position System (GPS) readings that contain position and velocity information. In addition, the system has access to a historical database of several weeks of readings. The historic data is necessary for the system to learn and achieve full functionality. Before the historic data is recorded, the device functions as an ordinary GPS-enabled personal digital assistant (PDA).

OUTPUT

The Activity Compass chooses when to alert, and outputs an arrow and an icon that instructs the user on how to reach a destination.

PROCESSING

User Model

The Activity Compass maintains a basic model of a user's actions that we call *Activity Paths,* and generates them by abstracting the sensor readings from routine user behavior. For example, the Activity Compass might maintain an Activity Path that corresponds to finding your parked car after work. Generating Activity Paths requires segmentation of the input data stream into semantically coherent pieces, finding corresponding segments, and abstracting the details that they share. The server, while off-line, periodically does these computations. This processing corresponds to training the device and is only done incrementally to incorporate new knowledge into the Activity Compass.

Activity Paths capture relationships between time, user location and mode of transportation, and can be partially abstracted to capture concepts such as "home", "bus stop", or "morning." The Activity Compass only maintains Activity Paths for behaviors that it has learned are typical of its user.

User Monitoring

When active, the Activity Compass monitors which Activity Paths it believes are in progress at the current time, and what constraints are in place that might prevent Activity Path completion. Determining constraints involves integrating a calendar, real-time bus and traffic information, perceived user-preferences and knowledge about the transportation domain. Once constraints have been identified, the Activity Compass can choose a destination that satisfies the most important constraints, and direct a user toward it. For example, the Activity Compass might reason that in order to get home, a user must take a bus, and so it is currently more important to direct the user to walk to an appropriate bus stop than it is to direct him to walk toward home.

User Feedback

In order to identify which constraints are the most important, the Activity Compass incorporates *implicit* feedback from the user. Observing which Activity Paths the user follows and which suggestions the user ignores generates positive and negative feedback.

An important human computer interface (HCI) question we are studying is the most appropriate method for *negative* feedback. The least intrusive approach would simply take the user's lack of compliance with the system's suggestion as negative feedback. However, if the user's disease is such that he is sometimes unaware that the device is trying to give him advice, it may be more appropriate to require the user to explicitly indicate that he is ignoring the suggestion, for example, by tapping on the screen.

Implicit feedback allows the Activity Compass to self-regulate its training period. By silently predicting which Activity Paths are in progress and noting which are completed, the Activity Compass is able to validate its predictions without any user intervention. Only when a sufficient threshold of accuracy is crossed does the Activity Compass begin to emerge from its training phase by alerting the user of its high-confidence recommendations. The training period should be palatable to the user because the "standard" features of the hand-held device such as calendar management, cellular telephony and contact management are always functional.

IMPLEMENTATION STATUS

The Activity Compass is currently implemented with a client-server architecture. A Palm i705 hand-held computer functions as the client and a 1.5Ghz Pentium II networked computer running the Linux operating system functions as the server.

Despite the abundance of consumer grade hand-held computers, GPS receivers, cellular telephones, and wireless devices, there are very few options for devices that combine all of these technologies. The convergence of these devices, coupled with a suitable battery life, is a prerequisite for a consumer-grade Activity Compass. In the meantime, we have developed a custom-built battery pack and GPS receiver (shown in Figure 1). Together with the handheld computer the entire device is approximately the size of a large novel. As the technology emerges, future versions will incorporate more accurate location information (< 15m) and telephone capabilities.

The Activity Compass utilizes a very low bandwidth cellular telephone network connection to communicate current sensor readings to a server. The server responds with a user interface task that the client carries out. A typical task might be to direct the user to a latitude and longitude location and show the car icon. By utilizing the local GPS sensor and the computational ability of the client, the system is moderately robust to temporary communication failures with the server.

Our current training set consists of 18,328 data points taken at two and ten second intervals over a span of three months. Our current research is focused on four issues:

- Automatically determining appropriate thresholds for intervention that balance annoying the user and providing important information [4].
- Determining how to interpret various methods of negative feedback.
- Exploring the trade-offs between expressiveness and tractability in the representation of the Activity Paths.
- The application of Dynamic Bayes Nets and particle filters to data smoothing and labeling.
- Methods of abstracting Activity Paths that utilize Version Space algorithms.

RELATED WORK

There are several projects that are concerned with similar user interface issues as the Activity Compass; they include the Nursebot project and the Aware Home Research Initiative.

The Nursebot project at Carnegie Mellon University [5] aims to develop personal robots to help elderly people during their everyday lives. The Nursebot project focuses on indoor environments and does not specifically address cognitive decline, but shares many similarities to the Activity Compass by targeting an elderly user population and their associated interface concerns.

The Aware Home Research Initiative at Georgia Tech, like the Nursebot project, is aimed at developing an indoor environment that can perceive intentions and assist elderly occupants [6]. The Aware Home is not targeting a mobile robot interface, and as such most of their work has focused on the low-level sensing infrastructure [7], but they have also built several prototype applications to help occupants with general problems of aging [8].

CONCLUSIONS

The Activity Compass is an example of an intelligent user interface for Alzheimer's disease sufferers and others with mild cognitive deficits. As a result of this unique user group, the interface must be simple and self-explanatory and rely on concepts that are already well understood. Part of the challenge of working with the Alzheimer's community is to address the needs of the elderly as well as the needs of users with mild cognitive decline. The Activity Compass strives to provide a suitable interface by requiring zero configuration, and using the common metaphor of a navigation compass. Despite, or perhaps because, of its very simple user interface, we anticipate that the Activity Compass will provide significant value for users.

REFERENCES

1. Federal Interagency Forum on Aging-Related Statistics (http://www.agingstats.gov). Older Americans 2000: Key Indicators of Well-Being. Federal Interagency Forum on Aging-Related Statistics, Washington, DC: U.S. Government Printing Office. August 2000.
2. National Institute on Aging 2000. Progress Report on Alzheimer s Disease: Taking the Next Steps. Bethesda, Maryland: National Institute on Aging.
3. B. Reisberg, S. H. Ferris, J. J. Leon, and T. Crook. The global deterioration scale for the assessment of primary degenerative dementia. *American Journal of Psychiatry*, 1982.
4. E. Horvitz, A. Jacobs, and D. Hovel. Attention-sensitive alerting. In *Proceedings of UAI '99 Conference on Uncertainty and Artificial Intelligence*, pages 305–313, 1999.
5. G. Baltus, D. Fox, F. Gemperle, J. Goetz, T. Hirsh, D. Magaritis, M. Montemerlo, J. Pineau, N. Roy, J. Schulte, and S. Thrun. Towards personal service robots for the elderly. In *Proc. of the Workshop on Interactive Robotics and Entertainment (WIRE-2000)*, 2000.
6. A. K. Dey, D. Salber, and G. D. Abowd. A context-based infrastructure for smart environments. In *Proceedings of the 1st International Workshop on Managing Interactions in Smart Environments (MANSE '99*), Dublin, Ireland, Dec. 1999.
7. J. M. Sanders. Sensing the subtleties of everyday life. *Research Horizon*s, Winter 2000.
8. R. J. Orr and G. D. Abowd. The smart floor: a mechanism for natural user identification and tracking. In *Proceedings of the 2000 Conference on Human Factors in Computing Systems (CHI 2000*), The Hague, Netherlands, 2000.