

An Overview of the Assisted Cognition Project

Henry Kautz, Dieter Fox, Oren Etzioni, Gaetano Borriello, and Larry Arnstein

Dept. Computer Science & Engineering
University of Washington
Seattle, WA 98195

Abstract

The rise of Alzheimer's disease is one of the greatest health crises facing the industrialized world. Today, approximately four million Americans suffer from Alzheimer's disease; by 2050, the number is expected to rise to 15 million people. As a result of the increasing longevity of the elderly, many sufferers are now aware that their capacities to remember, to learn, and to carry out the tasks of everyday life are slowly being lost.

The Assisted Cognition Project is a new joint effort between the University of Washington's Department of Computer Science, Medical Center, and Alzheimer's Disease Research Center that is exploring the use of AI systems to support and enhance the independence and quality of life of Alzheimer's patients. The goal of the Assisted Cognition project is to develop novel computer systems that will enhance the quality of life of people suffering from Alzheimer's Disease and similar cognitive disorders. Assisted Cognition systems use ubiquitous computing and artificial intelligence technology to replace some of the memory and problem-solving abilities that have been lost by an Alzheimer's patient. Two concrete examples of the Assisted Cognition systems we are developing are an *ACTIVITY COMPASS* that helps reduce spatial disorientation both inside and outside the home, and an *ADAPTIVE PROMPTER* that helps patients carry out multi-step everyday tasks.

Introduction

In the last 50 years Alzheimer's disease has changed from a relatively rare ailment to one of the defining characteristics of industrialized societies. In 1950 at most 200,000 people in the US had Alzheimer's disease; this number increased to 500,000 by 1975, and stands at 4 million today (National Institute 2000). By 2050 the number of Alzheimer's patients in the US is expected to be 15 million, out of a world total of 80 million (Shenk 2001). The economic toll of Alzheimer's is staggering: caring for a single person with Alzheimer's disease costs more than \$47,000 a year whether a person is at home or in a nursing home (Rice, Fox, & Max 1993).

Alzheimer's disease is a specific medical condition distinct from the "normal" kinds of forgetfulness that

come with aging (Hooyman & Kiyak 2001). Reisberg's seven stage theory describes the typical progress of the disease. In stage 1, the disease has begun but is asymptomatic. In early Alzheimer's, stages 2–3, patients exhibit forgetfulness, easily become lost, have difficulty with word and name recognition, misplace objects, *etc.* In the middle stages, 4–5, the ability to perform multi-step tasks without assistance and frequent prompts is often impaired. Late Alzheimer's, stages 6–7, include loss of awareness of surroundings, loss of speech, and finally loss of basic psychomotor skills.

For centuries technology has been created to help people with physical limitations: consider eyeglasses, hearing aids, wheelchairs, and artificial limbs. We have coined the name "Assisted Cognition" for new synthesis of AI and ubiquitous computing technology designed to help people with the *cognitive* limitations associated with Alzheimer's disease and similar conditions.

With physical limitations it is obvious that an individual's level of function is determined by the total system of the individual and his or her environment. For example, a wheelchair bound person in a city without wheelchair ramps and wheelchair accessible buses has limited mobility, while that same individual in a wheelchair-friendly city can have high mobility. In judging cognitive function it is likewise important to take account of the environmental context.

For example, in a familiar, static environment, a person in early stages of Alzheimer's may show a high level of function in performing everyday tasks as dressing and housework. The familiar sights and sounds of the home provide *implicit prompts* to initiate appropriate behavior. For example, when the person wakes in the morning the sight of the light slanting in the window, the smell of coffee coming up the stairs, and the sounds of the person's spouse moving about in the kitchen below triggers the start of the person's morning routine: *e.g.*, getting out of bed and walking toward the bathroom. Place the same person in a novel environment and he or she have difficulty in performing such basic tasks (Lawton 1983).

Just as important as the physical environment is the *social context* of other people. For example, an Alzheimer's patient may come to rely on the memory, reminders, guidance, and prompts of a caregiver, such as a spouse or other family member (Reisberg *et al.* 1982;

Shenk 2001). Thus the social environment can provide *active interventions* that extend the patient's ability to handle the challenges of everyday life. But there is a limit to any person's capability to provide such help: physical and emotional 'burnout' of caregivers, often with serious health consequences, is a common phenomenon (Hooyman & Kiyak 2001).

The goal of Assisted Cognition is to develop computer systems that can provide such *active* assistance to an Alzheimer's patient. In brief, Assisted Cognition systems (i) sense aspects of an individual's location and environment, both outdoors and at home, relying on a wide range of sensors such as Global Positioning Systems (GPS), active badges, motion detectors, and other ubiquitous computing infrastructure; (ii) learn to interpret patterns of everyday behavior, and recognize signs of distress, disorientation, or confusion, using techniques from state estimation, plan recognition, and machine learning; and (iii) offer help to patients through various kinds of interventions, and alert human caregivers in case of danger.

The ultimate aim is for Assisted Cognition systems to become an integral part of a patient's physical and social contexts. After providing an overview of organization of the Assisted Cognition project at the University of Washington, we will briefly describe two specific Assisted Cognition systems we are developing, the ACTIVITY COMPASS and the ADAPTIVE PROMPTER.

The Assisted Cognition Project at UW

The Assisted Cognition Project is an interdisciplinary effort between many organizations with expertise in computer science and/or care of Alzheimer's patients. The first category includes the University of Washington's Dept. of Computer Science and Engineering, the Dept. of Electrical Engineering, and Intel Research. The second includes the Alzheimer's Disease Research Center (ADRC), the UW schools of Nursing, Health and Public Policy, and Medicine, and the UW Medical Center.

UW is a national center for research on Alzheimer's disease. The ADRC is one of 32 centers funded by the National Institute on Aging, and has a unique emphasis on the care of Alzheimer's patients in the home. The UW Medical Center runs the *Alzheimer's Disease Patient Registry*, which tracks a population of 3,500 people, both normally aging and those with Alzheimer's, and the school of Medicine hosts the *National Alzheimer Coordinating Center*.

An important partner in the Assisted Cognition Project is **EliteCare**, a private company whose work in building and running retirement homes is pushing the state of the art in ubiquitous computing applications for the elderly (ComputerWorld). Oatfield Estates is a living prototype home for 60 residents in Portland, OR. In addition to providing high-speed internet access and customized applications for its residents, the home senses and records nearly all activity going on in or around its campus. This includes the movements of all resi-

dents and staff, operation of all lights and appliances, movement of all doors, *etc.* This enormous and detailed real-world source provides the primary raw data for our project. In addition, we are also gathering other real-world data sources. For example, a major part of the ACTIVITY COMPASS system involves learning patterns of behavior from GPS data. We have built small GPS recorders from off-the-shelf hardware and are gathering initial data sets from the movements of students involved in the project.

Our rich set of collaborations allow us to take a principled approach to Assisted Cognition, wherein (i) research is based on real, not simulated data, (ii) applications are designed in collaboration with experts on care of the elderly, and (iii) developed systems are tested on real patient populations, rather than simply shown off as a one-time demo.

Examples of Assisted Cognition Systems

We begin by describing two concrete examples of Assisted Cognition systems we are developing, and then briefly describing some of the common underlying technical research.

The Activity Compass

The ACTIVITY COMPASS is an Assisted Cognition tool that helps direct a disoriented person towards their destination. The compass is based on our client/server architecture (see Figure 1) where the client handles interaction with the user, and the server stores sensor readings, constructed models, and background information. The ACTIVITY COMPASS client is a hand held device. We refer to it as a *compass* because the central user interface element is a large arrow that directs a person based on their current location and the system's assessment of their desired destination.

For example, when coupled with a Global Positioning System (GPS), the compass may detect that a patient is wandering around and appears to be lost. In this case, it may suggest that the patient head home by pointing the arrow in an appropriate direction, and displaying the graphic for "home". Inside the home, the compass can rely on indoor sensors to direct a person towards the kitchen when meal time is overdue, or towards the medicine cabinet (with appropriate graphics, again) when appropriate.

These examples may appear straight forward, but several capabilities have to be in place before the ACTIVITY COMPASS can function effectively. First, the compass needs to construct a high level model of a person's activities based on sensors readings and corresponding background information. Outdoors, for example, the sensors are simply GPS coordinates and the background information is a detailed map annotated with key locations and addresses (*e.g.*, "home", "doctor's office"). In this case, the compass needs to record the person's typical activities such as favorite walks, destinations, *etc.* and construct models of typical behavior. In addition, GPS readings can be augmented by additional sensors to com-

pute the person's orientation; outdoors, attaching a tiny "real" compass to the client can provide orientation information, and triangulation on IR signals can achieve the same effect indoors.

Next, plan recognition is essential for the compass to make sense of the stream of sensor readings about the patient current actions. Finally, the compass needs to assess whether a deviation from a normal pattern constitutes a cause for intervention or not. For example, if the patient is walking to a new location in a commercial area she may be going to a newly opened cafe; the compass should not attempt to direct her home. However, if the patient has been wandering in circles for hours, then intervention is clearly appropriate. Intermediate cases require appropriate reasoning under uncertainty.

The ACTIVITY COMPASS is designed to learn from its own experience of interacting with the patient. Which interventions are well received? Which are ignored? Is the patient's behavior changing over time? Are more interventions warranted? Frequently, the compass will be able to detect the answers to these questions automatically and update its behavior accordingly.

The Adaptive Prompter

A common problem in the early to middle stages of Alzheimer's is a difficulty in carrying out complex tasks, while the ability to perform simple actions is relatively unimpaired (Reisberg *et al.* 1982). Some complex tasks — such as driving — should certainly not be attempted by anyone with Alzheimer's, while many other activities of daily living — such as personal grooming, hobbies, and housework — should be supported as long as possible. For example, an Alzheimer's patient may be able to perform the individual steps in dressing him or herself, but be unable to recall the proper sequence of actions, *e.g.* that socks go on before shoes.

The *adaptive prompter* is a system that helps guide an impaired individual through the steps of such an activity. Input to the system comes from a sensor network embedded in the home environment. Data from the network, together with other information such as the time of day, is integrated to create a model that predicts *what the patient is trying to do*. The sensors may include ones for sound, motion, position of objects, movement of doors, operation of appliances, and so on. For example, the system might note that it is morning, and that the patient entered the bathroom and turned on the sink. Some time passes, and the patient remains motionless. The system predicts that the "morning toothbrushing and bathing" activity has begun but become stalled. Finally, the system decides to intervene by prompting the patient to pick up the toothbrush.

A prompt could be verbal or visual (*e.g.*, using a spotlight that can highlight any object in the room), or a combination. Note that a prompt is not given unless it is deemed necessary, and prompts are not pre-programmed for certain times of the time. An Assisted Cognition system must show "emotional intelligence" (Klein, Moon, & Picard 2002), because *unnecessary* prompting is likely to be more harmful than helpful.

Technical Components

Both the ACTIVITY COMPASS and ADAPTIVE PROMPTER are based on a layered architecture (see Fig. 1) linking sensor data to behaviors, behaviors to plans, and plans to potential interventions. Each layer takes in noisy and uncertain information, abstracts and fuses the data to reduce (but not always eliminate) uncertainty, and passes this higher-level view of the world to the next layer. Feedback from the effects of invention feed back down through the layers, in order to improve the accuracy and effectiveness of the underlying models. Below we briefly describe some of the ongoing research issues for each layer.

Sensor Layer

A fundamental and common subproblem to many Assisted Cognition applications is the issue of localization (Brumitt & Shafer 2001) — namely, knowing the position of objects and/or people with a certain degree of accuracy and precision. In this project, we will install and investigate a variety of sensors for location tracking, including infrared based active badges, GPS, Berkeley notes (attached to household objects), and accelerometers (Hightower & Borriello 2001). These sensors have very different capabilities, ranging from sub-centimeter distance accuracy (infrared ID tags) to purely incremental motion sensing (accelerometers). Other kinds of sensors we are employing include weight sensors (on beds and furniture) and monitors on all lights and appliances. Our Intel Research partners are now designing specialized sensors for use in a bathroom, including water and waste flow sensors and an instrumented medicine cabinet.

Data Fusion Layer

The ability to extract information from a continuous stream of data collected by a large number of sensors is one of the fundamental problems for Assisted Cognition systems. In order to deal with noisy sensor information we apply Bayes filters, which estimate posterior probability densities over the complete state of a dynamic system. In the Assisted Cognition project, the state vector can contain a variety of information such as the location of a person, the current weather condition, or the fact whether the person wears a rain coat. The *sensor model* describes the uncertainty of the sensor in connection with a model of the environment. Posterior densities over the state are updated using a version of Baye's rule. All variables can be high-dimensional vectors yielding a potentially exponential growth in complexity. Fortunately, independences between different parts of the state space can be used to reduce the complexity of the reasoning process (Koller & Lerner 2001). Figure 2 shows an example of the track of a user of an early prototype of the activity compass as he walks around the UW campus.

A major precondition for successful location tracking is the availability of a map of the environment describing the sensors and their locations. As discussed in (Brumitt & Shafer 2001), *geometric models* can significantly

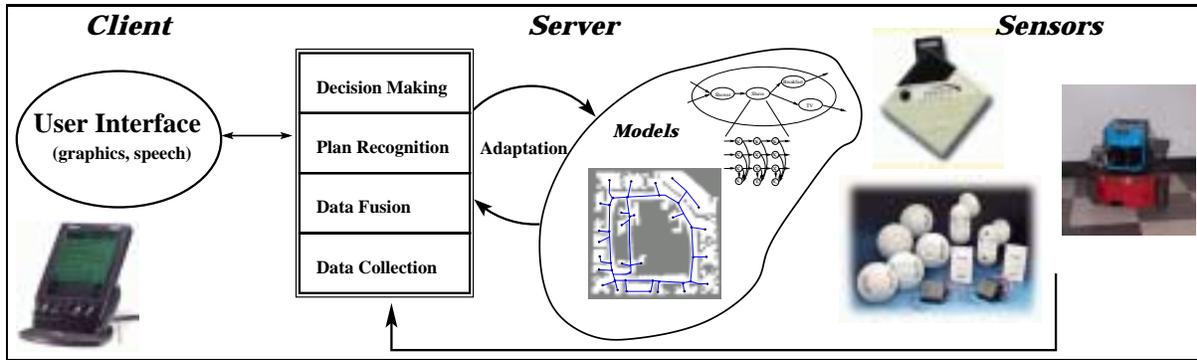


Figure 1: Architecture of Assisted Cognition systems. The server module consists of layers of increasing levels of abstraction. Probabilistic models of sensors and the environment are applied to process noisy sensor data. User activities are tracked using hierarchical Bayesian models. Probabilistic descriptions of these activities help to determine the user's plans and goals, which enable the system to choose which interactions to trigger. The models are continuously updated and refined based on information collected by the sensors.

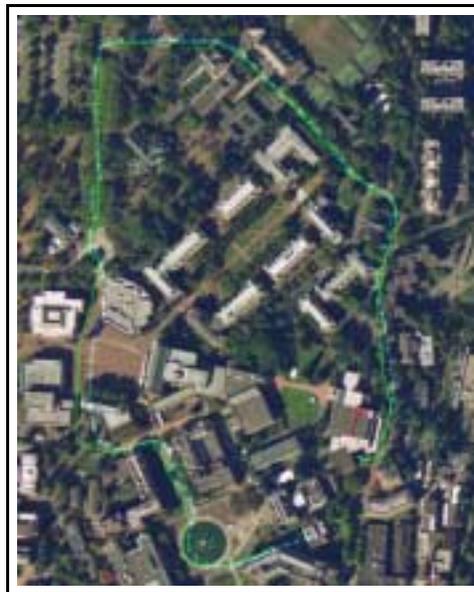


Figure 2: An prototype activity compass monitoring the movement of a user as he walks around the UW campus. The x'd line is based on the raw GPS readings. The smooth line is the actual trail of the user.

increase the performance of ubiquitous computing environments. Unfortunately, such models do not exist for most home environments, and a manual construction can be extremely difficult. We propose to build geometric maps using a *mobile robot as an active sensing device* (Thrun, Burgard, & Fox 2000). A map of a 50m \times 50m large environment and the mobile robot used to learn this map is shown in Figure 1.

Behavior Recognition Layer

The output of the data fusion layer are probability distributions over locations of people and objects at each point in time. These state sequences form the basis for the estimation of user *activities*. Over the last years, there has been tremendous progress in the development of state-based Bayesian approaches for user modeling (Horvitz *et al.* 1998) and human action recognition (Pentland 1995). These approaches connect seamlessly with sample-based Bayes filters and they are well-

suited to bridge the gap between low level state estimation and higher level behavior recognition. For example, a sequence of locations leading to the grocery store followed by a loss of GPS signal can be used to extract a high probability that the person is currently buying groceries.

The techniques discussed so far do not only apply to location tracking but extend directly to more complex scenarios like the ADAPTIVE PROMPTER. For example, the state can additionally contain information whether the faucet is running or whether the person took the toothbrush from the bathroom cabinet.

Plan and Intention Recognition Layer

Plan recognition is the task of identifying an actor's plans and goals given a partial view of that actor's behavior (Schmidt, Sridharan, & Goodson 1978; Kautz 1991). Solutions to this problem are essential, because if the computer is to provide appropriate assistance to a per-

son it needs at least a rough understanding of the goals the person is trying to achieve, and prototypical methods (*i.e.* plans) for achieving those goals. The computer needs to assess whether the plan is proceeding smoothly (no intervention is required) or, alternatively, whether the person is stymied and intervention is appropriate. A central idea in plan recognition is that plans are *hierarchical* in nature, with a lattice of more and more abstract plans connecting observable actions to high-level goals.

We are currently developing a probabilistic version of the version-space approach to plan recognition (Lesh & Etzioni 1995; Lau, Domingos, & Weld 2001), in order to handle uncertain observations and plans that may contain errors on the part of the performer. This is an open and challenging area of research.

Intervention Layer

One of the key decisions for any intelligent user interface, and for Assisted Cognition systems in particular, is under what circumstances should the system initiate an interaction with the user. If the system is overly reluctant to make suggestions, then it risks being ineffective; but if it is overly intrusive, then it is likely to be turned off by the user.

A key role for the plan recognition layer is to discriminate between the following cases: the patient is executing a plan successfully or merely pottering about and ought to be left in peace; the patient is executing a flawed plan but a specific suggestion or a gentle reminder will get her on track; the patient is disoriented or confused and needs to be re-directed; or the patient is in danger and an alert is appropriate. Making such discriminations reliably is one of the key challenges of this project. One source of information that will help is endowing our system with a modicum of “emotional intelligence” (Klein, Moon, & Picard 2002) in an attempt to track not only the patient’s actions but also her emotional state. Emotions as measured by voice timber and intonation, skin conductivity, and pulse rate can provide important cues as to whether the patient needs help, and how she is responding to specific suggestions.

We formulate the discrimination problem in a decision-theoretic framework (see (Etzioni 1991; Horvitz *et al.* 1998)) that weighs the following factors: (i) The benefit to the patient from the system’s suggestion; (ii) The probability the system’s model of the situation is accurate; and (iii) The cost of making a mistake, in terms of negative impact on both the patient’s health and perceived usefulness of the system.

We track each intervention (and each “close call” where a decision was made to not intervene) along with the factors that led to that decision, and subsequent readings of the patient’s and the environment’s state. This data forms the basis for negative and positive examples that help the Assisted Cognition system evolve its policies regarding when to intervene and when to hold back.

Related Work

There are a number of recent projects on AI systems for care of the elderly, but none (to our knowledge) targeted at Alzheimer’s patients, or based on large datasets from a potential user population. The *Aware Home* at Georgia Tech prototypes technologies to create a home that can perceive and assist its occupants (Sanders 2000). The *Nursebot* project at CMU, U. Pittsburgh, and U. Michigan (Baltus *et al.* 2000) aims at developing personal robots to help elderly people during their everyday lives. Research in *Learning Humans* project at the MIT Media Lab (Pentland 1995) has developed probabilistic algorithms for behavior recognition that are used in our work, although we are focusing on a distinct set of input modalities (GPS and motion rather than vision). Finally, general work on behavior recognition with the goal of creating intelligent user interfaces is the focus of several projects at Microsoft Research, including *Easy Living* (Brumitt & Shafer 2001) and *Attentional User Interfaces* (Horvitz & Paek 2001).

Project Status

Organization Assisted Cognition Project at the University of Washington took place in the Fall of 2001, and an initial team of faculty, students, and staff from Intel Research and EliteCare are now engaged in laying the technical groundwork: building the sensor infrastructure, collecting data, and implementing initial versions of the data fusion layer. In parallel we are working on a principled approach to the design and linkage of the behavior and plan recognition layers. In the Spring of 2002 we began an ongoing seminar on Assisted Cognition, bringing together medical and engineering faculty and students.

We believe that research on AI for cognitive aids has boundless opportunity for both practical applications and fundamental insights into perception, reasoning, and human-computer interaction.

References

- Baltus, G.; Fox, D.; Gemperle, F.; Goetz, J.; Hirsh, T.; Margaritis, D.; Montemerlo, M.; Pineau, J.; Roy, N.; Schulte, J.; and Thrun, S. 2000. Towards personal service robots for the elderly. In *Proc. of the Workshop on Interactive Robotics and Entertainment (WIRE-2000)*.
- Brumitt, B., and Shafer, S. 2001. Better living through geometry. *Personal and Ubiquitous Computing* 5(1).
- ComputerWorld. Oatfield estates elite care assisted living units. A Search for New Heroes: The 2001 ComputerWorld 21st Century Achievement Awards, online at <http://www.cwheroes.org/>.
- Etzioni, O. 1991. Embedding decision-analytic control in a learning architecture. *Artificial Intelligence* 49(1–3):129–160.
- Hightower, J., and Borriello, G. 2001. Location systems for ubiquitous computing. *Computer* 34(8).
- Hooymans, N. R., and Kiyak, H. A. 2001. *Social Gerontology: a Multidisciplinary Perspective (6th Edition)*. Allyn and Bacon.

- Horvitz, E., and Paek, T. 2001. Harnessing models of users' goals to mediate clarification dialog in spoken language systems. In *Proceedings of the Eighth International Conference on User Modeling*.
- Horvitz, E.; Breese, J.; Heckerman, D.; Hovel, D.; and Rommelse, K. 1998. The lumiere project: Bayesian user modeling for inferring the goals and needs of software users. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*.
- Kautz, H. 1991. A formal theory of plan recognition and its implementation. In Allen, J.; Kautz, H.; Pelavin, R.; and Tenenbergs, J., eds., *Reasoning About Plans*. Morgan Kaufmann Publishers. 69–126.
- Klein, J.; Moon, Y.; and Picard, R. W. 2002. This computer responds to user frustration: Theory, design, results, and implications. *Interacting with Computers – Special Issue on Cognition and Emotion* (to appear).
- Koller, D., and Lerner, U. 2001. Sampling in factored dynamic systems. In Doucet, A.; de Freitas, N.; and Gordon, N., eds., *Sequential Monte Carlo in Practice*. New York: Springer-Verlag.
- Lau, T.; Domingos, P.; and Weld, D. S. 2001. Programming by demonstration using version space algebra. *Machine Learning*. In submission.
- Lawton, M. P. 1983. Environment and other determinants of well-being in older people. *Gerontologist* 23(4):349–57.
- Lesh, N., and Etzioni, O. 1995. A sound and fast goal recognizer. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1704–1710. San Francisco, CA: Morgan Kaufmann.
2000. Progress report on alzheimer's disease 2000. online at <http://www.alzheimers.org>.
- Pentland, A. 1995. Machine understanding of human action. In *Proceedings 7th International Forum on Frontier of Telecom Technology*.
- Reisberg, B.; Ferris, S. H.; Leon, J. J.; and Crook, T. 1982. The global deterioration scale for the assessment of primary degenerative dementia. *American Journal of Psychiatry*.
- Rice, D.; Fox, P. J.; and Max, W. 1993. The economic burden of alzheimer's disease care. *Health Affairs* 12(2):164–176.
- Sanders, J. M. 2000. Sensing the subtleties of everyday life. *Research Horizons*.
- Schmidt, C. F.; Sridharan, N. S.; and Goodson, J. L. 1978. The plan recognition problem: an intersection of psychology and artificial intelligence. *Artificial Intelligence* 11:45–83.
- Shenk, D. 2001. *The Forgetting: Alzheimer's: Portrait of an Epidemic*. New York, NY: Doubleday.
- Thrun, S.; Burgard, W.; and Fox, D. 2000. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*.