

Intelligent Technology for an Aging Population

The Use of AI to Assist Elders with Cognitive Impairment

Martha E. Pollack

■ Today, approximately 10 percent of the world's population is over the age of 60; by 2050 this proportion will have more than doubled. Moreover, the greatest rate of increase is amongst the "oldest old," people aged 85 and over. While many older adults remain healthy and productive, overall this segment of the population is subject to physical and cognitive impairment at higher rates than younger people. This article surveys new technologies that incorporate artificial intelligence techniques to support older adults and help them cope with the changes of aging, in particular with cognitive decline.

We are in the midst of a profound demographic shift, moving from a world in which the majority of the population is relatively young to one in which a significant proportion of people are over the age of 65. This change poses both a challenge and an opportunity for the design of intelligent technology. While many older adults will remain healthy and productive, overall this segment of the population is subject to physical and cognitive impairment at higher rates than younger people. It is important to keep in mind that there is growth not just in the absolute number of older adults, but also in the proportion of the population that is over the age of 65; there will thus be fewer young people to help older adults cope with the challenges of aging. While human caregiving cannot and

will not be replaced, assistive technologies that can supplement human caregiving have the potential to improve the quality of life for both older adults and their caregivers. In particular, assistive technologies now being developed may enable older adults to "age in place," that is, remain living in their homes for longer periods of time. A large body of research has shown that older Americans prefer to maintain independent households as long as possible (Hareven 2001). Additionally, institutionalization has an enormous financial cost, not only for elders and their caregivers, but also for governments. In the United States, the federal government, under the auspices of Medicaid and Medicare, pays for nearly 60 percent of the nation's \$132 billion annual nursing home bill (CMS Statistics 2003), and similar expenses are incurred throughout other nations. Thus technology that can help seniors live at home longer provides a "win-win" effect, both improving quality of life and potentially saving enormous amounts of money. Other technology can help those elders who are in assisted living or skilled nursing care facilities maintain more independence there.

A range of artificial intelligence techniques has been used in the design of advanced assistive technologies. This article surveys these technologies, focusing on systems that support older adults who are grappling with cognitive decline.




Figure 1. Population Pyramid for the United States in 1950.

Source: U.S. Census Bureau

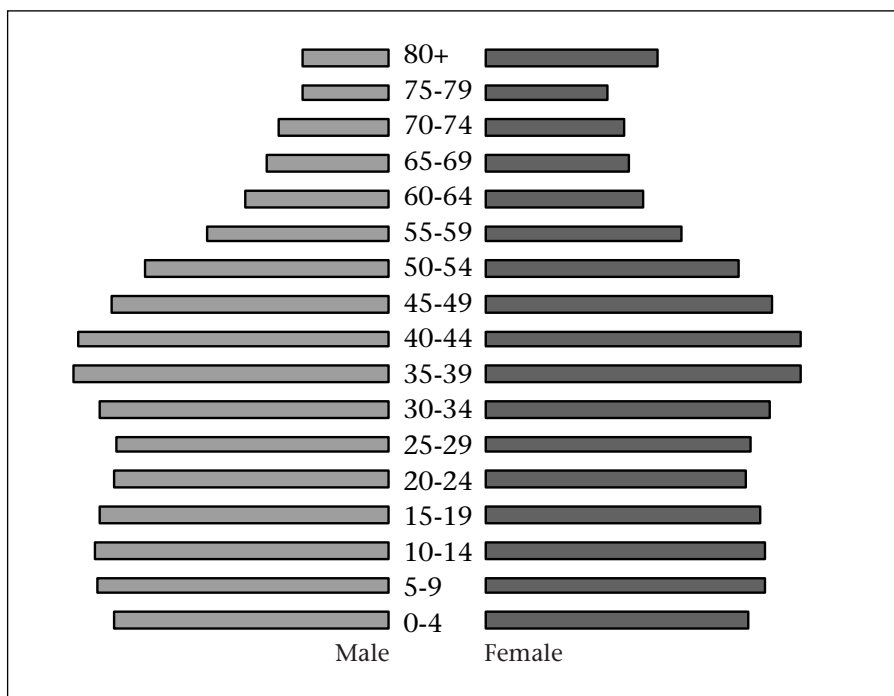


Figure 2. Population Pyramid for the United States in 2000.

Source: U.S. Census Bureau

An Aging World

To see just how dramatic the current demographic shift is, it is useful to look at population pyramids: diagrams traditionally used by demographers to visualize the

composition of a population. Population pyramids include a sequence of horizontal bars in which each bar represents the number of the people within a particular age cohort; the number of females within each cohort is shown to the right or the center line, while the number of males is shown to the left. For example, figure 1 shows the population pyramid for the United States in 1950. The source of the term *population pyramid* is clear from the figure, which has a pyramidal (or at least triangular) form, illustrating the fact that there tend to be relatively few people in the oldest age cohorts, with increasing numbers in successively younger cohorts.

As we begin the twenty-first century, population pyramids have become distinctly less pyramidal. Figure 2 shows the population pyramid for the United States in 2000, while figure 3 shows the projected pyramid for 2030. The change in demographics is immediately clear: older adults make up an increasingly greater proportion of the population. In 2000, people aged 65 and older made up 12.3 percent of the U.S. population, while by 2030, they will constitute 19.2 percent, after which growth is projected to level off so that this cohort represents 20.0 percent of the population in 2050. The most rapid growth will occur within a subgroup of this cohort—the so-called “oldest old,” or people over the age of 80. Today this group makes up 3.2 percent of the U.S. population, while by 2030 that number will increase to 5.0 percent, and by 2050, to 7.2 percent. Similar trends can be found worldwide, as illustrated in figure 4, which presents data for typical, selected countries. Although the shift is most dramatic in the more industrialized regions of the world, a significant growth in the percentage of older adults is expected in virtually every country.

A number of systems have been developed to help people compensate for the physical and sensory deficits that may accompany aging. Many of these do not rely on computer technology, for example liftchairs, which help people rise from seats, and ergonomic handles, which make it easier to open doors. However, an increasing number

of devices rely on AI and other advanced computer-based technologies. Examples include text-to-speech systems for people with low vision; a digital programmable hearing aid that incorporates a rule-based AI system to make real-time decisions among alternative signal-processing techniques based on current conditions (Flynn 2004); and a jewelrilylike device that allows people with limited mobility to control household appliances using simple hand gestures (Starner et al. 2000). In addition, significant research has been done to design obstacle-avoiding wheelchairs (see, for example, Yanco [2001] or Levine et al. [1999]).

Older adults can also be supported with technology that helps alleviate the social isolation that may stem from mobility issues or from the need to care full time for a seriously ill spouse or partner. For example, interactions between an older person and her family members and friends can be facilitated by elder-friendly email¹ and devices such as Dude's Magic Box (Siio, Rowan, and Mynatt 2002), an innovative system that allows children to interact with grandparents who are far away by placing their toys or other objects of interest into a box that projects a picture onto a touch-sensitive screen at the grandparents' home.

In addition to sensory-motor and psychosocial issues, a third area of concern for an aging population is cognitive decline. One widely cited study done in 1989 found that 10 percent of people over the age of 65 and 50 percent of those over 85 had Alzheimer's disease, probably the best-known cause of severe cognitive impairment in elders (Hebert et al. 2003). A 2001 study that was restricted to community-dwelling older adults—that is, excluding those living in nursing homes and assisted-living facilities—showed that 23.4 percent of the people over 65 exhibited cognitive impairment that was short of dementia (CIND: cognitive impairment/no dementia), while another 4.8 percent exhibited full-fledged dementia. The same study found that of those over 85, 38 percent exhibited CIND, while 27 percent exhibited dementia (Unverzagt et al. 2001). While the exact




Figure 3. Population Pyramid for the United States in 2030.

Source: U.S. Census Bureau

rates of cognitive impairment reported differ somewhat from study to study, there is no question that cognitive impairment is a serious problem for many older adults and that the prevalence of this problem increases significantly with age, the “oldest old” having a much greater rate of cognitive impairment than those in their late 60s and 70s.

This article will focus on the use of AI techniques in technology designed to support older adults with cognitive impairment. It is not intended to be a complete survey of the relevant projects but will instead provide a description of the major goals of this class of technology, along with examples that illustrate the role of AI for each class. LoPresti, Mihailidis, and Kirsch (2004) provide a comprehensive survey, including discussion of many relatively simple devices that do not include AI techniques. Haigh and Yanco (2002) and the proceedings of a National Research Council workshop (Pew and Hemel 2004) provide information about technologies that support sensory-motor and psychosocial as well as cognitive problems in elders. Sixsmith (2002) gives a brief introduction to “telecare” for cognitively impaired older people, including a description of some of the ethical issues raised by this technology. Finally, Czaja (1990) and Fisk et al. (2004) provide excellent discussions of human-factors concerns in the design of technology for seniors.

> Age 60	2000	2050	China	6.9%	22.7%
World	10.0%	21.4%	India	7.5%	20.1%
			Japan	23.3%	42.4%
Belarus	19.3%	37.6%	Myanmar	6.8%	20.5%
Germany	23.2%	34.5%			
Italy	24.1%	40.6%	Australia	16.4%	29.9%
Netherlands	18.2%	30.7%	Fiji	5.7%	22.7%
Slovenia	19.2%	41.5%			
			Egypt	6.8%	18.7%
United States	16.1%	25.5%	Iran	6.4%	24.8%
Mexico	6.9%	26.2%	Jordan	4.6%	19.0%
Brazil	7.8%	25.9%			
Colombia	6.9%	22.7%	Botswana	4.2%	6.0%
			Ethiopia	4.6%	7.7%
			Mali	3.9%	5.3%

Figure 4. Population Statistics for Typical Selected Countries Worldwide.

Source: United Nations Population Division

Goals for Assistive Technology for Cognition

Assistive technology can assist older people² with cognitive impairment in one or more of the following ways: (1) by providing assurance that the elder is safe and is performing necessary daily activities, and, if not, alerting a caregiver; (2) by helping the elder compensate for her impairment, assisting in the performance of daily activities; and (3) by assessing the elder's cognitive status.

Assurance systems aim primarily at ensuring safety and well-being and at reducing caregiver burden, by tracking an elder's behavior and providing up-to-date status reports to a caregiver. Compensation systems provide guidance to people as they carry out their daily activities, reminding them of what they need to do and how to do it. Assessment systems attempt to infer how well a person is doing—what her current cognitive level of functioning is—based on

continual observation of her performance of routine activities.

In all three cases, it is essential that the system be able to observe and reason about the elder's performance of daily activities.³ This is probably most obvious in the case of assurance systems, which must recognize whether someone has fallen, has eaten, has taken her medicine, and so on. But, as we will see later on, the ability to recognize the performance of routine activities is also essential for compensation systems, so that they can provide useful assistance that is tailored to the current needs of the user. Similarly, assessment systems work by reasoning about how and when the user performs her daily activities.

Activity monitoring is currently a very active research topic. Work has been done to use sensors to recognize the execution status of particular types of activities, such as handwashing (Mihailidis, Fernie, and Barbenel 2001), meal preparation (Barger et al. 2002), and move-

ments around town (Liao, Fox, and Kautz 2004). Additionally, several projects have attempted to do more general activity recognition, using radio frequency identification (RFID) tags attached to household objects and gloves (Philipose et al. 2004) or networks of heterogeneous sensors (Haigh et al. 2003, Glascock and Kutzik 2000). Figure 5 illustrates the range of sensor technologies that are being investigated for activity monitoring. As shown there, researchers are exploring both environmental sensors and biosensors. The former class include motion detectors and RFID readers that determine a person's location, contact switches on cabinets and refrigerator doors that indicate whether they have been opened, pressure sensors that indicate whether a person is sitting in a bed or chair, and thermometers that indicate whether a stove has been turned on. Biosensors are generally worn by a person to measure vital signs such as heart rate and body temperature. This range of sensors can be used to determine where a person is and what household objects she has used, as well as to get a general sense of her activity level. This information can be used to infer specific daily activities performed, and in turn, that knowledge, perhaps combined with biometric information, leads to a general assessment of health and well-being.

In general, Bayesian networks are the principal technology used for performing activity recognition. A typical approach is that taken in the PROACT system (Philipose et al. 2004), which employs a dynamic Bayesian network (DBN) that represents daily activities such as making tea, washing, brushing teeth, and so on. The user of PROACT wears a specially designed glove that includes an RFID reader, which can then sense household objects like cups, toothbrushes, and socks that have RFID tags affixed to them.⁴

Each activity type that PROACT is intended to recognize is modeled as a linear series of steps, and each step is then associated with the objects involved and the probability of seeing each such object. For example, making tea is modeled as a three-stage activity, in which there is high probability of using the tea kettle in the first stage (in which water is boiled), a high probability of using the box of tea bags in the second stage (in which the tea is steeped), and medium probability of using milk, sugar, or lemon in the third step (in which flavoring is added to the tea). The probabilities are intended to capture three possible sources of error: sensor noise, objects that are unknown in the model, and objects that are only optionally used for each stage. Once the model has been designed, it is converted into a DBN. PROACT

then treats information about objects used and time elapsed as observed variables and treats the current activity as a hidden variable, employing Bayes filtering techniques to derive a probability distribution over possible current activities. In preliminary testing with 14 different activity types, 14 subjects, and 108 tagged objects, PROACT demonstrated an average precision of 88 percent and average recall of 73 percent.

Assurance Systems

Examples of assurance systems include research projects (Haigh et al. 2003, Kart et al. 2002, Chan et al. 1999, Ogawa et al. 2002), a demonstration system being used in an elder-care residential setting (Stanford 2002), and a handful of commercially marketed products. The typical architecture for an assurance system is depicted in figure 6. These systems include sensors placed in the user's home, communicating via a short-range protocol such as X10 to a base station, which may in turn communicate wirelessly to a controller; the controller then uses broadband or a "plain old telephone system" (POTS) to send information to a monitoring station or directly to the caregiver. Caregivers can get status reports on a regular basis, typically by checking a web page, and are also alerted to emergencies by phone, pager, or email.

In some cases, the sensor network used by an assurance system is extremely simple, consisting, say, of just contact switches on external doorways, so that a caregiver can be immediately notified if a cognitively impaired elder leaves her home: wandering is a significant problem for people with certain types of cognitive impairment. With these systems, very little inference is actually required; instead, an alarm is simply generated whenever the contact switch is triggered.

In other cases, the network may include a wide range of sensors, which are continually monitored both to recognize deviations from normal trends that may indicate problems (for example, failure to eat meals regularly, as determined by lack of motion in the kitchen) and to detect emergencies that require immediate attention (for example, falls, as indicated by cessation of motion above a certain height). The sophistication of the inference performed using the collected sensor data varies from system to system. In some systems, only a loose connection is established between the sensor signal and the activity being reported on: machine learning may be used to infer broad patterns of sensor firings, which then serve as

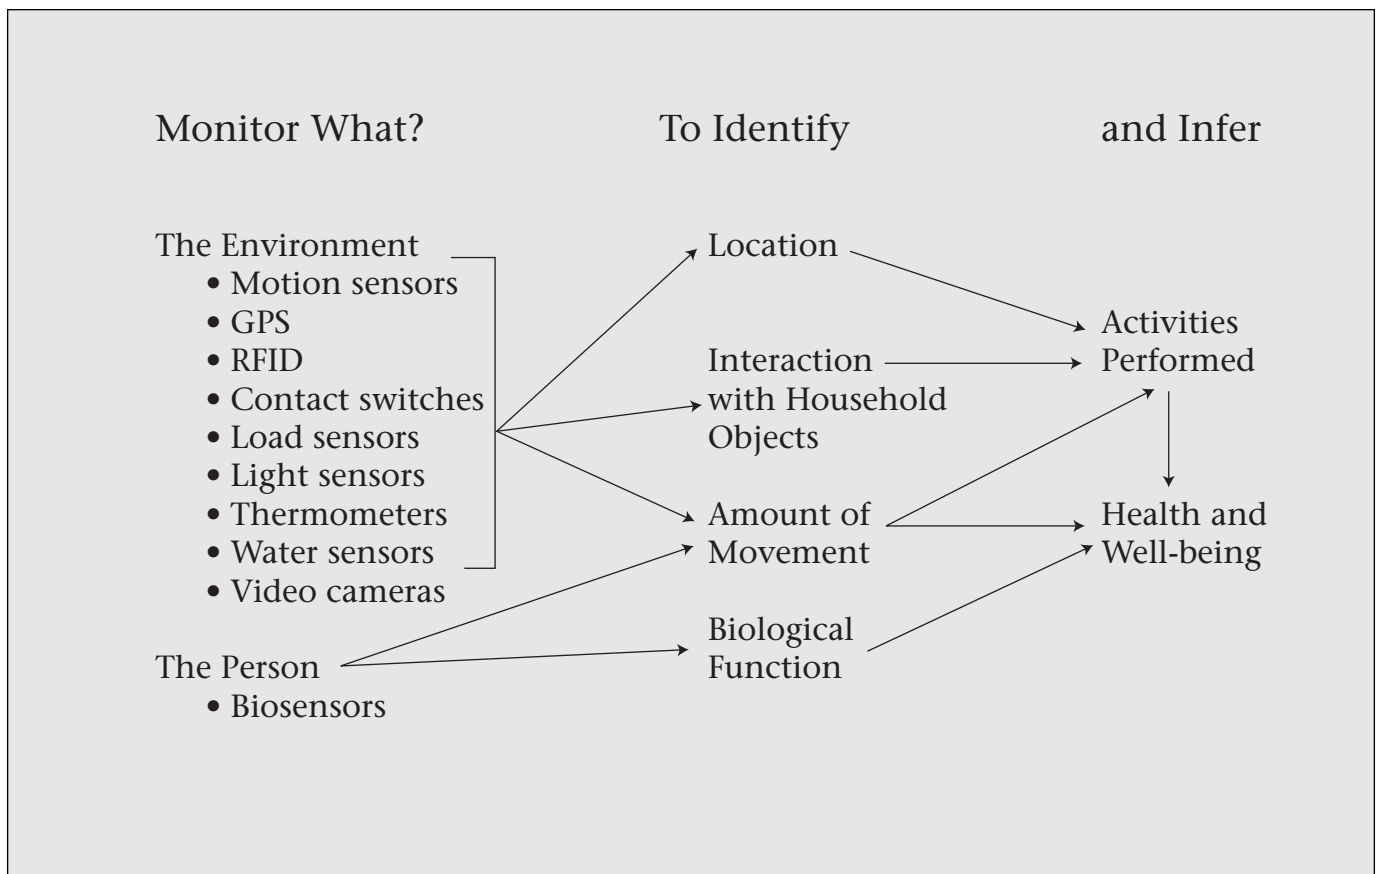


Figure 5. Sensors for Activity Monitoring.

the basis of establishing deviations from the norm that constitute grounds for issuing a warning. For example, the system may learn that a particular user is generally active in the kitchen between 7:30 and 8:15 AM, and may then indicate a potential problem if no such activity is detected one day. This contrasts with the kind of fine-grained activity recognition done by PROACT and related systems.

In addition to the question of how to determine what information to report, research has addressed the question of exactly how to present the information. One particularly interesting approach is that of the Digital Family Portrait project (Mynatt et al. 2001). The motivation behind this work is to have a continually up-to-date information display that is unobtrusive and provides a balance between the elder's privacy and the caregiver's need for information. In this system, the adult child of an elder keeps a picture of the elder in a digital picture frame, which has a pattern of 28 attractive icons (for example, butterflies or trees). The size of each icon provides a general sense of the elder's activity level on a given day; this way, the caregiver can immediately get a sense of the

overall trends in the elder's activity for the past month. Future plans are to extend the system so that when the caregiver touches one of the icons, more detailed information is provided about the elder's activities that day: whether she took her medicine, ate meals, and so on.

Compensation Systems

Where assurance systems monitor a person and provide alarms and status reports, they do not actually intervene and assist someone in accomplishing her daily activities. In contrast, a second class of systems is designed precisely to help an older adult compensate for any cognitive impairment he or she has experienced. These systems can help compensate for impairment in the ability to navigate, to manage a daily schedule, to complete a multistep task, to recognize faces, and to locate objects. Examples of systems providing the first three types of assistance are discussed below.

Navigational Support

Several systems have been developed to help older adults navigate successfully around their




Figure 6. Architecture of an Assurance System.

environments. Most of these systems aim to assist people who can no longer safely find their way around, because of either sensory difficulties such as diminished vision or mobility impairments that make walking difficult. A major goal of these systems is obstacle avoidance. Less work has been done on systems that provide navigational support to people with cognitive impairment, but one interesting example of such a system is the intelligent mobility platform (IMP) (Morris et al. 2003). IMP consists of a standard commercial walker augmented with a laser range-finder, a handheld computer providing a touchscreen interface for the user, an active drive mechanism, and intelligent navigation software. The goal of the IMP project has been to design a device that can help a potentially confused user find her way around a setting, such as a large assisted-living facility, in which she might otherwise become lost.

During an installation phase, IMP is driven through the user's environment via joystick and uses simultaneous localization and mapping (SLAM) technology (Leonard et al. 2002) to create a map based on the readings obtained by the laser range-finder. Individual regions of

the map such as the dining room, commons room, exercise room, and the user's bedroom are then hand-labeled. Subsequently, the user can select a destination from the on-board interface, and IMP will plan a path from the user's current location to that destination, guiding the user along that path by displaying a large red arrow that points the way; as the user moves along the path, the arrow shifts to show where she should turn.

Although IMP was designed primarily to provide navigational guidance, a study was also completed to determine whether the information that IMP has about a person's location throughout the day is sufficient to support activity recognition. The approach taken involves constructing a hierarchical semi-Markov model with three layers: a low-level positional layer, which represents the person's metric position (x , y , z -coordinates); an intermediate topological layer, which represents her location in terms of the mapped regions (for example, in the dining room); and a top-level activity layer, which represents the activity in which she is currently engaged (for example, eating dinner). Model parameters are learned from la-

beled data, and Bayesian filtering is then used to recognize activities from sensed location information. Preliminary results are very impressive, although more extensive testing, with a greater range of activities is required (Glover, Thrun, and Matthews 2004).

While IMP guides a person in an indoor facility, another recent system, called Opportunity Knocks, has been designed to provide outdoor navigational guidance to less severely impaired people who may still be traveling around their communities (Liao, Fox, and Kautz 2004). Opportunity Knocks, which is deployed on a cell phone enhanced with a global positioning system (GPS) chip and Bluetooth, learns its user's standard routes around town. Like IMP, it makes use of a hierarchical probabilistic model (a dynamic Bayesian network) to attempt to infer where its user is currently traveling, and, having done so, it can then detect whether the user has made an error, for example, by getting on the wrong bus. In that case, it alerts the user to her error, by making a knocking sound—hence the system's name—and provides information about how to get back to where the user is supposed to be. The user can also manually enter her destination and have the system guide her there.

Schedule Management

A second type of compensation system helps users who have suffered from memory decline that makes them prone to forgetfulness about routine daily activities. In particular, schedule-management systems remind people when to take their medicine, when to eat meals, when to take care of personal hygiene, when to check in with their adult children, and so on. Early schedule-management systems used alarm clocks, calendars, and buzzers (Harris 1978, Jones and Adams 1979, Wilson and Moffat 1994), while later systems employed wireless devices including pagers, cell phones, and palmtop computers (Hersh and Treadgold 1994, Kim et al. 2000, Wilson et al. 1997). Regardless of the delivery platform, these early systems—like most commercially available reminder systems today—function like glorified alarm clocks: they issue fixed reminders for activities at prespecified times. Unfortunately, this characteristic greatly limits their effectiveness, because older adults, like younger ones, do not live their lives according to unchanging schedules. To be useful, schedule-management systems need to be much more flexible.

The PEAT system (Levinson 1997) was the first system to use AI planning techniques to introduce flexibility into a schedule-management system. More recently, research on the

Autominder system (Pollack et al. 2003, Pollack 2002) has aimed at enhancing schedule management in several ways. Autominder provides a useful example of the value of AI technologies in assistive technology, so we describe it in some detail.

Autominder. To understand how Autominder works, it is instructive to consider an example of its interaction with a typical user: a forgetful, 80-year-old diabetic woman we'll call Mrs. Jones, who is supposed to eat a meal or snack every four hours, and who currently has an infection that requires her to take antibiotics on a full stomach. With an alarm-clock-like system, one would have to specify the exact time at which Mrs. Jones had to take her medicine. In contrast, using Autominder, Mrs. Jones or her caregiver would simply specify that she has to take her medicine within an hour of eating breakfast and dinner. Once Autominder recognizes that Mrs. Jones has eaten breakfast, it will know to remind her to take her medicine within the next hour, should she forget to do so. Similarly, rather than telling Autominder that Mrs. Jones has to eat at, say 7 AM, 11 AM, 3 PM, and 7 AM, it would instead be given the upper bound of four hours between meals or snacks. If Autominder then recognizes that Mrs. Jones has eaten lunch at 11:10, it will know to remind her to eat again at 3:10—or maybe a little earlier if her favorite TV show is on from 3:00 to 3:30.

To achieve this kind of behavior, Autominder must maintain an accurate and up-to-date model of its user's daily plan, monitor the execution of that plan, and decide about issuing reminders accordingly. As depicted in figure 7, Autominder's architecture includes three main components, one dedicated to each of these tasks. The Plan Manager stores the user's plan of daily activities in the client plan, and is responsible for updating it and identifying and resolving any potential conflicts in it. It is the component of the system responsible for answering the question "What is the user supposed to do?" The second module, the Client Modeler, uses information about the user's observed activities to track the execution of the client plan, storing beliefs about the execution status in the client model. The Client Modeler addresses the question "What is the user doing?" Finally, the third main module is a reminder-generation component called the Intelligent Reminder Generator (IRG), which reasons about any disparities between what the user is supposed to do and what she is doing, and makes decisions about whether and when to issue reminders. The Intelligent Reminder Generator answers the question "What actions

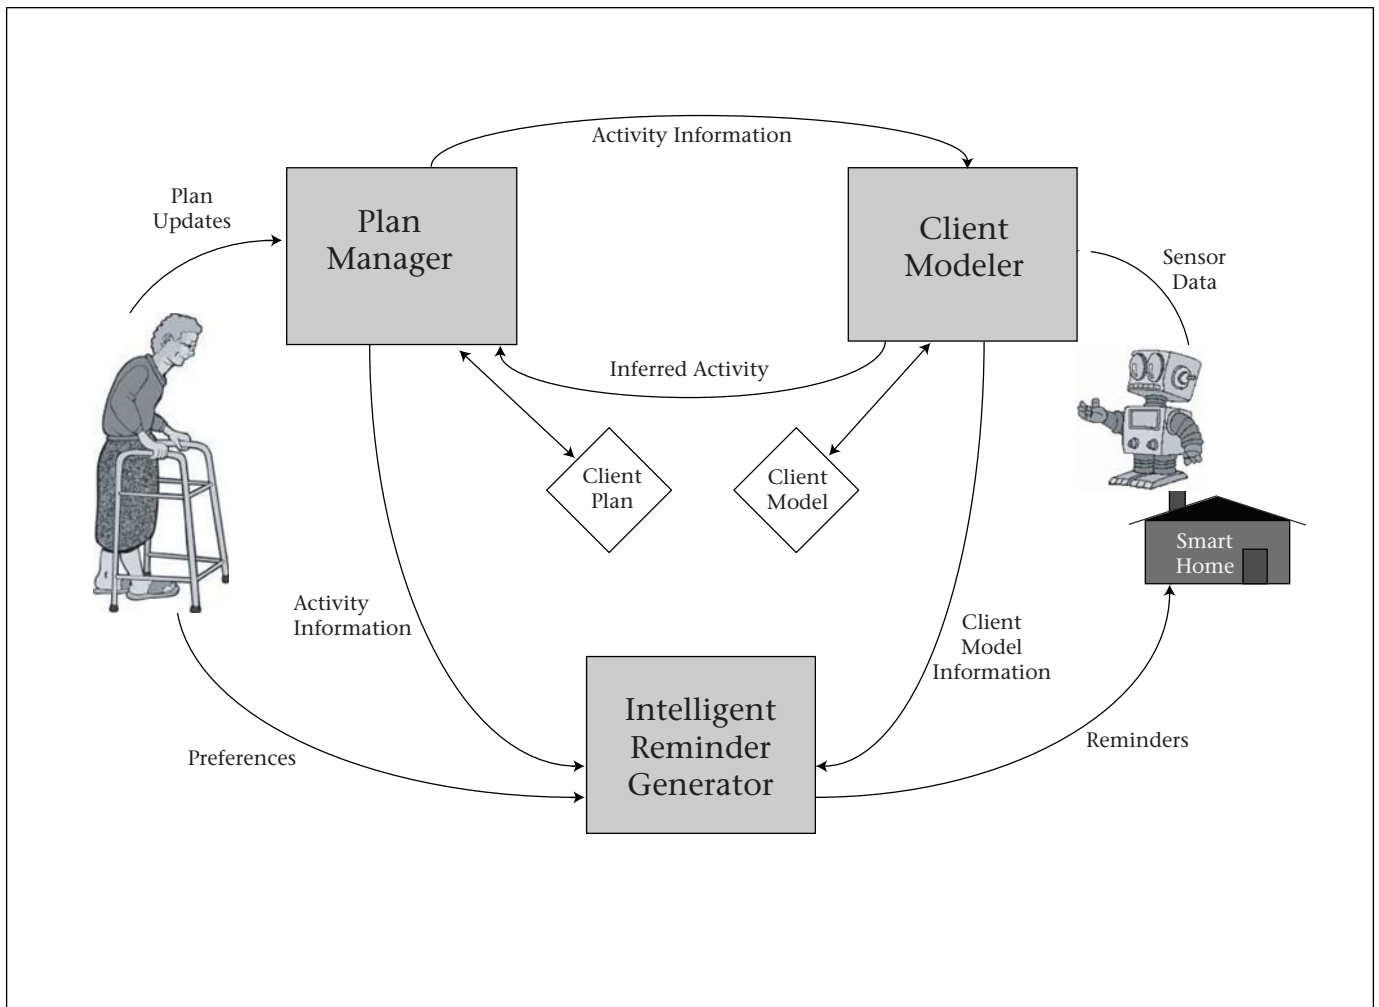


Figure 7. Autominder's Architecture.

should I (the Autominder system) take to ensure that the user successfully performs her daily activities?"

Each of Autominder's components makes heavy use of AI technology. At the core of the Plan Manager is a temporal constraint-satisfaction processing engine, which performs dynamic reasoning about the times at which activities are to be performed. (See sidebar.) Like the other systems described above for monitoring and recognizing an elder's activities, Autominder's Client Modeler relies on Bayesian inference techniques.⁵ In some regards, the most interesting component of Autominder is the Intelligent Reminder Generator. To decide whether and when to issue a reminder, the IRG must balance four goals: (1) ensuring that the user is aware of planned activities; (2) achieving a high level of user and caregiver satisfaction; (3) avoiding introducing inefficiency into the user activities; and (4) avoiding making the user overly reliant on the reminder system,

which would have the detrimental effect of decreasing, rather than increasing user independence.

It would be straightforward to generate reminders if only the first criterion were of concern: one could simply issue a reminder for every activity at its earliest possible start time, perhaps repeating the reminder at regular intervals if the activity is not performed. However, such a policy might do a potentially poor job of satisfying the other criteria. Two alternative approaches have thus been taken in Autominder to producing reminders that achieve all four goals. The first (McCarthy and Pollack 2002) adopts a local-search approach based on the planning-by-rewriting paradigm (Ambite and Knoblock 2001). It begins by creating an initial reminder plan that includes a reminder for each activity in the user plan at its earliest possible start time and then performs local search, using a set of plan rewrite rules to generate alternative candidate reminding plans. For ex-

Temporal Reasoning in Autominder's Plan Manager

To provide flexible, adaptive reminders to a user, the Autominder system must maintain an up-to-date model of the user's plan. This is done by the Plan Manager. The Plan Manager, like most automated planning systems, models plans as 4-tuples, $\langle S, O, L, B \rangle$, where S are steps in the plans, and O, L , and B are temporal ordering constraints, causal links, and binding constraints over those steps. However, we have found that temporal constraints are particularly important in providing schedule-management support, and so we augment the set of temporal constraints allowed: specifically, we make use of the language of disjunctive temporal problems (DTPs) (Stergiou and Koubarakis 2000, Tsamardinos and Pollack 2003). Formally, a DTP is a conjunctive set of disjunctive inequalities: each ordering constraint has the form

$$lb_1 \leq X_1 - Y_1 \leq ub_1 \vee \dots \vee lb_n \leq X_n - Y_n \leq ub_n$$

where the X_i and Y_i refer to the start or end points of steps in the plan, and the lower and upper bounds (lb_i and ub_i) are real numbers.

Figure A gives an example of four DTP constraints representing a typical

morning routine for an elderly user:

1. Mrs. Jones should breakfast between 7 and 8, reserving 20–30 minutes to eat.
2. She should take her medicine within an hour of eating.
3. She should bathe after 7 AM, reserving 30–40 minute for this activity.
4. She should finish breakfast and

her bath by 8:30 (perhaps because someone is coming to pick her up to drive her to the doctor).

The graph in figure A includes two nodes for each activity, one with subscript S to indicate its start, and one with E to indicate its end. The allowable interval between any two activities is indicated by $[l, u]$, denoting the lower (upper) bounds on the interval

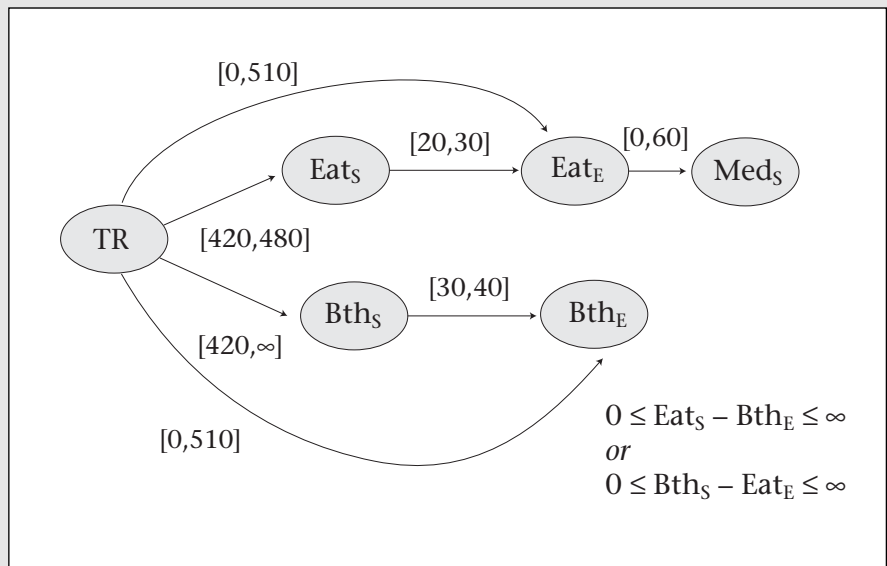


Figure A. An Example Plan as a DTP in Autominder.

ample, one rewrite rule deletes reminders for activities that have low importance and that are seldom forgotten by the user. Another rule shifts the time of a reminder for an activity to its expected time, that is, the time by which the user usually performs the activity. Yet another rule spaces out reminders for activities for the same type of action: for instance, instead of issuing eight reminders in a row to drink water, application of this rule would result in the reminders being spaced out through the day. The newly generated reminder plans are evaluated using a heuristic function that takes into account factors such as the number of reminders, their timing, and their relative spacing. The reminder plan with the highest

rank is selected, and the process iterates, with rewrite rules now being applied to the selected plan. Iteration continues until either some reminder plan is judged to have quality exceeding a prespecified threshold, or there is an interrupt indicating that there has been a change to the user plan or to the user model. In the latter case, the entire reminder plan generation process is restarted.

Although this approach has been satisfactory for the initial version of Autominder and is what is being used in the current field testing, it has three key limitations:

First, it is very difficult to handcraft good rules and evaluation functions.

Second, partly as a consequence of the first

in minutes. There is also a distinguished temporal reference point (TR), which is associated with a fixed clock time, midnight in our example; it is used to establish absolute (clock-time) constraints. In addition to the binary constraints shown in the graph, there is one additional constraint:

5. Eating breakfast and bathing cannot overlap.

This final constraint is inherently disjunctive, and is represented as

$$(0 \leq \text{Eat}_S - \text{Bath}_E \leq \infty) \\ \vee (0 \leq \text{Bath}_S - \text{Eat}_E \leq \infty)$$

This example illustrates one reason that disjunctive constraints are so important: they permit one to represent nonoverlapping activities. However, they are also useful for representing activities that simply have disjunctive temporal conditions associated with them, such as watching the TV news at either 6 or 11 PM.

Once a user plan has been specified, it is the job of Autominder's Plan Manager to update it in response to four types of events: (1) the addition of a new activity into the plan; (2) the deletion or modification of an activity already in the plan; (3) the observed execution of an activity in the plan; or (4) the passing of a critical time boundary without an activity being executed. In each case, new constraints are introduced into the client plan and propagated using a highly efficient DTP solver (Tsamardinis and Pollack 2003). As a very simple exam-

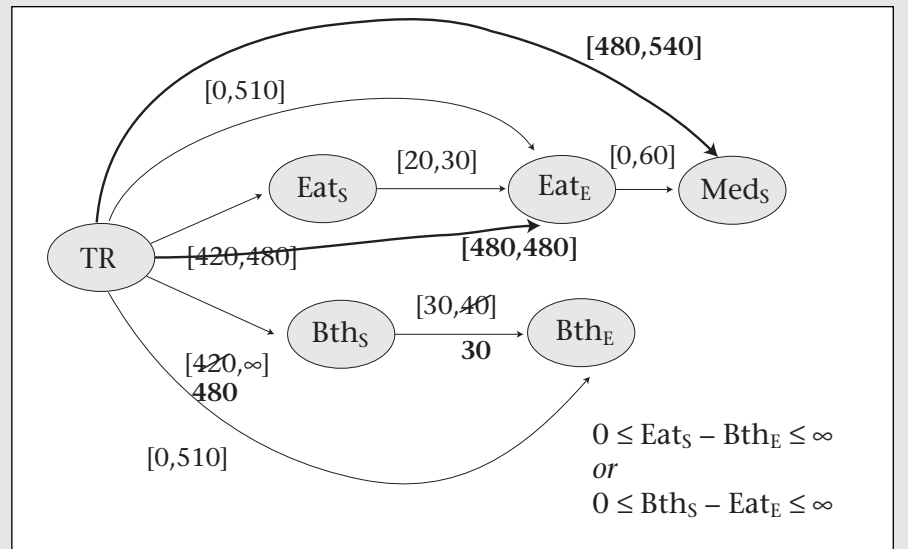


Figure B. The Plan in Figure A after Breakfast Ends.

ple, suppose that Autominder is working with the client plan in figure A and detects the end of breakfast at 8:00 AM. At this point, a new execution constraint is added (note that 480 is the number of minutes between midnight [the TR] and 8 AM.):

$$(480 \leq \text{Eat}_E - \text{TR} \leq 480)$$

Then propagation in the network representing the user's plan results in the set of constraints shown in figure B, in which, for example, the amount of time available for the user's bath and the time at which the medicine must be taken, have been updated.

The current version of Autominder's Plan Manager makes use of the

DTP representation as described above. However, a body of recent work has extended DTPs and related formalisms so that they can include preferences (Peintner and Pollack 2004; Khatib et al. 2003; Rossi, Venable, and Yorke-Smith 2004), temporal uncertainty (Morris, Muscettola, and Vidal 2001), and causal uncertainty (Tsamardinis, Vidal, and Pollack 2003). Introducing these extensions into Autominder's Plan Manager should make the system even more flexible and useful.

limitation, the rules and evaluation function that have been developed handle only a small subset of the types of interface control questions that might be addressed. To date, reasoning is only done about which activities in the client plan warrant reminders and, for those activities, when to issue the reminders. Reasoning is not done about other types of interaction with the client, for example, when to request confirmation about whether she has in fact completed some activity, nor about the best means of issuing reminders: all the reminders in the current version of Autominder are simple text strings.

Third, and perhaps most importantly, the IRG's rules and evaluation functions are fixed:

they do not change over time. Yet, typically the capabilities and needs of intended users of systems such as Autominder will change over time. As an elder's cognitive capacities diminish, she may require additional or more detailed reminders. Even over shorter time spans, a user's needs may change, for instance, during a period of illness.

The second approach being investigated in Autominder for intelligently generating reminders addresses these limitations. It uses reinforcement learning (RL) to induce an interaction policy, a function from features of the current state (for example, the time of day, the timing of the previous interaction, the user's mood, and the actions she is supposed to per-

form) to interface actions, including if and when to issue a reminder to perform a certain activity. Although this approach has so far only been tested in simulation, initial results show that it is possible to learn strategies that can personalize to an individual user and adapt to both short- and long-term changes in her needs and preferences (Rudary, Singh, and Pollock 2004).

Autominder has been deployed in prototype form on three platforms: a mobile robot, the IMP intelligent walker described above, and a handheld computer (a personal digital assistant) that communicates wirelessly with a base station. The mobile robot, which is named Pearl and is depicted in figure 8, includes a differential drive system, two on-board computers, laser range-finders, sonar sensors, microphones for speech recognition, speakers for speech synthesis, a touch-sensitive graphics display, a custom-designed actuated head unit, and stereo camera systems (Pineau et al. 2003).⁶

Preliminary field tests of Autominder have been conducted with both older adults and patients with traumatic brain injury, but systematic studies of its effectiveness have not yet been completed.

Activity-Guidance Systems

Where schedule-management systems provide a user with prompts to perform multiple activities during the course of her day, activity-guidance systems are geared towards providing reminders about consecutive steps in individual activities. As such, these systems tend to be aimed at people with more severe cognitive impairment.

An example is COACH (Mihailidis, Fernie, and Barbenel 2001), a system designed to assist a person with severe dementia who has difficulty remembering the proper sequence of everyday activities or how to use the tools that are part of those activities. The current version of COACH assists a person with handwashing; this task was chosen because studies have shown that caregivers of dementia patients find that providing assistance with bathroom activities is one of their most stressful tasks. The COACH system uses a video camera to observe the user as she attempts to wash her hands. The video image is processed to identify the two-dimensional (x, y) coordinates of a tracking bracelet worn by the subject, and a plan-recognition algorithm is then invoked to determine what step in the activity the subject is attempting (for example, turning on the water, using soap, drying hands). If the system

recognizes a problem—for instance, that the subject is using the towel before wetting her hands—a prerecorded verbal prompt is provided. In a field test with 10 subjects with moderate-to-severe dementia, COACH was shown to increase by 25 percent the number of handwashing steps that subjects could complete successfully without assistance from caregivers.

Assessment Systems

So far we have described assurance systems and compensation systems. A third use for AI technology in the care of older adults with cognitive impairment is to provide continual, naturalistic assessment of their cognitive status. Several studies have shown that cognitive impairment frequently remains undiagnosed for substantial periods of time (Boise, Neal, and Kaye 2004; Ross et al. 1997; Callahan, Hendrie, and Tierney 1995), an unfortunate situation since there exist a range of medical and behavioral means of helping cognitively impaired patients, as well as social support services for their families. Currently, most cognitive assessment is done in the clinical setting, when a person visits her physician. This means that assessment is only done infrequently and on the basis of limited information: it is quite possible that the person is having a particularly “good” or “bad” day when she happens to have a doctor’s appointment. Additionally, the very fact that the assessment is being done outside of the person’s normal living environment may bias the assessment, for instance if the person finds the doctor’s visit to be stressful.

Several researchers have thus begun to explore the possibility of using sensor-based monitoring, combined with sophisticated analysis algorithms, to assess a person’s level of functioning as she goes about her routine activities in her home. An example is Wired Independence Square, a project in which sensors are placed in a kitchen and used to collect timing data while a patient at risk for cognitive impairment performs a task such as making tea (Carter and Rosen 1999). The hypothesis, as yet unconfirmed, is that objective data such as this will be shown to correlate with assessments made with standard diagnostic batteries such as the Assessment of Motor and Process Skills (AMPS), a tool used by occupational therapists to measure the quality of a person’s performance of activities of daily living. Although the initial version of this system was installed in an in-hospital setting—a kitchen used by occupational therapists for observation and assessment of patients at the National Rehabilitation Hospital—the researchers hope to move to



Courtesy Carnegie Mellon University Robotics Institute.

Figure 8. "Pearl" the Robot.

installations within individual homes in the near future.

A related project attempts to measure cognitive performance by monitoring a person as she interacts with her home computer (Jimison, Pavel, and Pavel 2003). In initial studies, each user was monitored as she played an adapted version of the FreeCell solitaire game; this was selected both because interviews with older adults showed that they enjoyed this activity and because it is one that incorporates aspects of cognition such as short-term memory and strategic planning that are directly relevant to the performance of activities of daily living. Each time the user played FreeCell, her performance was compared to a standard established by an automated solver. Analysis was then done to track the user's relative performance over time; the goal is to identify declines in performance that may be indicative of more general cognitive decline.

In this latter project, one can already see the use of AI technology in the automated solver that determines the optimal solution to each FreeCell game. However, in both this and the preceding project, as well as in other efforts to do continual, naturalistic assessment, one can envision a larger and more central role for AI techniques: specifically, in the use of machine-learning methods to induce a person's normal level of functioning and to identify changes from that norm.

Conclusion

Interest in intelligent assistive technology for older adults is growing rapidly. Within the past five to eight years, research groups investigating the use of AI techniques for such technology have formed at more than a dozen different universities and industrial research laboratories. Workshops have been held at major AI conferences as well as at gerontology conferences (Haigh 2002; Rogers et al. 2002; Spry Foundation 2003; Gerontological Society 2003); the National Research Council sponsored a workshop resulting in a book on this topic (Pew and Hemel 2004); an online information clearinghouse has been created (Center for Aging Services Technology—CAST)⁷; and the United States Senate Special Committee has held a hearing on the assistive technology for elders (U. S. Senate 2004).

This article has provided a survey of intelligent technology to support elders with cognitive decline. Assurance systems are already available as commercial products; compensation systems mainly exist as research prototypes; and ideas about developing assessment

systems are still largely emerging. Research challenges abound, and in designing technology to support an aging population, there are compelling reasons to employ a number of different AI technologies, including plan generation and execution monitoring, reasoning under uncertainty, machine learning, natural language processing, and robotics and machine vision. Additional challenges will require collaboration with colleagues having expertise in sensor-network architectures, privacy and security, and human-machine interaction. The comedienne Lucille Ball is said to have remarked, "The secret to staying young is to live honestly, eat slowly, and lie about your age." With continued research, we can perhaps add a fourth injunction: "use AI technology."

Acknowledgements

This article was written in part with support from the National Science Foundation, Grant IIS-0085796, and from the Intel Corporation, Grant #13109. The Autominder system has been developed with the efforts of a number of past and present graduate students, including Laura Brown, Dirk Colbry, Colleen McCarthy, Michael Moffitt, Cheryl Orosz, Bart Peintner, Sailesh Ramakrishnan, Peter Schwartz, Joe Taylor, and Ioannis Tsamardinos. Thanks are also due to Jacqueline Dunbar-Jacob, Judith Matthews, and Sebastian Thrun, my fellow members of the original "Nursebot" project.

Notes

1. Generations on Line, www.generationsonline.com.
2. And others; although this paper concentrates on technology for older adults, this same technology can also be useful for younger people with cognitive impairment, including people with mental retardation and victims of traumatic brain injury.
3. The disability literature classifies certain fundamental activities such as eating, dressing, and bathing as activities of daily living (ADLs) while other, somewhat more complicated activities such as preparing meals, light housework, and paying bills are classified as instrumental activities of daily living (IADLs). Additional categories for classification are sometimes also introduced. In this article, we will simply refer to "daily" or "routine" activities to refer to the variety of things that a person must do to function autonomously in her home.
4. In the current prototype, the glove is somewhat awkward; it would need to be redesigned for an actually deployed system, but it suffices for a proof of concept.
5. In the versions of Autominder that we have been testing in the field to date, we have not employed sensors, instead relying on the user to tap a screen to indicate that she has completed an activity. This is obviously just a temporary approach; we are currently investigating sensor-based activity-recognition algorithms in the laboratory and aim to develop a ver-

sion of Autominder that fully integrates sensor networks in the near future.

6. Pearl was designed and built at Carnegie Mellon University by Sebastian Thrun and his students.

7. Center for Aging Services Technology (CAST), www.agingtech.org/.

References

- Ambite, J. L., and Knoblock, C. A. 2001. Planning by Rewriting. *Journal of Artificial Intelligence Research* 15: 207–261.
- Barger, T.; Alwan, M.; Kell, S.; Turner, B.; Wood, S.; and Naidu, A. 2002. Objective Remote Assessment of Activities of Daily Living: Analysis of Meal Preparation Patterns. Medical Automation Research Center, University of Virginia Health System, Charlottesville, VA (<http://marc.med.virginia.edu/pdfs/library/ADL.pdf>).
- Boise, L.; Neal, M. B.; and Kaye, J. 2004. Dementia Assessment in Primary Care: Results from a Study in Three Managed Care Systems. The *Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 59: M621–M626.
- Callahan, C. M.; Hendrie, H. C.; and Tierney, W. M. 1995. Documentation and Evaluation of Cognitive Impairment in Elderly Primary Care Patients. *Annals of Internal Medicine* 122(6): 422–429.
- Carter, J., and Rosen, M. 1999. Unobtrusive Sensing of Activities of Daily Living: A Preliminary Report. In *Proceedings of the First Joint Biomedical Engineering Society and Engineers in Medicine and Biology Conference*, 678. Piscataway, NJ Institute of Electrical and Electronics Engineers.
- Chan, M.; Bocquet, H.; Campo, E.; Val, T.; and Pous, J. 1999. Alarm Communication Network to Help Carers of the Elderly for Safety Purposes: A Survey of a Project. *International Journal of Rehabilitation Research* 22(2): 131–136.
- CMS Statistics. 2003. U.S. Department of Health and Human Services. Centers for Medicare and Medicaid Services 2003 CMS Statistics. Washington, D.C.: Government Printing Office (www.cms.hhs.gov/researchers/pubs/03cmsstats.pdf).
- Czaja, S. 1990. *Human Factors Research Needs for an Aging Population*. Washington, DC: National Academic Press.
- Fisk, A. D.; Rogers, W. A.; Charness, N.; Czaja, S. J.; and Sharit, J. 2004. *Designing for Older Adults: Principles and Creative Human Factors Approaches*. Boca Raton, FL: CRC Press.
- Flynn, M. 2004. Maximizing the Voice-to-Noise Ratio (VNR) via Voice Priority Processing. *Hearing Journal* (April 2004): 54–59.
- Gerontological Society. 2003. Gerontological Society of America Formal Interest Group on Technology and Aging Workshop. Washington, DC: Gerontological Society of America (www.stockton.edu/burdickd/tagupdates.html).
- Gluscock, A., and Kutzik, D. 2000. Behavioral Telemedicine: A New Approach to the Continuous Noninvasive Monitoring of Activities of Daily Living. *Telemedicine Journal* 6(1): 33–44.
- Glover, J.; Thrun, S.; and Matthews, J. T. 2004. Learning User Models of Mobility-Related Activities through Instrumented Walking Aids. In *Proceedings of the IEEE International Conference on Robotics and Automation*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Haigh, K., ed. 2002. *Automation as Caregiver: The Role of Intelligent Technology as Elder Care*. Papers from the AAAI Workshop. Technical Report WS-02-02, American Association for Artificial Intelligence, Menlo Park, CA.
- Haigh, K. Z.; Kiff, L. M.; Myers, J.; Guralnik, V.; Kirschbaum, K.; Phelps, J.; Plocher, T.; and Toms, D. 2003. *The Independent Lifestyle Assistant (ILSA): Lessons Learned*. Technical Report ACSPO3023, Honeywell Laboratories, Minneapolis, MN.
- Haigh, K. Z., and Yanco, H. 2002. Automation as Caregiver: A Survey of Issues and Technologies. In *Automation as Caregiver: The Role of Intelligent Technology in Elder Care*. Papers from the AAAI Workshop, 35–53. Technical Report WS-02-02, American Association for Artificial Intelligence, Menlo Park, CA.
- Hareven, T. K. 2001. Historical Perspectives on Aging and Family Relations. In *Handbook of Aging and the Social Sciences*, 5th ed., ed. R. Binstock and L. George, 141–159. New York: Elsevier Science.
- Harris, J. 1978. External Memory Aids. In *Practical Aspects of Memory*, ed. M. M. Gruneberg, P. N. Morris, and R. N. Sykes. London: Academic Press.
- Hebert, L. E.; Scherr, P. A.; Bienias, J. L.; Bennett, D.; and Evans, D. A. 2003. Alzheimer Disease in the U.S. Population: Prevalence Estimates Using the 2000 Census. *Archives of Neurology* 60(8): 1119–1122.
- Hersh, N. A., and Treadgold, L. 1994. Neuroprognosis: The Rehabilitation of Memory Dysfunction by Prosthetic Memory and Cueing. *NeuroRehabilitation* 4(3):187–197.
- Jimison, H. B.; Pavel, M.; and Pavel, J. 2003. Adaptive Interfaces for Home Health. Paper presented at the Second International Workshop on Ubiquitous Computing for Pervasive Healthcare. Seattle, WA, October 12 (www.healthcare.pervasive.dk/ubicomp2003/papers/).
- Jones, M., and Adams, J. 1979. Toward a Prosthetic Memory. *Bulletin of the British Psychological Society* 32(2): 165–167.
- Kart, C.; Kinney, J.; Murdoch, L.; and Ziemba, T. 2002. *Crossing the Digital Divide: Family Caregiver's Acceptance of Technology*. Technical Report, Scripps Gerontology Center, Miami University, Oxford, OH.
- Khatib, L.; Morris, P.; Morris, R.; and Venable, K. B. 2003. Tractable Pareto Optimal Optimization of Temporal Preferences. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 1289–1294. San Francisco: Morgan Kaufmann Publishers.
- Kim, H.; Burke, D.; Dowds, M.; Boone, K. A.; and Parks, G. J. 2000. Electronic Memory Aids for Outpatient Brain Injury: Follow-up Findings. *Brain Injury* 14(2):187–196.
- Leonard, J.; Tardós, J. D.; Thrun, S.; and Choset, H., eds. 2002. *IEEE International Conference on Robotics and Automation Workshop on Concurrent Mapping and Localization for Mobile Robots*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Levine, S. P.; Bell, D. A.; Jaros, L. A.; Simpson, R. C.; Koren, Y.; and Borenstein, J. 1999. The Navchair Assistive Wheelchair Navigation System. *IEEE Transactions of Rehabilitation Engineering* 7(4): 443–451.
- Levinson, R. 1997. PEAT—The Planning and Execution Assistant and Trainer. *Journal of Head Trauma Rehabilitation* 12(2)(April): 85–91.
- Liao, L.; Fox, D.; and Kautz, H. 2004. Learning and Inferring Transportation Routines. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, 348–353. Menlo Park, CA: AAAI Press.
- LoPresti, E. F.; Mihailidis, A.; and Kirsch, N. 2004. Assistive Technology for Cognitive Rehabilitation: State of the Art. *Neuropsychological Rehabilitation* 14(1/2):5–39.
- McCarthy, C. E., and Pollack, M. E. 2002. A Plan-Based Personalized Cognitive Orthotic. In *Proceedings of the Sixth International Conference on AI Planning and Scheduling*, 213–222. Menlo Park, CA: AAAI Press.
- Mihailidis, A.; Fernie, G. R.; and Barbenel, J. 2001. The Use of Artificial Intelligence in the Design of an Intelligent Cognitive Orthosis for People with Dementia. *Assistive Technology* 13(1): 23–39.
- Morris, A.; Donamukkala, R.; Kapuria, A.; Steinfeld, A.; Matthews, J.; Dunbar-Jacobs, J.; and Thrun, S. 2003. A Robotic Walker That Provides Guidance. In *Proceedings of the IEEE International Conference on Robotics and Automation*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Morris, P.; Muscettola, N.; and Vidal, T. 2001. Dynamic Control of Plans with Temporal Uncertainty. In *Proceedings of the Sev-*

- enteenth International Joint Conference on Artificial Intelligence, 494–499. San Francisco: Morgan Kaufmann Publishers.
- Mynatt, E. D.; Rowan, J.; Craighill, S.; and Jacobs, A. 2001. Digital Family Portraits: Providing Peace of Mind of Extended Family Members. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 333–340. New York: Association for Computing Machinery.
- Ogawa, M.; Suzuki, R.; Izutsu, T.; Iwaya, T.; and Togawa, T. 2002. Long Term Remote Behavioral Monitoring of Elderly by Using Sensors Installed in Ordinary Housing. In *Proceedings of the Second Annual International IEEE-EMBS Special Topic Conference on Microtechnologies in Medicine and Biology*, 322–325. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Peintner, B., and Pollack, M. E. 2004. Low-Cost Addition of Preferences to DTPs and TCSPs. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, 723–728. Menlo Park, CA: AAAI Press.
- Pew, R., and Hemel, S. V., eds. 2004. *Technology for Adaptive Aging*. Washington, DC: National Academies Press.
- Philipose, M.; Fishkin, K. P.; Perkowitz, M.; Patterson, D. J.; Hahnel, D.; Fox, D.; and Kautz, H. 2004. Inferring ADLs from Interactions with Objects. *IEEE Pervasive Computing* 3(4): 50–57.
- Pineau, J.; Montemerlo, M.; Pollack, M. E.; Roy, N.; and Thrun, S. 2003. Towards Robotic Assistants in Nursing Homes: Challenges and Results. *Robotics and Autonomous Systems* 42(3–4): 273–282.
- Pollack, M. E. 2002. Planning Technology for Intelligent Cognitive Orthotics. In *Proceedings of the Sixth International Conference on AI Planning and Scheduling*, 322–331. Menlo Park, CA: AAAI Press.
- Pollack, M. E.; Brown, L.; Colbry, D.; McCarthy, C. E.; Orosz, C.; Peintner, B.; Ramakrishnan, S.; and Tsamardinos, I. 2003. Autominder: An Intelligent Cognitive Orthotic System for People with Memory Impairment. *Robotics and Autonomous Systems* 44(3–4): 273–282.
- Rogers, W. A.; Charness, N.; Fisk, A. D.; and Czaja, S., eds. 2002. The Fourth International Conference on Gerontechnology. *Gerontechnology* 2(1) (www.gerontechnology.info/).
- Ross, G. W.; Abbott, R. D.; Petrovitch, H.; Morens, D. M.; Grandinetti, A.; Tung, K.; Tanner, C.; Masaki, K.; Blanchette, P.; Curb, J. D.; Popper, J. S.; and White, L. R. 1997. Frequency and Characteristics of Silent Dementia among Elderly Japanese-American Men: The Honolulu Study. *Journal of the American Medical Association* 277(10): 800–805.
- Rossi, F.; Venable, K. B.; and Yorke-Smith, N. 2004. Controllability of Soft Temporal Constraint Problems. In *Proceedings of the Tenth International Conference on Principles and Practice of Constraint Programming*. Springer Lecture Notes in Computer Science Series, Volume 3258, 588–603. Berlin: Springer-Verlag.
- Rudary, M.; Singh, S. B.; and Pollack, M. E. 2004. Adaptive Cognitive Orthotics: Combining Reinforcement Learning and Constraint-Based Temporal Reasoning. In *Proceedings of the Twenty-First International Conference on Machine Learning*, 719–726. New York: Association for Computing Machinery.
- Siiio, I.; Rowan, J.; and Mynatt, E. 2002. Peek-a-Drawer: Communication by Furniture. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, 582–583. New York: Association for Computing Machinery.
- Sixsmith, A. J. 2002. New Technology for the Care of Cognitively Impaired Older People. In *Geriatric Psychiatry*, 2nd ed., ed. J. R. M. Copeland, M. T. Abou-Saleh, and D. G. Blazer, 685–687. New York: Wiley.
- Spry Foundation 2003. *Computer Based Technology and Caring for Older Adults*. Technical Report. Spry Foundation National Conference, Washington, DC, October 2–3 (www.spry.org/).
- Stanford, V. 2002. Using Pervasive Computing to Deliver Elder Care. IEEE Distributed Systems Online. New York: Institute of Electrical and Electronics Engineers (dsonline.computer.org/0203/departments/bp1a_pp.htm).
- Starner, T.; Auxier, J.; Ashbrook, D.; and Gandy, M. 2000. The Gesture Pendant: A Self-Illuminating, Wearable, Infrared Computer Vision System for Home Automation Control and Medical Monitoring. In *Proceedings of the IEEE International Symposium on Wearable Computing*, 87–94. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Stergiou, K., and Koubarakis, M. 2000. Backtracking Algorithms for Disjunctions of Temporal Constraints. *Artificial Intelligence* 120(1): 81–117.
- Tsamardinos, I., and Pollack, M. E. 2003. Efficient Solution Techniques for Disjunctive Temporal Problems. *Artificial Intelligence* 151(12): 43–89.
- Tsamardinos, I.; Vidal, T.; and Pollack, M. E. 2003. CTP: A New, Constraint-based Formalism for Conditional, Temporal, Planning. *Constraints: An International Journal* 8(4):365–388.
- United Nations. 2002. Department of Economic and Social Affairs. Population Division. World Population Prospects: The 2002 Revision Population Database. New York, NY (<http://esa.un.org/unpp/>).
- Unverzagt, F. W.; Gao, G.; Baiyewu, O.; Ogunniyi, A. O.; Gureje, O.; Perkins, A.; Emsley, C. L.; Dickens, J.; Evans, R.; Musick, B.; Hall, K. S.; Hui, S. L.; and Hendrie, H. C. 2001. Prevalence of Cognitive Impairment: Data from the Indianapolis Study of Health and Aging. *Neurology* 57(9): 1655–1662.
- United States Department of Commerce. Bureau of the Census. Population Division, International Programs Center (IPC) International Data Base (IDB). ([/www.census.gov/ipc/www/idbpyr.html](http://www.census.gov/ipc/www/idbpyr.html)).
- United States Senate, Special Committee on Aging 2004. Assistive Technologies and Independent Aging: Opportunities and Challenges. A Hearing of the United States Senate Special Committee on Aging, Serial No. 108-33. Washington, DC: U.S. Government Printing Office.
- Wilson, B. A., and Moffat, M., eds. 1994. *Clinical Management of Memory Problems*. Rockville, MD: Aspen Systems Corporation.
- Wilson, B. A.; Evans, J. J.; Emslie, H.; and Malinek, V. 1997. Evaluation of Neupage: A New Memory Aid. *Journal of Neurology, Neurosurgery and Psychiatry* 63(1): 113–115.
- Yanco, H. A. 2001. Development and testing of a robotic wheelchair system for outdoor navigation. In *Proceedings of the 2001 Conference of the Rehabilitation Engineering and Assistive Technology Society of North America*.



Martha E. Pollack is a professor in the Department of Electrical Engineering and Computer Science at the University of Michigan, where she is also associate chair for computer science and engineering. Pollack was

previously professor of computer science and professor and director of the Intelligent Systems Program at the University of Pittsburgh, and, before that, she was a senior computer scientist at the Artificial Intelligence Center at SRI International in Menlo Park, California. She received her B.A. degree from Dartmouth College and her Ph.D. from the University of Pennsylvania. A Fellow of the AAAI, Pollack recently completed a term as editor-in-chief of the *Journal of Artificial Intelligence Research*. She has done research and published papers in the areas of automated planning and execution monitoring, temporal reasoning, and natural-language processing, as well as on assistive technology for cognitively impaired people. In April of 2004, she testified before the United States Senate Special Committee on Aging, discussing the potential value of assistive technology in an aging world.